

Binary Hypothesis Testing Game With Training Data

Mauro Barni, *Fellow, IEEE*, and Benedetta Tondi, *Student Member, IEEE*

Abstract—We introduce a game-theoretic framework to study the hypothesis testing problem in the presence of an adversary aiming to prevent a correct decision. Specifically, this paper considers a scenario in which an analyst has to accept or reject the null hypothesis H_0 characterized by a probability mass function (pmf) P_X based on the evidence provided by a test sequence. In turn, the goal of the adversary is to take a sequence generated according to a different pmf and modify it in such a way to induce a decision error. P_X is known only through one or more training sequences. We derive the asymptotic equilibrium of the game under the assumption that the analyst relies only on first order statistics of the test and training sequences, and compute the asymptotic payoff of the game when the length of the sequences tends to infinity. We introduce the concept of indistinguishability region, defined as the set of pmfs that can not be distinguished reliably from P_X in the presence of attacks. Two different scenarios are considered: in the first one the analyst and the adversary share the same training sequence, in the second scenario, they rely on independent sequences. The obtained results are compared with a version of the game in which the pmf P_X is perfectly known to both the analyst and the adversary.

Index Terms—Hypothesis testing, adversarial signal processing, cybersecurity, game theory, source identification, multimedia forensics, counter-forensics.

I. INTRODUCTION

HYPOTHESIS testing is a widely studied topic with applications in virtually all technological and scientific fields. In its most basic form, an analyst is asked to decide which among two hypotheses, usually referred to as null hypothesis (or H_0) and alternative hypothesis (H_1), is true based on a set of observables, say $x^n = (x_1, x_2 \dots x_n)$. Several versions of the problem are obtained according to the knowledge that the analyst has on the probability distribution of the observables when one of the two hypotheses holds. In some cases, the probability mass function (pmf)¹ conditioned to the two hypotheses is known, in other cases only the pmf under H_0 is known, in yet other cases only a number of sample observables (hereafter referred to as training sequences) obtained under H_0 and H_1 are available.

Manuscript received April 8, 2013; revised February 4, 2014; accepted April 30, 2014. Date of publication June 2, 2014, date of current version July 10, 2014. This work was supported in part by the REWIND Project (Future and Emerging Technologies Program within the 7FP of the European Commission) under Grant 268478.

The authors are with the Department of Information Engineering, University of Siena, Siena 53100, Italy (e-mail: barni@dii.unisi.it; benedettatondi@gmail.com).

Communicated by V. Borkar, Associate Editor for Communication Networks.

Digital Object Identifier 10.1109/TIT.2014.2325571

¹Since in the rest of the paper we will assume that the elements of x^n belong to a finite alphabet, we prefer to use the term probability mass function instead of probability density function, even if at this stage we do not need to restrict our attention to the discrete case.

Due to its importance, hypothesis testing has been extensively studied and a solid theoretical framework has been built permitting to analyze and understand its many facets. In the last years, though, many applications have emerged in which hypothesis testing is given a new twist, due to the presence of an adversary aiming at making the test fail [1]. In multimedia forensics [2], for instance, a forensic analyst may be asked to decide whether an image has been acquired by a given camera, notwithstanding the presence of an adversary aiming at deleting the acquisition traces left by the camera. In the same way, the analyst may be asked to decide whether a signal has undergone a certain processing or not, by taking into account the possibility that someone deliberately tried to delete the traces left during the processing phase. Another popular example comes from spam filtering [3], wherein an anti-spam filter is presented with a test e-mail and must decide whether the e-mail contains a genuine or a spam message. It is obvious that such a test can not neglect the presence of an adversary trying to shape the message in such a way to fool the filter. Biometric authentication provides a further example. In this case the authentication system must decide whether a biometric template belongs to a certain individual, despite the opposite efforts of an attacker aiming at building a fake template that passes the authentication test [4], [5]. Other examples include: watermarking, where the detector is asked to decide whether a document contains a given watermark or not [6], steganalysis, in which the steganalyzer has to distinguish between cover and stego images [7], network intrusion detection, wherein anomalous traffic conditions must be distinguished from normal ones [8], reputation systems [9], for which it is essential to distinguish between genuine and malevolent scores, anomaly detection, cognitive radio [10], [11], and many others. In all these fields, the system designer has to take into account the presence of one or more adversaries explicitly aiming at system failure.

In the framework depicted above, the goal of this paper is to move a first step towards the construction of a general theoretical framework to analyze the binary hypothesis testing problem by taking into account the presence of an adversary aiming at impeding a correct decision. More specifically, we introduce and analyze an adversarial version of the Neyman-Pearson setup in which an analyst and an adversary, hereafter referred to as the Defender (D) and the Attacker (A), face each other in a rigorously defined context. Given a null hypothesis H_0 and a test sequence x^n , D must decide whether to accept hypothesis H_0 characterized by the pmf P_X , or not. As in the classical Neyman-Pearson scenario, the defender must ensure that the type-I error probability (i.e., the probability of rejecting H_0 when H_0 holds) is lower than a certain value. In turn, the attacker takes a sequence y^n generated under an




Fig. 1. The adversarial hypothesis testing set up considered in this paper.

alternative hypothesis H_1 , characterized by a different pmf P_Y , and transforms it into a modified sequence z^n so that when presented with the modified sequence D still accepts H_0 . In other words, A aims at maximizing the type-II error probability (i.e. the probability that the defender accepts H_0 when H_1 holds), while D 's goal is to minimize it by taking into account the presence of the attacker. A schematic representation of the above framework is depicted in Fig. 1. Even if the above scenario fits many real-world applications, throughout the paper we will use multimedia forensics as a leading example. Specifically, we will often refer to the case of a forensic analyst (playing the role of the defender), who wants to decide whether a given signal, say a video or an image, has been produced by an authorized acquisition device, for instance a video surveillance camera, or not. The analyst wants to do that in presence of a forger (playing the role of the attacker) aiming at taking a sequence drawn from a non authorized source and modify it in such a way that it seems to have been produced by the authorized device.

In order to analyze the framework reported in Fig. 1, we adopt a game-theoretic approach, in which the defender and the attacker have opposite goals and operate by satisfying a different set of requirements, all together specifying the nature of the game. The final goal will be the derivation of the optimum strategies for the defender and the attacker in terms of game equilibrium points, and the study of the achievable performance at the equilibrium.

A. Prior Work

The use of game theory to model the impact that the presence of an adversary has on (binary) hypothesis testing is not an absolute novelty. In many security oriented fields in which hypothesis testing plays a central role, game theory has been advocated to avoid entering a cat and mouse loop in which researchers alternatively play the role of the defender and the attacker, and continuously develop new countermeasures, each time by attacking a specific algorithm or strategy.

By referring to multimedia forensics, in [12] Böhme and Kirchner cast the forensic problem in a hypothesis testing framework. Several versions of the problem are defined according to the particular hypothesis being tested, including distinction between natural and computer generated images, manipulation detection, source identification. Counter-forensics is then defined as a way to degrade the performance

of the hypothesis test envisaged by the analyst. By relying on arguments similar to those used in steganography and steganalysis [13], Böhme and Kirchner argue that the divergence between the probability density functions of the observed signals after the application of the counter-forensic attack is a proper way to measure the effectiveness of the attack. Noticeably, such measure does not depend on the particular investigation technique adopted by the analyst. Even if Böhme and Kirchner do not explicitly use a game-theoretic formulation, their attempt to decouple the counter-attack from a specific forensic strategy can be seen as a first - implicit - step towards the definition of the equilibrium point of a general multimedia forensics game. Another work loosely related to the present paper is [14], where the authors introduce a game-theoretic framework to evaluate the effectiveness of a given attacking strategy and derive the optimum countermeasures. As opposed to our analysis, in [14] the attacker's strategy is fixed and the game-theoretic framework is used only to determine the optimum parameters of the forensic analysis and the attack, thus failing to provide a complete characterization of the game between the attacker and the analyst.

Game theory and information theory have been used in watermarking to model the interplay between the watermarker and the attacker. In [15]–[17], the game is played between the watermark embedder/decoder and an attacker who attempts to degrade the embedded message by modifying the watermarked signal, e.g. by adding some noise. The payoff of the game is usually the capacity of the watermark channel. A problem that is closer to the one addressed in this paper is the one considered in [18], where the jointly optimum embedding and detection strategies for a detector with limited resources are derived. Indeed, the approach used in the present paper is reminiscent of the analysis carried out in [18], given that in both cases the analysis focuses on an asymptotic version of the problem in which the resources available to the defender are limited. As opposed to the present work, however, the analysis in [18] is carried out under the assumption that no attack is present or that the attack channel is fixed, and the resort to a min-max optimization is due only to the necessity of finding the jointly optimum watermark embedding and detection strategies.

The work that is most closely related to the present paper is [19], where the source identification game with known statistics is introduced and the corresponding asymptotic Nash equilibrium derived. As a matter of fact, even if [19] restricts the analysis to multimedia forensics, the framework adopted to model the game between the forensic analyst and the adversary is very general and can also be used to model a binary hypothesis testing problem in which the statistics of the observables under the null hypothesis are perfectly known to the defender and the attacker. In contrast to [19], here we focus on a different version of the game in which the statistics of the observables under H_0 are known only through training data. This is a very important step that permits to get closer to real applications wherein an accurate statistical model of H_0 is rarely available, and the analyst can rely only on the knowledge of some sample outputs produced when H_0 holds. Also from a technical point of view, knowing P_X only

through training data represents a major deviation from [19], since it requires a thorough reformulation of the problem and a new derivation of the equilibrium point. In order to make the current paper self-contained and allow the reader to better appreciate the difference between the results obtained in [19] and the new findings provided here, we summarize the main definitions and results of [19] in Section III.

B. Contribution

With the above ideas in mind, in this paper we address the following problem, hereafter referred to as the binary hypothesis testing problem with training data (HT_{tr}). Let $X \sim P_X$ be a discrete memoryless source (DMS) ruling the emission of observables under the null hypothesis H_0 , and let x^n be a test sequence, i.e., a sequence of observables. The goal of the defender is to accept or reject hypothesis H_0 , that is, to decide whether x^n was drawn from X or not.² In doing so, D must ensure that the type-I error probability does not exceed a predefined value. Let then $Y \sim P_Y$ be a second discrete memoryless source and let $y^n = (y_1, y_2 \dots y_n)$ be a sequence generated by Y . It is the aim of A to transform y^n into a new sequence z^n in such a way that when presented with z^n the defender accepts H_0 . We also impose that A satisfies a distortion constraint requiring that the *distance* between y^n and z^n is below a certain threshold. As opposed to [19], D and A do not know the exact statistics of P_X and P_Y , since all they know is a training sequence drawn from X .

Given the above scenario, the goal of this paper is to propose a rigorous game-theoretic framework to cast the HT_{tr} problem in, and derive the asymptotic equilibrium point of the game under a simplifying hypothesis about the kind of analysis which D can carry out. We will do so for two different versions of the game stemming from different assumptions about the relationship between the training sequence available to the defender and that available to the attacker. In a first case, we will assume that A and D share the same training sequence, while in the second part of the paper we will assume that the two sequences are generated independently from each other. The main results proven in the paper can be summarized as follows:

- 1) We show that under the *limited resources* assumptions [18], the HT_{tr} game admits an asymptotic equilibrium. We also prove that the asymptotic equilibrium is the only rationalizable equilibrium of the game [20], [21]. Such an equilibrium is much stronger than the usual notion of Nash equilibrium, since the strategies corresponding to such an equilibrium are the only ones that two rational players may adopt (Theorem 3, Section V);
- 2) We compute the payoff at the equilibrium for the defender and the attacker, and introduce the notion of indistinguishability region, defined as the region with the P_Y 's that can not be distinguished reliably (i.e. with a vanishing type II error probability) from P_X (Theorem 4, Section VI);

- 3) We compare the achievable payoff of the HT_{tr} game with the results obtained in [19], where D and A have a perfect knowledge of P_X , showing that the HT_{tr} game is more favorable to the attacker with respect to the situation analyzed in [19] (Theorem 5, Section VI);
- 4) We show that the indistinguishability region is the same when D and A share the same training sequence and when they rely on independent sequences (Theorem 6, Section VII).

With regard to 1), the asymptotic equilibrium point when D and A rely on the same training sequence was already derived in [22], without realizing that the equilibrium is indeed stronger than a Nash equilibrium. The case of independent training sequences has never been studied before. The methodology used to derive the equilibrium point goes along the same lines used in [18] to derive the jointly optimum watermark embedding and detection strategy. Finally, the optimum strategy of the defender can be paralleled to the results obtained in [23], even if in a completely different context. As to point 2), our results are closely related to Sanov's theorem [24], [25], however, to the best of our knowledge, their derivation in a game-theoretic context is not trivial and represents an original contribution of the present paper.

The rest of this work is organized as follows. In Section II we introduce the notation and definitions used throughout the paper. In Section III, we recall the main results proved in [19] by casting them into a hypothesis testing framework. Such results represent the baseline against which we will compare the new results obtained in this paper. In Section IV, we formally introduce the HT_{tr} game and lay the basis for the analysis carried out in the subsequent sections. Section V is devoted to the derivation of the equilibrium point of the HT_{tr} game when A and D share the same training sequence. The payoff at the equilibrium is analyzed in Section VI, where we also introduce the notion of indistinguishability region. The case of independent training sequences is analyzed in Section VII. The paper ends in Section VIII, with some conclusions and hints for future research. Some of the most technical proofs are given in the the appendix, so as to avoid interrupting the main flow of ideas.

II. NOTATION AND DEFINITIONS

In the rest of this work we will use capital letters to indicate discrete memoryless sources (e.g. X). Sequences of length n drawn from a source will be indicated with the corresponding lowercase letters (e.g. x^n). In the same way, we will indicate with x_i , $i = 1, \dots, n$ the i -th element of a sequence x^n . The alphabet of an information source will be indicated by the corresponding calligraphic capital letter (e.g. \mathcal{X}). The pmf of a discrete memoryless source X will be denoted by P_X . The same notation will be used to indicate the probability measure ruling the emission of sequences from a source X , so we will use the expressions $P_X(a)$ and $P_X(x^n)$ to indicate, respectively, the probability of symbol $a \in \mathcal{X}$ and the probability that the source X emits the sequence x^n , the exact meaning of P_X being always clearly recoverable from the context wherein it is used. Given an event A (be it a subset

²In order to keep the notation as light as possible, we use the symbol x^n to indicate the test sequence even if, in principle, it is not known whether x^n originated from X or not.

of \mathcal{X} or \mathcal{X}^n), we will use the notation $P_X(A)$ to indicate the probability of the event A under the probability measure P_X .

Our analysis relies extensively on the concepts of type and type class defined as follows (see [24] and [26] for more details). Let x^n be a sequence with elements belonging to an alphabet \mathcal{X} . The type P_{x^n} of x^n is the empirical pmf induced by the sequence x^n , i.e. $\forall a \in \mathcal{X}, P_{x^n}(a) = \frac{1}{n} \sum_{i=1}^n \delta(x_i, a)$, where $\delta(x_i, a) = 1$ if $x_i = a$ and zero otherwise. In the following we indicate with \mathcal{P}_n the set of types with denominator n , i.e. the set of types induced by sequences of length n . Given $P \in \mathcal{P}_n$, we indicate with $T(P)$ the type class of P , i.e. the set of all the sequences in \mathcal{X}^n having type P .

The Kullback-Leibler (KL) divergence between two distributions P and Q on the same finite alphabet \mathcal{X} is defined as:

$$\mathcal{D}(P||Q) = \sum_{a \in \mathcal{X}} P(a) \log \frac{P(a)}{Q(a)}, \quad (1)$$

where, as usual, $0 \log 0 = 0$ and $p \log p/0 = \infty$ if $p > 0$.

A. Hypothesis Testing Framework

Given a sequence $x^n \in \mathcal{X}^n$, as a result of the test, \mathcal{X}^n is partitioned into two complementary regions Λ and $\bar{\Lambda}$, such that for $x^n \in \Lambda$ the defender decides in favor of H_0 , while for $x^n \in \bar{\Lambda}$ H_1 is preferred. We say that a Type-I error occurs if H_1 is chosen when H_0 holds. In the same way, we say that a Type-II error occurs when H_1 holds but H_0 is chosen. In the following, we will refer to Type-I errors as false positive errors (or false alarms) and to Type-II as false negative (or missed detection), and will indicate the probability of such events as P_{FP} and P_{FN} respectively. The motivation for such a terminology comes from applications in which H_0 is seen as a standard situation and its rejection in favor of H_1 raises an alarm since something unusual happened. This is the case, for instance, of multimedia forensics applications, where H_0 corresponds to the hypothesis that x^n was produced by a legitimate source, and an alarm is raised whenever this is not the case. It goes without saying that our derivation remains valid even in different scenarios where the false positive and false negative terms may not be appropriate. In our analysis we are mainly interested in the asymptotic behavior of P_{FP} and P_{FN} when n tends to infinity. In particular we define the false positive (λ) and false negative (ε) error exponents as follows:

$$\lambda = - \limsup_{n \rightarrow \infty} \frac{\log P_{FP}}{n}; \quad \varepsilon = - \limsup_{n \rightarrow \infty} \frac{\log P_{FN}}{n}, \quad (2)$$

where the log's are taken in base 2. Note that when the classical limit exists the above definitions can be simplified by avoiding the use of \limsup . Given that this is the case in all our derivations, for sake of simplicity, throughout the paper we will use \lim instead of \limsup .

B. Game Theory

As we said, the goal of this paper is to model the binary hypothesis testing problem in an adversarial setting as a 2-player game. More formally, a 2-player game is defined as a 4-uple $G(\mathcal{S}_1, \mathcal{S}_2, u_1, u_2)$, where $\mathcal{S}_1 = \{s_{1,1} \dots s_{1,n_1}\}$ and

$\mathcal{S}_2 = \{s_{2,1} \dots s_{2,n_2}\}$ are the sets of strategies (actions) the first and the second player can choose from, and $u_l(s_{1,i}, s_{2,j})$, $l = 1, 2$, is the payoff of the game for player l , when the first player chooses the strategy $s_{1,i}$ and the second chooses $s_{2,j}$. A pair of strategies $(s_{1,i}, s_{2,j})$ is called a profile. In a zero-sum competitive game, the two payoff functions are strictly related to each other since for any profile we have $u_1(s_{1,i}, s_{2,j}) + u_2(s_{1,i}, s_{2,j}) = 0$. In other words, the win of a player is equal to the loss of the other. In the particular case of a zero-sum game, then, only one payoff function needs to be defined. Without loss of generality we can specify the payoff of the first player (generally indicated by u), with the understanding that the payoff of the second player u_2 is equal to $-u$. In the most common formulation, the sets $\mathcal{S}_1, \mathcal{S}_2$ and the payoff functions are assumed to be known to both players. In addition, it is assumed that the players choose their strategies before starting the game so that they have no hints about the strategy actually chosen by the other player (strategic game).

A common goal in game theory is to determine the existence of equilibrium points, i.e. profiles that, in *some sense* represent a *satisfactory* choice for both players. While there are many definitions of equilibrium, the most famous and commonly adopted is the one due by Nash [27], [28]. For the particular case of a 2-player game, a profile (s_{1,i^*}, s_{2,j^*}) is a Nash equilibrium if:

$$\begin{aligned} u_1(s_{1,i^*}, s_{2,j^*}) &\geq u_1(s_{1,i}, s_{2,j^*}) \quad \forall s_{1,i} \in \mathcal{S}_1 \\ u_2(s_{1,i^*}, s_{2,j^*}) &\geq u_2(s_{1,i^*}, s_{2,j}) \quad \forall s_{2,j} \in \mathcal{S}_2, \end{aligned} \quad (3)$$

where for a zero-sum game $u_2 = -u_1$. In practice, a profile is a Nash equilibrium if each player does not have any interest in changing his choice assuming the other does not change its strategy.

Despite its popularity, the practical meaning of Nash equilibrium is difficult to grasp, since there is no guarantee that the players will end up playing at the equilibrium. This is particularly evident when more than one Nash equilibrium exists. A definition of game equilibrium with a more practical meaning can be obtained by relying on the notion of dominant and dominated strategies. A strategy is said to be strictly dominant for one player if it is the best strategy for the player, no matter how the other player may play. In a similar way, we say that a strategy $s_{l,i}$ is strictly dominated by strategy $s_{l,j}$, if the payoff achieved by player l choosing $s_{l,i}$ is always lower than that obtained by playing $s_{l,j}$ regardless of the choice made by the other player. The recursive elimination of dominated strategies is one common technique for solving games. In the first step, all the dominated strategies are removed from the set of available strategies, since no rational player would ever play them. In this way a new smaller game is obtained. At this point, some strategies, that were not dominated before, may be dominated in the remaining game, and hence are eliminated. The process goes on until no dominated strategy exists for any player. A *rationalizable equilibrium* is any profile which survives the iterated elimination of dominated strategies. If at the end of the process only one profile is left, the remaining profile is said to be the *only rationalizable equilibrium* of the game, which is also the only Nash equilibrium point.

The corresponding strategies are the only rational choice for the two players and we say that the game is *dominance solvable*. Dominance solvable games are easy to analyze since, under the assumption of rational players, we can anticipate that the players will choose the strategies corresponding to the unique rationalizable equilibrium [29].

III. BINARY HYPOTHESIS TESTING GAME WITH KNOWN SOURCES

In this section, we use the framework introduced in [19] to define a version of the hypothesis testing game in which the pmf's ruling the emission of sequences from X and Y are known to both D and A. We also summarize the main results proven in [19], so to ease the comparison with the new results that will be proven in the rest of the paper.

Let H_0 and H_1 be two hypotheses characterized by the pmf's P_X and P_Y respectively, defined over the same alphabet \mathcal{X} . Let y^n be a sequence drawn under hypothesis H_1 and let z^n be a modified version of y^n produced by A in the attempt to deceive D. The binary hypothesis testing game under the known source assumption (HT_{ks}) is defined as follows.

Definition 1: The $HT_{ks}(S_D, S_A, u)$ game is a zero-sum, strategic, game played by D and A, defined by the following strategies and payoff.

- The set of strategies D can choose from is the set of acceptance regions Λ for which the false positive probability is below a certain threshold:

$$S_D = \{\Lambda : P_X(x^n \notin \Lambda) \leq P_{FP}\}. \quad (4)$$

- The set of strategies A can choose from is formed by all the functions that map a sequence $y^n \in \mathcal{X}^n$ into a new sequence $z^n \in \mathcal{X}^n$ subject to a distortion constraint:

$$S_A = \{f : d(y^n, z^n) \leq nL\}, \quad (5)$$

where $d(\cdot, \cdot)$ is a proper distortion measure and L is the maximum allowed per-letter distortion.

- The payoff function is defined as the false negative error probability (P_{FN}), namely:

$$u(\Lambda, f) = -P_{FN} = - \sum_{y^n: f(y^n) \in \Lambda} P_Y(y^n). \quad (6)$$

Given the difficulty of studying the game defined above, a simplified version of the game is introduced in [19] wherein the set of strategies available to D is limited. More specifically, the so-called limited resources assumption is introduced, forcing D to base its analysis only on first order statistics of x^n . Stated in another way, it is required that Λ is a union of type classes. Since a type class is univocally defined by the empirical probability mass function of the sequences contained in it, the acceptance region Λ can be seen as a union of types $P \in \mathcal{P}_n$. As an additional simplification, the constraint on the false positive probability is defined in asymptotic terms, requiring that P_{FP} decreases exponentially fast with a given decay rate. All these considerations lead to the following definition:

Definition 2: The $HT_{ks}^{lr}(S_D, S_A, u)$ game is a game between D and A defined by the following strategies and payoff:

$$S_D = \{\Lambda \in 2^{\mathcal{P}_n} : P_{FP} \leq 2^{-\lambda n}\}, \quad (7)$$

$$S_A = \{f : d(y^n, f(y^n)) \leq nL\}, \quad (8)$$

$$u(\Lambda, f) = -P_{FN}, \quad (9)$$

where $2^{\mathcal{P}_n}$ indicates the power set of \mathcal{P}_n . From the analysis given in [19] we know the following results.

Theorem 1: The profile $(\Lambda_{ks}^, f_{ks}^*)$ with*

$$\Lambda_{ks}^* = \left\{ P \in \mathcal{P}_n : \mathcal{D}(P||P_X) < \lambda - |\mathcal{X}| \frac{\log(n+1)}{n} \right\}, \quad (10)$$

and

$$f_{ks}^*(y^n) = \arg \min_{z^n: d(z^n, y^n) \leq nL} \mathcal{D}(P_{z^n}||P_X). \quad (11)$$

defines an asymptotic Nash equilibrium for the HT_{ks}^{lr} game.

As a matter of fact, from the proof given in [19], it is easy to see that $(\Lambda_{ks}^*, f_{ks}^*)$ is the only rationalizable equilibrium of the game and hence HT_{ks}^{lr} is a dominance solvable game.

A fundamental consequence of Theorem 1 is that the optimum strategies for D and A do not depend on P_Y hence making the assumption that P_Y is known irrelevant. With a few modifications, then, Theorem 1 can be applied to a composite hypothesis testing scenario in which only the pmf conditioned to H_0 is known [30].

The second main result proven in [19] regards the payoff at the equilibrium, and specifies under which conditions it is possible for D to devise a decision strategy such that P_{FN} tends to zero exponentially fast when n tends to infinity. Let Γ_{ks}^n be defined as follows:

$$\Gamma_{ks}^n = \{P \in \mathcal{P}_n : \forall y^n \in T(P), \exists z^n \in \Lambda_{ks}^* \text{ s.t. } d(y^n, z^n) \leq nL\} \quad (12)$$

and let the asymptotic version of Γ_{ks}^n be defined as

$$\Gamma_{ks}^\infty = cl\left(\bigcup_n \Gamma_{ks}^n\right), \quad (13)$$

where $cl(S)$ indicates the *closure* of the set S . The following theorem holds:

Theorem 2: For the HT_{ks}^{lr} game, the error exponent of the false negative error probability at the equilibrium is given by:

$$\varepsilon_{ks} = \min_{P \in \Gamma_{ks}^\infty} \mathcal{D}(P||P_Y), \quad (14)$$

leading to the following cases:

- 1) $\varepsilon_{ks} = 0$, if $P_Y \in \Gamma_{ks}^\infty$;
- 2) $\varepsilon_{ks} \neq 0$, if $P_Y \notin \Gamma_{ks}^\infty$.

where ε_{ks} indicates the false negative error exponent at the equilibrium and where the use of \min instead of \inf is justified by the compactness of Γ_{ks}^∞ . Given two pmf's P_X and P_Y , a maximum distortion L and the desired false positive error exponent λ , Theorem 2 permits to understand whether D may ever succeed to make the false negative error probability vanishingly small and thus *win* the game. As a matter of fact, this is possible only if $P_Y \notin \Gamma_{ks}^\infty$, since otherwise

$\varepsilon_{ks} = 0$. We will call Γ_{ks}^∞ the indistinguishability region for P_X , i.e. set of pmf's characterizing H_1 for which, under certain conditions (summarized by L and λ), hypothesis H_0 can not be distinguished from H_1 .

IV. BINARY HYPOTHESIS TESTING GAME WITH TRAINING DATA

The analysis carried out in [19] requires that P_X and P_Y are known to D and A (as we have seen in the previous section, in the asymptotic case only the knowledge of P_X is required); in real applications, however, this is rarely the case. As an example, let us consider again a multimedia forensics scenario in which D is asked to verify that a signal has been generated by a given acquisition device. It is very unlikely that a good statistical model of the device is available. On the contrary, it is likely that the analyst will build a suitable model to characterize H_0 by relying on a number of signals produced by the same acquisition device [31]. To get closer to a realistic scenario, then, we remove the assumption that P_X and P_Y are known and introduce the hypothesis testing game with training data.

Let \mathcal{C} be the class of discrete memoryless sources with alphabet \mathcal{X} , and let $X \sim P_X$ be a source in \mathcal{C} characterizing H_0 . As for the HT_{ks} game, D accepts or rejects H_0 by deciding whether a test sequence x^n was drawn from X or not. To make his decision, D relies on the knowledge of a training sequence t_D^N drawn from X . On his side, A takes a sequence y^n emitted by another source $Y \sim P_Y$ still belonging to \mathcal{C} and tries to modify it in such a way that D thinks that the modified sequence was generated by the same source that generated t_D^N . As usual, the attacker must satisfy a distortion constraint stating that the distance between the modified sequence and y^n must be lower than a threshold. Like the defender, A derives his knowledge about the statistics of the sequences generated under H_0 through a training sequence t_A^K drawn from P_X , that in general may not coincide with t_D^N . We assume that t_D^N , t_A^K , x^n and y^n are generated independently. With regard to P_Y , we could also assume that it is known through two training sequences, one available to A and one to D, however we will see that - as for known sources and at least in the asymptotic case - such an assumption is not necessary, and hence we take the simplifying assumption that P_Y is known to neither D nor A. In this framework, H_0 is equivalent to the hypothesis that the test sequence has been generated by the same source that generated t_D^N . As usual Λ indicates the acceptance region for H_0 , however now it is more convenient to think of Λ as a subset of $\mathcal{X}^n \times \mathcal{X}^N$, i.e., as the set of all the pairs of sequences (x^n, t_D^N) that the defender considers to be drawn from the same, unknown, source. With the above ideas in mind, and by paralleling the definition given in Section III, we define a first version of the binary hypothesis testing game with training sequences as follows:

Definition 3: The $HT_{r,a}(\mathcal{S}_D, \mathcal{S}_A, u)$ game is a zero-sum, strategic, game played by D and A, defined by the following strategies and payoff.

- The set of strategies D can choose from is the set of acceptance regions Λ for which the maximum false

positive probability across all possible $P_X \in \mathcal{C}$ is lower than a given threshold:

$$\mathcal{S}_D = \{\Lambda : \max_{P_X \in \mathcal{C}} P_X\{(x^n, t_D^N) \notin \Lambda\} \leq P_{FP}\}, \quad (15)$$

where P_{FP} is a prescribed maximum false positive probability, and where $P_X\{(x^n, t_D^N) \notin \Lambda\}$ indicates the probability that two independent sequences generated by X do not belong to Λ . Note that the acceptance region is defined as a union of pairs of sequences, and hence $\Lambda \subset \mathcal{X}^n \times \mathcal{X}^N$.

- The set of strategies A can choose from is formed by all the functions that map a sequence y^n generated by Y into a new sequence z^n subject to a distortion constraint:

$$\mathcal{S}_A = \{f : d(y^n, f(y^n, t_A^K)) \leq nL\}, \quad (16)$$

where $d(\cdot, \cdot)$ is a proper distortion function and L is the maximum allowed per-letter distortion. Note that the function $f(\cdot)$ depends on t_A^K , since when performing his attack A can exploit the knowledge of his training sequence.

- The payoff function is defined in terms of the false negative error probability, namely:

$$u(\Lambda, f) = -P_{FN} = -\sum_{\substack{t_D^N \in \mathcal{X}^n, t_A^K \in \mathcal{X}^K \\ y^n : (f(y^n, t_A^K), t_D^N) \in \Lambda}} P_Y(y^n) P_X(t_D^N) P_X(t_A^K), \quad (17)$$

where the error probability is averaged across all possible y^n and training sequences and where we have exploited the independence of y^n , t_D^N and t_A^K .

A. Discussion

Before going on with the analysis, we pause to discuss some of the choices we implicitly made with the above definition.

A first observation regards the payoff function. As a matter of fact, the expression in (17) looks problematic, since its evaluation requires that the pmf's P_X and P_Y are known, however this is not the case in our scenario since we have assumed that P_X is known only through t_D^N and t_A^K , and that P_Y is not known at all. As a consequence it may seem that the players of the game are not able to compute the payoff associated to a given profile and hence have no arguments upon which they can base their choice. While this is indeed a problem in a generic setup, we will show later on in the paper that asymptotically (when n , N and K tend to infinity) the optimum strategies of D and A are uniformly optimum across all P_X and P_Y and hence the ignorance of P_X and P_Y is not a problem. One may wonder why we did not define the payoff under a worst case assumption (from D's perspective) on P_X and/or P_Y . The reason is that doing so would result in a meaningless game since the worst case for D would always correspond to $P_Y = P_X$ for which no decision is possible.³

As a second remark, we stress that we decided to limit the strategies available to A to deterministic functions of y^n . This may seem a limiting choice, however we will see in the

³Alternatively, we could assume that X and Y belong to two disjoint source classes \mathcal{C}_X and \mathcal{C}_Y . We leave this analysis for further research.

subsequent sections that, at least asymptotically, the optimum strategy of D depends neither on the strategy chosen by A nor on P_Y , then, it does not make sense for A to adopt a randomized strategy to confuse D.

A last, even more basic, comment regards the overall structure of the game. In our definition we assumed that the attacker does not intervene when H_0 holds, since we restricted his interest to the false negative error probability. An alternative approach could be to let the attacker modify also the sequences generated under H_0 in the attempt to increase the false positive rate. We could also depart from the Neyman-Pearson set up and define the payoff in terms of the overall error probability, or the overall Bayes risk defined on the basis of suitable cost functions associated to the two kinds of errors.⁴ While these are interesting research directions, in this paper we restrict our analysis to the game specified by Definition 3, and leave the alternative approaches for future research.

B. A Variant of the Game

An interesting variant of the $HT_{ir,a}$ game is obtained by assuming that the training sequence available to A is equal to that available to D, leading to the following definition.

Definition 4: The $HT_{ir,b}(S_D, S_A, u)$ game is a zero-sum, strategic, game defined as the $HT_{ir,a}$ game with the only difference that $K = N$ and $t_A^K = t_D^N$ (simply indicated as t^N in the following). The set of strategies of D and A are the same as in the $HT_{ir,a}$ game, while the payoff is redefined as:

$$u(\Lambda, f) = -P_{FN} = -\sum_{\substack{t^N \in \mathcal{X}^N \\ y^n: (f(y^n, t^N), t^N) \in \Lambda}} P_Y(y^n) P_X(t^N). \quad (18)$$

Due to its simplicity, in the rest of the paper we will first focus on version b of the game, and then extend our results to cover version a as well.

C. Hypothesis Testing Game With Limited Resources ($HT_{ir,b}^{lr}$)

Studying the existence of an equilibrium point for the $HT_{ir,b}$ game is a prohibitive task, hence we use the same approach adopted in [18] and [19] and consider a simplified version of the game in which D can only base his decision on a limited set of statistics computed on the test and training sequences. Specifically, we require that D relies only on the relative frequencies with which the symbols in \mathcal{X} appear in x^n and t^N , i.e. P_{x^n} and P_{t^N} . Following [18] and [19], we call this version of the game *hypothesis testing with limited-resources*, and we refer to it as the $HT_{ir,b}^{lr}$ game. Note that P_{x^n} and P_{t^N} are not sufficient statistics for D, since even if Y is a memoryless source, the attacker could introduce some memory within the sequence as a result of the application of f . In the same way, he could introduce some dependencies between the attacked sequence z^n and t^N . It is then necessary to treat the assumption that D relies only on P_{x^n} and P_{t^N} as an

explicit requirement. At first sight, requiring that D relies only on the relative frequencies of the symbols of \mathcal{X} in x^n and t^N may seem too much a limiting assumption, however this is not the case. First of all, such an assumption is often met in several practical applications. In image forensics, for instance, it is common to rely on first order statistics of either frequency or spatial image samples. This is the case with all the systems proposed so far for the detection of double JPEG compression (which are commonly based on the first order statistics of DCT coefficients [32]), and the systems for the detection of contrast enhancement in grey-level images [33]. As a second observation, we note that our analysis (which strongly relies on the method of types) can be extended to include a wider class of acceptance regions, not necessarily defined in terms of relative frequencies only. In fact, even if the method of types was initially developed to work with memoryless sources, it can be extended to more complicated models, e.g. to sources with memory [34]. This is the case, for instance, of Markov sources with finite order [34] and renewal processes [35].

As a consequence of the limited resources assumption, Λ can only be a union of Cartesian products of pairs of type classes, i.e. if the pair of sequences (x^n, t^N) belongs to Λ , then any pair of sequences belonging to the Cartesian product $T(P_{x^n}) \times T(P_{t^N})$ will also be contained in Λ . Since a type class is univocally defined by the empirical pmf of the sequences contained in it, we can redefine Λ as a union of pairs of types (P, Q) with $P \in \mathcal{P}_n$ and $Q \in \mathcal{P}_N$. In the following, we will use the two interpretations of Λ (as a set of pairs of sequences or pairs of types) interchangeably, the exact meaning being always recoverable from the context.

We are interested in studying the asymptotic behavior of the game when n and N tend to infinity. To avoid the necessity of considering two limits with n and N tending to infinity independently, we will express N as a function of n , and study what happens when n tends to infinity. This assumption does not reduce the generality of our analysis, however it destroys the symmetry of the hypothesis testing problem with respect to the two sequences x^n and t^N . The consequences of this loss of symmetry will be discussed in Section V-A.

We are now ready to define the asymptotic $HT_{ir,b}^{lr}$ game. Specifically, we have:

Definition 5: The $HT_{ir,b}^{lr}(S_D, S_A, u)$ game is a zero-sum, strategic, game played by D and A, defined by the following strategies and payoff:

$$S_D = \{\Lambda \subset \mathcal{P}_n \times \mathcal{P}_N : \max_{P_X \in \mathcal{C}} P_X\{(x^n, t^{N(n)}) \notin \Lambda\} \leq 2^{-\lambda n}\}, \quad (19)$$

$$S_A = \{f : d(y^n, f(y^n, t^{N(n)})) \leq nL\}, \quad (20)$$

$$u(\Lambda, f) = -P_{FN} = -\sum_{\substack{t^{N(n)} \in \mathcal{X}^{N(n)} \\ y^n: (f(y^n, t^{N(n)}), t^{N(n)}) \in \Lambda}} P_Y(y^n) P_X(t^{N(n)}). \quad (21)$$

Note that we ask that the false positive error probability decays exponentially fast with n , thus opening the way to the asymptotic solution of the game. A similar definition can be given for version a of the game.

⁴In this case it would be necessary that the a priori probabilities of the two hypotheses are known.

V. ASYMPTOTIC EQUILIBRIUM OF THE $HT_{tr,b}^{lr}$ GAME

We start the analysis of the asymptotic equilibrium point of the $HT_{tr,b}^{lr}$ game by determining the optimum acceptance region for D. To do so we will use an analysis similar to that carried out in [23] to study hypothesis testing with observed statistics. The main difference between our analysis and [23] is the presence of the attacker, i.e. the game-theoretic nature of our problem. The derivation of the optimum strategy for D passes through the definition of the generalized log-likelihood ratio function $h(x^n, t^N)$. Given the test and training sequences x^n and t^N , the generalized log-likelihood ratio function is defined as ([23], [36])⁵:

$$h(x^n, t^N) = \mathcal{D}(P_{x^n} || P_{r^{n+N}}) + \frac{N}{n} \mathcal{D}(P_{t^N} || P_{r^{n+N}}), \quad (22)$$

where $P_{r^{n+N}}$ indicates the empirical pmf of the sequence r^{n+N} , obtained by concatenating x^n and t^N , i.e.

$$r_i = \begin{cases} x_i & i \leq n \\ t_{i-n} & n < i \leq n + N. \end{cases} \quad (23)$$

Observing that $h(x^n, t^N)$ depends on the test and the training sequences only through their empirical pmf, we can also use the notation $h(P_{x^n}, P_{t^N})$. The study of the equilibrium for the $HT_{tr,b}^{lr}$ game passes through the following lemmas.

Lemma 1: For any P_X we have:

$$\begin{aligned} n\mathcal{D}(P_{x^n} || P_{r^{n+N}}) + N\mathcal{D}(P_{t^N} || P_{r^{n+N}}) \\ \leq n\mathcal{D}(P_{x^n} || P_X) + N\mathcal{D}(P_{t^N} || P_X), \end{aligned} \quad (24)$$

with equality holding if only if $P_X = P_{r^{n+N}}$.

The proof of Lemma 1 is given in Appendix A.

Lemma 2: $\Lambda_{tr,b}^$ being defined as follows:*

$$\Lambda_{tr,b}^* = \left\{ (P_{x^n}, P_{t^N}) : \begin{aligned} h(P_{x^n}, P_{t^N}) < \lambda - |\mathcal{X}| \frac{\log(n+1)(N+1)}{n} \end{aligned} \right\} \quad (25)$$

with

$$\lim_{n \rightarrow \infty} \frac{\log(N(n)+1)}{n} = 0, \quad (26)$$

then:

- 1) $\max_{P_X} P_X \{(x^n, t^N) \notin \Lambda_{tr,b}^*\} \leq 2^{-n(\lambda - v_n)}$, with $v_n \rightarrow 0$, for $n \rightarrow \infty$,
- 2) $\forall \Lambda \in \mathcal{S}_D$, we have $\bar{\Lambda} \subseteq \bar{\Lambda}_{tr,b}^*$.

Proof: Being $\Lambda_{tr,b}^*$ a union of pairs of types (or, equivalently, a union of Cartesian products of type classes), we have:

$$\begin{aligned} \max_{P_X} P_{FP} &= \max_{P_X} \sum_{(x^n, t^N) \in \bar{\Lambda}_{tr,b}^*} P_X(x^n, t^N) \\ &= \max_{P_X} \sum_{(P_{x^n}, P_{t^N}) \in \bar{\Lambda}_{tr,b}^*} P_X(T(P_{x^n}) \times T(P_{t^N})). \end{aligned} \quad (27)$$

For the class of discrete memoryless sources, the number of types with denominators n and N is bounded by $(n+1)^{|\mathcal{X}|}$

⁵To simplify the notation, when it is not strictly necessary, we omit to indicate explicitly the dependence of N on n .

and $(N+1)^{|\mathcal{X}|}$ respectively [24], so we can write:

$$\begin{aligned} \max_{P_X} P_{FP} &\leq \max_{P_X} \max_{(P_{x^n}, P_{t^N}) \in \bar{\Lambda}_{tr,b}^*} \\ &[(n+1)^{|\mathcal{X}|} (N+1)^{|\mathcal{X}|} P_X(T(P_{x^n}) \times T(P_{t^N}))] \\ &\leq (n+1)^{|\mathcal{X}|} (N+1)^{|\mathcal{X}|} \cdot \max_{P_X} \\ &\max_{(P_{x^n}, P_{t^N}) \in \bar{\Lambda}_{tr,b}^*} 2^{-n[\mathcal{D}(P_{x^n} || P_X) + \frac{N}{n} \mathcal{D}(P_{t^N} || P_X)]}, \end{aligned} \quad (28)$$

where in the second inequality we have exploited the independence of x^n and t^N and the property of types according to which for any sequence x^n we have $P_X(T(P_{x^n})) \leq 2^{-n\mathcal{D}(P_{x^n} || P_X)}$ (see [24]). By exploiting Lemma 1, we can write:

$$\begin{aligned} \max_{P_X} P_{FP} &\leq (n+1)^{|\mathcal{X}|} (N+1)^{|\mathcal{X}|} \\ &\cdot \max_{(P_{x^n}, P_{t^N}) \in \bar{\Lambda}_{tr,b}^*} 2^{-n[\mathcal{D}(P_{x^n} || P_{r^{n+N}}) + \frac{N}{n} \mathcal{D}(P_{t^N} || P_{r^{n+N}})]} \\ &\leq (n+1)^{|\mathcal{X}|} (N+1)^{|\mathcal{X}|} 2^{-n(\lambda - |\mathcal{X}| \frac{\log(n+1)(N+1)}{n})} \\ &= 2^{-n(\lambda - 2|\mathcal{X}| \frac{\log(n+1)(N+1)}{n})}, \end{aligned} \quad (29)$$

where the last inequality derives from the definition of $\Lambda_{tr,b}^*$. Together with (26), equation (29) proves the first part of the lemma with $v_n = 2|\mathcal{X}| \frac{\log(n+1)(N+1)}{n}$.

For any $\Lambda \in \mathcal{S}_D$, let (x^n, t^N) be a generic pair of sequences contained in $\bar{\Lambda}$, due to the limited resources assumption the cartesian product between $T(P_{x^n})$ and $T(P_{t^N})$ will be entirely contained in $\bar{\Lambda}$. Then we have:

$$\begin{aligned} 2^{-\lambda n} &\geq \max_{P_X} P_X(\bar{\Lambda}) \\ &\stackrel{(a)}{\geq} \max_{P_X} P_X(T(P_{x^n}) \times T(P_{t^N})) \\ &\stackrel{(b)}{\geq} \max_{P_X} \frac{2^{-n[\mathcal{D}(P_{x^n} || P_X) + \frac{N}{n} \mathcal{D}(P_{t^N} || P_X)]}}{(n+1)^{|\mathcal{X}|} (N+1)^{|\mathcal{X}|}} \\ &\stackrel{(c)}{=} \frac{2^{-n[\mathcal{D}(P_{x^n} || P_{r^{n+N}}) + \frac{N}{n} \mathcal{D}(P_{t^N} || P_{r^{n+N}})]}}{(n+1)^{|\mathcal{X}|} (N+1)^{|\mathcal{X}|}}, \end{aligned} \quad (30)$$

where (a) is due to the limited resources assumption, (b) follows from the independence of x^n and t^N and a lower bound on the probability of a pair of type classes [24], and (c) derives from Lemma 1. By taking the logarithm of both sides we find that $(x^n, t^N) \in \bar{\Lambda}_{tr,b}^*$, thus completing the proof. ■

The first part of Lemma 2 shows that, at least asymptotically, $\Lambda_{tr,b}^*$ belongs to \mathcal{S}_D , while the second part implies the optimality of $\Lambda_{tr,b}^*$. An important observation is that the optimum strategy of D is univocally determined by the false positive constraint. This solves the apparent problem that we pointed out when defining the payoff of the game, namely that the payoff depends on P_X and P_Y and hence it is not fully known to D. We also observe that $\Lambda_{tr,b}^*$ does not depend on t_A^K , hence it is the optimum defender's strategy even for version a of the HT_{tr}^{lr} game. For this reason, from now on we will simply indicate it as Λ_{tr}^* .

The most important consequence of Lemma 2 is that the optimum strategy of D does not depend on the strategy chosen

by the attacker, that is Λ_{tr}^* is a strictly dominant strategy for D. In turn this simplifies the analysis of the optimum attacking strategy. In fact, a rationale defender will surely play the dominant strategy Λ_{tr}^* , hence A can choose his strategy by assuming that D chooses Λ_{tr}^* . The derivation of the optimum attacking strategy is now an easy task. We only need to observe that the goal of A is to take a sequence y^n drawn from Y and modify it in such a way that:

$$h(z^n, t^N) < \lambda - |\mathcal{X}| \frac{\log(n+1)(N+1)}{n}, \quad (31)$$

with $d(y^n, z^n) \leq nL$. The optimum attacking strategy, then, can be expressed as a minimization problem, i.e.:

$$f_{tr,b}^*(y^n, t^N) = \arg \min_{z^n: d(y^n, z^n) \leq nL} h(z^n, t^N). \quad (32)$$

Note that to implement this strategy A needs to know t^N , i.e., (32) determines the optimum strategy only for version b of the game.

Having determined the optimum strategies for D and A, we can state the first main result of the paper, summarized in the following theorem.

Theorem 3 (Asymptotic Equilibrium of the $HT_{tr,b}^{lr}$ Game): *The $HT_{tr,b}^{lr}$ game is a dominance solvable game and the profile $(\Lambda_{tr}^*, f_{tr,b}^*)$ is the only rationalizable equilibrium.*

Proof: Lemma 2 says that Λ_{tr}^* is a strictly dominant strategy for D, thus permitting us to eliminate all the other strategies in \mathcal{S}_D (since they are strictly dominated by Λ_{tr}^*). The theorem, then, follows by observing that $f_{tr,b}^*$ maximizes the false negative error probability when Λ_{tr}^* is fixed. In fact, if the minimum in (32) is not lower than the acceptance threshold, no other strategy will succeed in deceiving D. ■

A. Discussion

As a first remark, we observe that $(\Lambda_{tr}^*, f_{tr,b}^*)$ is the unique Nash equilibrium of the game. In addition to the properties of Nash equilibria, however, $(\Lambda_{tr}^*, f_{tr,b}^*)$ has the desirable characteristic of being the only possible choice if the two players behave rationally. In fact, a rational defender will surely adopt the acceptance region Λ_{tr}^* , since any other choice will lead to a (asymptotically) higher P_{FN} , regardless of the choice made by A. On his side, a rational attacker, knowing that D will behave rationally, will adopt the strategy $f_{tr,b}^*$, since this is the strategy that optimizes his payoff when D plays Λ_{tr}^* .

To get a better insight into the meaning of the equilibrium point of the $HT_{tr,b}^{lr}$ game, it is instructive to compare it with the equilibrium of the corresponding game with known sources, namely the HT_{ks}^{lr} game. To start with, we observe that the use of the h function instead of the divergence \mathcal{D} derives from the fact that D must ensure that the false positive probability stays below the desired threshold for all possible DMS's. To do so, he has to estimate the pmf that better explains the evidence provided by both x^n and t^N , that is the pmf maximizing the probability of observing x^n and t^N . As shown in the Appendix (see (A6)), such a maximizing pmf corresponds to $P_{r^{n+N}}$, and the generalized log-likelihood ratio corresponds to 1 over n the log of the (asymptotic) probability




Fig. 2. Geometric interpretation of the difference between \mathcal{D} (left) and h (right) functions.

that a source with pmf equal to $P_{r^{n+N}}$ outputs the sequences x^n and t^N . A geometric illustration of the difference between the \mathcal{D} and the h functions used, respectively, in the optimum acceptance regions for the HT_{ks}^{lr} and the $HT_{tr,b}^{lr}$ games is given in Fig. 2. Another observation regards the optimum strategy of the attacker. As a matter of fact, the functions $h(P_{x^n}, P_{t^N})$ and $\mathcal{D}(P_{x^n} || P_{t^N})$ share a similar behavior: both are positive and convex functions achieving the absolute minimum when $P_{x^n} = P_{t^N}$, so one may be tempted to think that from A's point of view minimizing $\mathcal{D}(P_{x^n} || P_{t^N})$ is equivalent to minimizing $h(P_{x^n}, P_{t^N})$. While this is the case in some situations, e.g. when the absolute minimum can be reached, in general the two minimization problems yield different solutions.

To further compare the $HT_{tr,b}^{lr}$ and the HT_{ks}^{lr} games, it is useful to rewrite the generalized likelihood function in a more convenient way. By applying some algebra, it is easy to prove the following equivalent expression for h :

$$h(P_{x^n}, P_{t^N}) = \mathcal{D}(P_{x^n} || P_{t^N}) - \frac{N+n}{n} \mathcal{D}(P_{r^{n+N}} || P_{t^N}), \quad (33)$$

showing that $h(P_{x^n}, P_{t^N}) \leq \mathcal{D}(P_{x^n} || P_{t^N})$ with the equality holding only in the trivial case $P_{x^n} = P_{t^N}$. This suggests that, at least for large n , it should be easier for A to bring a sequence generated by Y within Λ_{tr}^* than to bring it within Λ_{ks}^* . This is indeed the case, as it will be shown in Section VI-A, where we will provide a rigorous proof that the $HT_{tr,b}^{lr}$ game is actually more favorable to the attacker than the HT_{ks}^{lr} game.

We conclude this section by investigating the behavior of the optimal acceptance strategy for different values of the ratio $\frac{N}{n}$. To do so we introduce the two quantities $c_x = \frac{n}{n+N}$ and $c_t = \frac{N}{n+N}$, representing the weights of the sequences x^n and t^N in r^{n+N} . It is easy to show, in fact, that

$$P_{r^{n+N}} = c_x P_{x^n} + c_t P_{t^N}. \quad (34)$$

In the simplest case, n and N will tend to infinity with the same speed, hence we can assume that the ratio between N and n is fixed, namely, $\frac{N}{n} = c \neq 0$ (we obviously have $c_x = \frac{1}{1+c}$ and $c_t = \frac{c}{1+c}$). Under this assumption, the decision of D is dictated by (25) and no particular behavior can be noticed. This is not the case when N/n tends to 0 or ∞ .

If $N/n \rightarrow 0$, then $P_{r^{n+N}} \rightarrow P_{x^n}$ and $h(P_{x^n}, P_{t^N}) \rightarrow 0$. This means that the defender will always decide in favor of H_0 . This makes sense since when the test sequence is infinitely longer than the training sequence, the evidence provided by

the training sequence is not strong enough to let the defender reject hypothesis 0.

If $N/n \rightarrow \infty$, the analysis is slightly more involved. In this case $c_t \rightarrow 1$ and $P_{r^{n+N}} \rightarrow P_{t^N}$, hence the first term in (22) tends to $\mathcal{D}(P_{X^n}||P_{t^N})$. To understand the behavior of the second term of (22) when $n \rightarrow \infty$, we can use the Taylor expansion of $\mathcal{D}(P||Q)$ when P approaches Q (see [25, Ch. 4]), which applied to the second term of the h function yields:

$$\begin{aligned} \frac{N}{n} \cdot \mathcal{D}(P_{t^N}||P_{r^{n+N}}) &\approx \frac{N}{2n} \cdot \sum_x \frac{(P_{t^N}(x) - P_{r^{n+N}}(x))^2}{P_{r^{n+N}}(x)} \\ &= \frac{N}{2n} \cdot \sum_x \frac{(c_x P_{t^N}(x) - c_x P_{X^n}(x))^2}{P_{r^{n+N}}(x)} \\ &= \frac{\frac{n}{N}}{2(\frac{n}{N} + 1)^2} \sum_x \frac{(P_{t^N}(x) - P_{X^n}(x))^2}{P_{r^{n+N}}(x)}. \end{aligned} \quad (35)$$

When $N/n \rightarrow \infty$, the above expression clearly tends to 0, and hence $h(P_{X^n}, P_{t^N}) \rightarrow \mathcal{D}(P_{X^n}||P_{t^N})$. In other words, the optimum acceptance region tends to be equal to the one obtained for the case of known sources with P_X replaced by P_{t^N} . This is also an intuitively reasonable result: when the training sequence is much longer than the test sequence, the empirical pmf of the training sequence provides such a reliable estimate of P_X that the defender can treat it as the true pmf.

One may wonder the reason behind the asymmetric behavior of the optimum decision strategy when the length of one between the two sequences under analysis grows much faster than the other. This apparent anomaly derives from the choice of analyzing the asymptotic behavior by letting n tend to infinity, a choice that breaks the symmetry between the test and training sequences. If we had defined the false positive and false negative error exponents in terms of N , the situation would have been completely reversed.

In the following we will always assume that $N/n = c$, since from the above analysis this turns out to be most interesting case.

VI. ANALYSIS OF THE PAYOFF AT THE EQUILIBRIUM

Now that we have derived the equilibrium point of the $HT_{tr,b}^{lr}$ game, we are ready to analyze the payoff at the equilibrium to understand who, between the defender and the attacker is going to *win* the game. Our aim is to derive a result similar to Theorem 2, so that given two pmf's P_X and P_Y , a false positive error exponent λ and a distortion constraint L , we can derive the best achievable (for the defender) false negative error exponent $\varepsilon_{tr,b}$. Specifically, we would like to know whether it is possible for D to obtain a strictly positive value of $\varepsilon_{tr,b}$, thus ensuring that the false negative error probability tends to zero exponentially fast for increasing values of n .

In our proofs we will find it necessary to generalize the h function so that it can be applied to general pmf's not necessarily belonging to \mathcal{P}_n or \mathcal{P}_N . By remembering that $N/n = c$, we introduce the following definition:

$$h_c(P, Q) = \mathcal{D}(P||U) + c\mathcal{D}(Q||U), \quad (36)$$

with

$$U = \frac{1}{1+c}P + \frac{c}{1+c}Q. \quad (37)$$

Note that when $P \in \mathcal{P}_n$ and $Q \in \mathcal{P}_N$, the above definition is equivalent to (22). By using h_c instead of h , we can generalize the expression of the optimum acceptance region Λ_{tr}^* , so to make it possible to apply it to any pair of pmf's P and Q (of course when P and Q are not empirical pmf's the meaning of Λ_{tr}^* as acceptance region for H_0 is lost):

$$\Lambda_{tr}^* = \left\{ (P, Q) : h_c(P, Q) < \lambda - |\mathcal{X}| \frac{\log(n+1)(N+1)}{n} \right\}. \quad (38)$$

With these ideas in mind, let us introduce the set $\Gamma_{tr,b}^n$ containing all the pairs of sequences (y^n, t^N) , for which A is able to bring y^n within Λ_{tr}^* (for sake of clarity we use the apex n to explicitly indicate that Λ_{tr}^* refers to pairs of sequences respectively of length n and $N = cn$):

$$\begin{aligned} \Gamma_{tr,b}^n &= \{(y^n, t^N) : \exists z^n \\ &\text{s.t. } (z^n, t^N) \in \Lambda_{tr}^* \text{ and } d(y^n, z^n) \leq nL\}. \end{aligned} \quad (39)$$

By observing that Λ_{tr}^* depends on t^N only through P_{t^N} and by reasoning as in the proof of [19, Property 1] (we need to assume that the distortion measure d is permutation-invariant), we can show that $\Gamma_{tr,b}^n$ is still a union of pairs of type classes, and hence we can redefine it as:

$$\begin{aligned} \Gamma_{tr,b}^n &= \{(P_{y^n}, P_{t^N}) : \forall y^n \in T(P_{y^n}) \exists z^n \\ &\text{s.t. } (P_{z^n}, P_{t^N}) \in \Lambda_{tr}^* \text{ and } d(y^n, z^n) \leq nL\}. \end{aligned} \quad (40)$$

Note that, by adopting the generalized version of Λ_{tr}^* in which h_c is used instead of h , the above definition can also be applied when P_{t^N} is replaced by a generic pmf Q not necessarily belonging to \mathcal{P}_N . We will also find it convenient to fix Q and consider the set of types P_{X^n} for which (P_{X^n}, Q) belongs to Λ_{tr}^* and $\Gamma_{tr,b}^n$, that is:

$$\begin{aligned} \Lambda_{tr}^*(Q) &= \{P_{X^n} : (P_{X^n}, Q) \in \Lambda_{tr}^*\}, \\ \Gamma_{tr,b}^n(Q) &= \{P_{y^n} : \forall y^n \in T(P_{y^n}) \exists z^n \\ &\text{s.t. } P_{z^n} \in \Lambda_{tr}^*(Q) \text{ and } d(y^n, z^n) \leq nL\}. \end{aligned} \quad (41)$$

The derivation of the false negative error exponent at the equilibrium passes through the following asymptotic extension of $\Gamma_{tr,b}^n(Q)$:

$$\Gamma_{tr,b}^\infty(Q) = cl \left(\bigcup_n \Gamma_{tr,b}^n(Q) \right). \quad (43)$$

The importance of the above definition is that for any source P_X , given λ and the maximum allowed per-letter distortion L , the set $\Gamma_{tr,b}^\infty(P_X)$ corresponds to the indistinguishability region of the $HT_{tr,b}^{lr}$ game, i.e. the set of all the pmf's for which D does not succeed in distinguishing between H_0 and H_1 ensuring a false negative error probability that tends to zero exponentially fast. In other words, if $P_Y \in \Gamma_{tr,b}^\infty(P_X)$, no strictly positive false negative error exponent can be achieved by D . To prove that this is indeed the case, we need to prove the following theorem.

Theorem 4 (Asymptotic Payoff of the $HT_{tr,b}^{lr}$ Game): For the $HT_{tr,b}^{lr}$ game, with $N/n = c$ and assuming an additive distortion measure, the false negative error exponent at the equilibrium is given by

$$\varepsilon_{tr,b} = \min_Q [c \cdot \mathcal{D}(Q||P_X) + \min_{P \in \Gamma_{tr,b}^\infty(Q)} \mathcal{D}(P||P_Y)]. \quad (44)$$

Proof: By using the definitions given in this section, the false negative error probability at the equilibrium, for a given n , can be written:

$$\begin{aligned} P_{FN} &= \sum_{t^N} P_X(t^N) \sum_{P \in \Gamma_{tr,b}^n(P_{t^N})} P_Y(T(P)) \\ &= \sum_{Q \in \mathcal{P}_N} P_X(T(Q)) \sum_{P \in \Gamma_{tr,b}^n(Q)} P_Y(T(P)). \end{aligned} \quad (45)$$

We start by deriving an upper-bound of the false negative error probability. By exploiting some well-known bounds on the probability of a type class and the number of types in \mathcal{P}_n [24], we can write:

$$\begin{aligned} P_{FN} &\leq \sum_{Q \in \mathcal{P}_N} P_X(T(Q)) \sum_{P \in \Gamma_{tr,b}^n(Q)} 2^{-n\mathcal{D}(P||P_Y)} \\ &\leq \sum_{Q \in \mathcal{P}_N} P_X(T(Q)) (n+1)^{|\mathcal{X}|} 2^{-n \min_{P \in \Gamma_{tr,b}^n(Q)} \mathcal{D}(P||P_Y)} \\ &\leq \sum_{Q \in \mathcal{P}_N} P_X(T(Q)) (n+1)^{|\mathcal{X}|} 2^{-n \min_{P \in \Gamma_{tr,b}^\infty(Q)} \mathcal{D}(P||P_Y)} \\ &\leq (n+1)^{|\mathcal{X}|} (N+1)^{|\mathcal{X}|} \\ &\quad \cdot 2^{-n \min_{Q \in \mathcal{P}_N} [\frac{N}{n} \mathcal{D}(Q||P_X) + \min_{P \in \Gamma_{tr,b}^\infty(Q)} \mathcal{D}(P||P_Y)]} \\ &\leq (n+1)^{|\mathcal{X}|} (N+1)^{|\mathcal{X}|} \\ &\quad \cdot 2^{-n \min_Q [c\mathcal{D}(Q||P_X) + \min_{P \in \Gamma_{tr,b}^\infty(Q)} \mathcal{D}(P||P_Y)]}, \end{aligned} \quad (46)$$

where the last inequality is obtained by minimizing over all Q without requiring that $Q \in \mathcal{P}_N$ and where the use of the minimum instead of the infimum is justified by the fact that $\Gamma_{tr,b}^n(Q)$ and $\Gamma_{tr,b}^\infty(Q)$ are compact sets. By taking the log and dividing by n we find:

$$\begin{aligned} -\frac{\log P_{FN}}{n} &\geq \min_{Q \in \mathcal{C}} [c\mathcal{D}(Q||P_X) \\ &\quad + \min_{P \in \Gamma_{tr,b}^\infty(Q)} \mathcal{D}(P||P_Y)] + \alpha_n, \end{aligned} \quad (47)$$

with $\alpha_n = |\mathcal{X}| \frac{\log(n+1)(N+1)}{n}$ tending to 0 when n tends to infinity.

We now turn to the analysis of a lower bound for P_{FN} . Let Q^* be the pmf achieving the minimum in (44). Due to the density of rational numbers within real numbers, we can find a sequence of pmf's $Q_n \in \mathcal{P}_n$ that tends to Q^* when n tends to infinity. By remembering that $N = nc$, the subsequence $Q_N = Q_{nc}$ will also tend to Q^* when n (and hence N) tends to infinity.⁶ Let us now consider the following sequence of

⁶In order to simplify the analysis, we assume that c is a non-null integer value, the extension of the proof to non-integer values of c is tedious but straightforward.

inequalities:

$$\begin{aligned} P_{FN} &\stackrel{(a)}{\geq} \sum_{Q \in \mathcal{P}_N} P_X(T(Q)) \sum_{P \in \Gamma_{tr,b}^n(Q)} \frac{2^{-n\mathcal{D}(P||P_Y)}}{(n+1)^{|\mathcal{X}|}} \\ &\geq \sum_{Q \in \mathcal{P}_N} P_X(T(Q)) \frac{2^{-n \min_{P \in \Gamma_{tr,b}^n(Q)} \mathcal{D}(P||P_Y)}}{(n+1)^{|\mathcal{X}|}} \\ &\stackrel{(b)}{\geq} \sum_{Q \in \mathcal{P}_N} \frac{2^{-N\mathcal{D}(Q||P_X)}}{(N+1)^{|\mathcal{X}|}} \frac{2^{-n \min_{P \in \Gamma_{tr,b}^n(Q)} \mathcal{D}(P||P_Y)}}{(n+1)^{|\mathcal{X}|}} \\ &= \sum_{Q \in \mathcal{P}_N} \frac{2^{-n[c\mathcal{D}(Q||P_X) + \min_{P \in \Gamma_{tr,b}^n(Q)} \mathcal{D}(P||P_Y)]}}{(N+1)^{|\mathcal{X}|} (n+1)^{|\mathcal{X}|}} \\ &\stackrel{(c)}{\geq} \frac{2^{-n[c\mathcal{D}(Q_N||P_X) + \min_{P \in \Gamma_{tr,b}^n(Q_N)} \mathcal{D}(P||P_Y)]}}{(N+1)^{|\mathcal{X}|} (n+1)^{|\mathcal{X}|}}, \end{aligned} \quad (48)$$

where inequalities (a) and (b) derive from a known lower bound on the probability of a type class [24], and in (c) we have replaced the sum with a single element of the subsequence Q_N defined previously. By taking the log and dividing by n , we obtain

$$-\frac{\log P_{FN}}{n} \leq c\mathcal{D}(Q_N||P_X) + \min_{P \in \Gamma_{tr,b}^n(Q_N)} \mathcal{D}(P||P_Y) + \alpha_n, \quad (49)$$

where, as in (47), $\alpha_n = |\mathcal{X}| \frac{\log(n+1)(N+1)}{n}$ tends to 0 when n tends to infinity. To continue, let P^* be defined as follows

$$P^* = \arg \min_{P \in \Gamma_{tr,b}^\infty(Q^*)} \mathcal{D}(P||P_Y). \quad (50)$$

In Appendix B, we show that it is possible to find a sequence P_n , where each P_n belongs to $\Gamma_{tr,b}^n(Q_N)$, that tends to P^* when n tends to infinity. By starting from (49) and by exploiting the continuity of the divergence function, for n large enough we can write

$$\begin{aligned} -\frac{\log P_{FN}}{n} &\leq c\mathcal{D}(Q^*||P_X) + \beta'_n + \mathcal{D}(P_n||P_Y) + \alpha_n, \\ &\leq c\mathcal{D}(Q^*||P_X) + \beta'_n + \mathcal{D}(P^*||P_Y) + \beta''_n + \alpha_n, \end{aligned} \quad (51)$$

where all the sequences α_n , β'_n and β''_n tend to zero when n tends to infinity.

By coupling equations (47) and (51) and by letting $n \rightarrow \infty$, we eventually obtain:

$$-\lim_{n \rightarrow \infty} \frac{\log P_{FN}}{n} = \min_Q [c \cdot \mathcal{D}(Q||P_X) + \min_{P \in \Gamma_{tr,b}^\infty(Q)} \mathcal{D}(P||P_Y)], \quad (52)$$

thus proving the theorem. \blacksquare

According to Theorem 4, we can distinguish two cases depending on the relationship between P_X and P_Y .

- 1) $P_Y \in \Gamma_{tr,b}^\infty(P_X)$ then $\varepsilon_{tr,b} = 0$;
- 2) $P_Y \notin \Gamma_{tr,b}^\infty(P_X)$ then $\varepsilon_{tr,b} > 0$.

In the former case, for which the minimum in (44) is obtained by letting $Q^* = P_X$, it is not possible for D to obtain a strictly positive false negative error exponent while ensuring




Fig. 3. Geometric interpretation of the sets $\Lambda_{tr}^{*,\infty}$ and $\Gamma_{tr,b}^{\infty}$. When $P_Y \in \Gamma_{tr,b}^{\infty}$, a reliable distinction between H_0 and H_1 is not possible and the attacker wins the game.

that the false positive error exponent is at least equal to λ . In the latter case, it is not possible that the two divergences in (44) are simultaneously equal to zero, hence P_{FN} tends to 0 exponentially fast. In other words, given λ and L , the condition $P_Y \notin \Gamma_{tr,b}^{\infty}(P_X)$ ensures that the distance between P_Y and P_X is large enough to allow a reliable discrimination between H_0 and H_1 despite the presence of the adversary. As anticipated, then, $\Gamma_{tr,b}^{\infty}(P_X)$ is the indistinguishability region of the $HT_{tr,b}^{lr}$ game. A pictorial representation of the sets $\Lambda_{tr}^{*,\infty}$ (obtained by letting $n \rightarrow \infty$ in (38)) and $\Gamma_{tr,b}^{\infty}$ is given in Fig. 3.

A. Comparison Between the HT_{ks}^{lr} and $HT_{tr,b}^{lr}$ Games

In this section we compare the asymptotic performance achievable by D for the HT_{ks}^{lr} and $HT_{tr,b}^{lr}$ games. We start the analysis by comparing the indistinguishability regions of the two games, namely $\Gamma_{ks}^{\infty}(P_X)$ and $\Gamma_{tr,b}^{\infty}(P_X)$ (where, as opposed to Section III, we now explicitly indicate the dependence of Γ_{ks}^{∞} on P_X).

The comparison between the two regions relies on the comparison between the divergence and the generalized likelihood function. In particular, the starting point of our analysis is the following lemma.

Lemma 3 (Relationship Between h_c and \mathcal{D}): Let $N/n = c$, with $c \neq 0$, $c \neq \infty$, for any $P \neq P_X$ we have,

$$h_c(P, P_X) < \mathcal{D}(P||P_X). \quad (53)$$

Proof: By rewriting $h_c(P, P_X)$ as in (33), we have:

$$h_c(P, P_X) = \mathcal{D}(P||P_X) - (1+c)\mathcal{D}(U||P_X), \quad (54)$$

with $U = P/(1+c) + cP_X/(1+c)$, which is equal to P_X if and only if $P = P_X$, when we have $\mathcal{D}(U||P_X) = 0$ thus yielding $h_c(P, P_X) = \mathcal{D}(P||P_X) = 0$. ■

In the subsequent proofs we will refer to the way the mapping function f operates on y^n to produce z^n .⁷ Specifically, we will indicate with $n_f(i \rightarrow j)$ the number of times that f transforms the i -th symbol of \mathcal{X} into the j -th one. The main result of our analysis is stated in the following theorem.

⁷To keep the notation as light as possible, we will not distinguish between the mapping functions used for the known source case and those applying to hypothesis testing with training sequences, even if, rigorously speaking, these are quite different functions since the latter also depend on the training sequence t^N .

Theorem 5 ($HT_{tr,b}^{lr}$ vs HT_{ks}^{lr}): For any finite, non-null value of c , any P_X , $\lambda > 0$ and L we have

$$\Gamma_{ks}^{\infty}(P_X) \subset \Gamma_{tr,b}^{\infty}(P_X). \quad (55)$$

Proof: We will prove the theorem by first showing that $\Gamma_{ks}^{\infty}(P_X) \subseteq \Gamma_{tr,b}^{\infty}(P_X)$, and then finding at least one point (actually an infinite set of points) that belongs to $\Gamma_{tr,b}^{\infty}(P_X)$ but does not stay in $\Gamma_{ks}^{\infty}(P_X)$.

The idea behind the proof that $\Gamma_{ks}^{\infty}(P_X) \subseteq \Gamma_{tr,b}^{\infty}(P_X)$ is that since $h_c(P, P_X) < \mathcal{D}(P||P_X)$ and given that the thresholds defining the acceptance regions for both the $HT_{tr,b}^{lr}$ and the HT_{ks}^{lr} games tend to the same value λ for large values of n , if the attacker is able to bring a sequence within $\Lambda_{ks}^{*,n}$ he is also able to bring it within $\Lambda_{tr}^{*,m}$ (for some large values of n and m). We now turn the above idea into a rigorous proof. Let y^n be a sequence such that $P_{y^n} \in \Gamma_{ks}^{\infty}(P_X)$, this means that a mapping f_n exists that transforms y^n into a sequence z^n belonging to $\Lambda_{ks}^{*,n}(P_X)$, while satisfying the distortion constraint. Let now m be a multiple of n ($m = kn$). For k large enough we have

$$\lambda - |\mathcal{X}| \frac{\log(n+1)}{n} < \lambda - |\mathcal{X}| \frac{\log(m+1)(cm+1)}{m}. \quad (56)$$

Since $P_{z^n} \in \Lambda_{ks}^{*,n}(P_X)$, Lemma 3 permits us to write:

$$\begin{aligned} h_c(P_{z^n}, P_X) &\leq \mathcal{D}(P_{z^n}||P_X) \\ &< \lambda - |\mathcal{X}| \frac{\log(n+1)}{n} \\ &< \lambda - |\mathcal{X}| \frac{\log(m+1)(cm+1)}{m}. \end{aligned} \quad (57)$$

Given that m is a multiple of n , any $P \in \mathcal{P}_n$ also belongs to \mathcal{P}_m , permitting us to conclude that $P_{z^n} \in \Lambda_{tr}^{*,m}(P_X)$. Let now y^m be an m -long sequence having the same type of y^n (this is surely possible since m is a multiple of n). If we apply to y^m a mapping function v_m for which $n_{v_m}(i \rightarrow j) = kn_{f_n}(i \rightarrow j)$, the sequence $z^m = v_m(y^m)$ will have the same type of z^n , and hence by virtue of (57) $P_{z^m} \in \Lambda_{tr}^{*,m}(P_X)$. In addition, the mapping v_m introduces the same per-letter distortion of f_n for any additive distortion measure, permitting us to conclude that $P_{y^m} \in \Gamma_{tr,b}^{\infty}(P_X)$. In summary, we have shown that for any $P \in \Gamma_{ks}^{\infty}(P_X)$ an $m = kn$ exists such that $P \in \Gamma_{tr,b}^{\infty}(P_X)$, and hence:

$$\begin{aligned} \Gamma_{ks}^{\infty}(P_X) &= cl \left(\bigcup_n \Gamma_{ks}^n(P_X) \right) \\ &\subseteq cl \left(\bigcup_m \Gamma_{tr,b}^m(P_X) \right) = \Gamma_{tr,b}^{\infty}(P_X). \end{aligned} \quad (58)$$

We now prove that there is at least one point that belongs to $\Gamma_{tr,b}^{\infty}(P_X)$ but does not belong to $\Gamma_{ks}^{\infty}(P_X)$ (actually we will prove that there is an infinite number of such points). The main idea behind the proof is the following. For the types on the border of $\Gamma_{tr,b}^{\infty}(P_X)$, or arbitrarily close to the border but still within $\Gamma_{tr,b}^{\infty}(P_X)$, the optimum attacking strategy will only succeed in modifying the sequences within the type class so that the h function is lower than the acceptance threshold (which for large n is essentially equal to λ) but arbitrarily close to it. Given that the divergence is strictly larger than the h_c function (see Lemma 3) and that the thresholds of the two

games tend both to λ , it is impossible for the attacker to modify the sequences belonging to the type classes on the border of $\Gamma_{tr,b}^\infty(P_X)$, or very close to it, so to satisfy the acceptance conditions of the HT_{ks}^{lr} game, thus proving that such types do not belong to $\Gamma_{ks}^\infty(P_X)$. More rigorously, let us consider a point P^* belonging to the boundary of $\Gamma_{tr,b}^\infty(P_X)$. Since by definition $\Gamma_{tr,b}^\infty(P_X)$ is a closed set, $P^* \in \Gamma_{tr,b}^\infty(P_X)$. Due to Lemma 4 (Appendix B), we can find a sequence of types P_n^i in $\Gamma_{tr,b}^\infty(P_X)$ that tends to P^* from the inside of $\Gamma_{tr,b}^\infty(P_X)$. On the other hand, due to the density of rational numbers in real numbers, it is possible to find also an outer sequence of types P_n^o , i.e. lying inside $\overline{\Gamma_{tr,b}^\infty(P_X)}$, tending to P^* for $n \rightarrow \infty$, and for which $P_n^o \in \mathcal{P}_n$ for each n . By construction, for each n , no mapping function exists that when applied to the sequences in the type class of P_n^o moves them into $\Lambda_{tr}^{*,n}(P_X)$ with a per-letter distortion equal or lower than L . In other words, for any sequence $y^n \in P_n^o$ and any (distortion-limited) mapping $f(y^n) = w^n$ we have

$$h_c(P_{w^n}, P_X) \geq \lambda - |\mathcal{X}| \frac{\log(n+1)(cn+1)}{n}. \quad (59)$$

Let f^n be the sequence of optimum mapping functions that applied to the sequences in $T(P_n^i)$ result in a sequence z^n for which $h_c(P_{z^n}, P_X) < \lambda - |\mathcal{X}|[\log(n+1)(cn+1)]/n$. We argue that

$$\lim_{n \rightarrow \infty} h_c(P_{z^n}, P_X) = \lambda. \quad (60)$$

To show that this is indeed the case, we observe that

$$\limsup_{n \rightarrow \infty} h_c(P_{z^n}, P_X) \leq \lambda, \quad (61)$$

because by construction the types P_n^i belong to $\Gamma_{tr,b}^\infty(P_X)$. On the other side, we can prove that

$$\liminf_{n \rightarrow \infty} h_c(P_{z^n}, P_X) \geq \lambda. \quad (62)$$

In fact, if such a limit were lower than λ , say λ' , a type P_n^i would always exist such that n is arbitrarily large and for any sequences in $T(P_n^i)$ we have $h_c(P_{z^n}, P_X) = \lambda' + \beta$ with an arbitrarily small β (e.g. lower than $\lambda - \lambda'$). Then, by choosing n large enough, we could find a type P_m^o (with large m) which is arbitrarily close to P_n^i . By applying the optimum mapping f_n to the sequences in $T(P_m^o)$,⁸ by the continuity of the h_c function and the mapping f_n , we would obtain a sequence w^m satisfying the distortion constraint and for which $h_c(P_{w^m}, P_X)$ is arbitrarily close to λ' and hence lower than λ , which contradicts the hypothesis that the types of the sequence P_m^o do not belong to $\Gamma_{tr,b}^\infty(P_X)$.

Let us now consider the following positive quantity

$$\tau = \mathcal{D}(P^* || P_X) - h_c(P^*, P_X). \quad (63)$$

⁸Rigorously speaking, there's no guarantee that f_n can be applied as is to the sequences in P_m^o . To ensure that this is the case, we should first of all choose $m = kn$ for some large k , define a new mapping v_m such that $n_{v_m}(i \rightarrow j) = kn f_n(i \rightarrow j)$ and then modify v_m so that $n_{v_m}(i \rightarrow j)$ is never larger than the number of times that the symbol i appears in the to-be-mapped sequence. Given that the types P_n^i and P_m^o are arbitrarily close, and given that n and m can be taken arbitrarily large, the application of the modified mapping will result in a sequence z^m for which $h_c(P_{z^m}, P_X)$ is arbitrarily close to $h_c(P_{z^n}, P_X)$.

Due to (60), for large n , for any mapping function, and for all the sequences in $T(P_n^i)$, by the continuity of the \mathcal{D} and the h_c functions we, have

$$\begin{aligned} \mathcal{D}(P_{z^n} || P_X) &= h_c(P_{z^n}, P_X) + \tau + \beta_n \\ &> \lambda - \frac{\tau}{2} + \tau + \beta_n > \lambda, \end{aligned} \quad (64)$$

where z^n indicates the output of the optimum mapping function, β_n is a vanishingly small sequence and where the last inequality holds for sufficiently large n , hence proving that for large n the types in the sequence P_n^i do not belong to $\Gamma_{ks}^\infty(P_X)$. ■

Theorem 5 has two simple corollaries.

Corollary 1: For any pmf P belonging to the boundary of $\Gamma_{ks}^\infty(P_X)$ there exists a positive value ε such that $B(P, \varepsilon) \subset \Gamma_{tr,b}^\infty(P_X)$, where $B(P, \varepsilon)$ is a ball centered in P with radius ε . In the same way, for any pmf P belonging to the boundary of $\Gamma_{tr,b}^\infty(P_X)$ there exists a positive value ε such that $B(P, \varepsilon) \cap \Gamma_{ks}^\infty(P_X) = \emptyset$.

Proof: It follows immediately from the proof of Theorem 5. ■

Corollary 2: Let ε_{ks} and $\varepsilon_{tr,b}$ denote the error exponents at the equilibrium for the HT_{ks}^{lr} and $HT_{tr,b}^{lr}$ games. Then we have:

$$\varepsilon_{tr,b} \leq \varepsilon_{ks}, \quad (65)$$

where the equality holds if and only if $P_Y \in \Gamma_{ks}^\infty(P_X)$, when both error exponents are equal to 0.

Proof: The corollary is obvious when $P_Y \in \Gamma_{tr,b}^\infty(P_X)$, since in this case $\varepsilon_{tr,b} = 0$ while ε is equal to zero if $P_Y \in \Gamma_{ks}^\infty(P_X)$ and nonzero otherwise. When $P_Y \notin \Gamma_{tr,b}^\infty(P_X)$, by considering the expression of the error exponent for the $HT_{tr,b}^{lr}$ game we have:

$$\begin{aligned} \varepsilon_{tr,b} &= \min_Q [c \cdot \mathcal{D}(Q || P_X) + \min_{P \in \Gamma_{tr,b}^\infty(Q)} \mathcal{D}(P || P_Y)] \\ &\leq c \mathcal{D}(P_X || P_X) + \min_{P \in \Gamma_{tr,b}^\infty(P_X)} \mathcal{D}(P || P_Y) \\ &\stackrel{(a)}{=} \min_{P \in \Gamma_{tr,b}^\infty(P_X)} \mathcal{D}(P || P_Y) \\ &< \min_{P \in \Gamma_{ks}^\infty(P_X)} \mathcal{D}(P || P_Y) = \varepsilon_{ks}, \end{aligned} \quad (66)$$

where the last strict inequality is justified by observing that the absolute minimum of $\mathcal{D}(P || P_Y)$ is obtained for $P = P_Y$ which we have assumed to lie outside $\Gamma_{tr,b}^\infty(P_X)$ and hence, due to the convexity of \mathcal{D} , the value P satisfying the minimization on the right-hand side of equality (a) belongs to the boundary of $\Gamma_{tr,b}^\infty(P_X)$ and then, by Corollary 1, lies outside the closed set $\Gamma_{ks}^\infty(P_X)$. ■

Theorem 5 and Corollary 2 permit us so to conclude that hypothesis testing with training data is more favorable to the attacker than hypothesis testing with known sources. The reason behind such a result is the use of the h function instead of the divergence, which in turn stems from the need for the defender to ensure that the constraint on the false positive error probability is satisfied for all $P_X \in \mathcal{C}$. It is such a worst case assumption that ultimately favors the attacker in the $HT_{tr,b}^{lr}$ game.

VII. BINARY HYPOTHESIS TESTING GAME WITH INDEPENDENT TRAINING SEQUENCES ($HT_{tr,a}^{lr}$)

We now pass to the analysis of version a of the HT_{tr}^{lr} game. We remind that in this case D and A rely on independent training sequences, namely t_D^N and t_A^K . As for version c , we assume that both N and K grow linearly with n and that the asymptotic analysis is carried out by letting n go to infinity. As a matter of fact, assuming that K grows more than linearly with n is not reasonable in practical applications, since usually the defender has a better knowledge of the system than the attacker. This is the case, for instance, in source identification for multimedia forensics, where we can assume that the analyst has a better knowledge of the statistics of authorized sources. On the contrary, one could consider the case where K grows less than linearly with n . While this could be an interesting research direction, we leave it for future research.

Given the above, in the following, we assume that $N = cn$ and $K = dn$. As we already noted in Section V, the strategy Λ_{tr}^* identified by Lemma 2 is optimum regardless of the relationship between t_D^N and t_A^K , hence the only difference between versions a and b of the game is in the strategy of the attacker. In fact, now the attacker does not have a perfect knowledge of the acceptance region adopted by the defender, since such a region depends on the empirical pmf of t_D^N which A does not know.

A reasonable strategy for the attacker could be to use the empirical pmf of t_A^K instead of the one derived from t_D^N . More precisely, by using the notation introduced in Section VI (equation (42)), the attacker could try to move y^n into $\Lambda_{tr}^{*,n}(P_{t_A^K})$, while the acceptance region adopted by the defender is $\Lambda_{tr}^{*,n}(P_{t_D^N})$. Given that t_D^N and t_A^K are generated by the same source, their empirical pmf's will both tend to P_X when n goes to infinity, and hence using $\Lambda_{tr}^{*,n}(P_{t_A^K})$ should be *in some way* equivalent to using $\Lambda_{tr}^{*,n}(P_{t_D^N})$. In fact, in the following we will show that, given P_X , L and λ , the indistinguishability region for version a of the game, let us call it $\Gamma_{tr,a}^\infty(P_X)$, is identical to the indistinguishability region of version b . Of course, this does not mean that the achievable payoff for the $HT_{tr,a}^{lr}$ game is equal to that of the $HT_{tr,b}^{lr}$ game, since, even if the indistinguishability region is the same, outside it the false negative error exponent for case a may be different (actually larger) than that of case b .

We start our analysis by assuming that $c = d$ (and hence $N = K$), i.e. the training sequences available to the defender and the attacker have the same length. Our goal is to investigate the asymptotic behavior of the payoff of the $HT_{tr,a}^{lr}$ game for the profile $(\Lambda_{tr}^{*,n}(P_{t_D^N}), \tilde{f})$, where the, not necessarily optimum, strategy \tilde{f} adopted by the attacker is defined as:

$$\tilde{f}(y^n, t_A^N) = \arg \min_{z^n: d(z^n, y^n) \leq nL} h(P_{z^n}, P_{t_A^N}). \quad (67)$$

A problem with the above strategy arises when the minimization has several solutions $\{z^n(1) \dots z^n(k)\}$, some of which, say the first m , fall within $\Lambda_{tr}^{*,n}(P_{t_D^N})$ while the others don't. If we consider a permutation of y^n , say $\sigma(y^n)$, instead of y^n , $\{\sigma(z^n(1)) \dots \sigma(z^n(k))\}$ will still be solutions of the minimization problem. In addition, the first m sequences will

still belong to $\Lambda_{tr}^{*,n}(P_{t_D^N})$, while the others will not. Since we want that $\tilde{\Gamma}_{tr,a}^n$ is a union of triple of type classes, it is necessary to require that when y^n is permuted, \tilde{f} continues to pick up a minimizer inside (or outside) $\Lambda_{tr}^{*,n}(P_{t_D^N})$. For this reason we also impose the additional constraint that, for any permutation σ ,

$$\tilde{f}(\sigma(y^n), t_A^N) = \sigma(\tilde{f}(y^n, t_A^N)). \quad (68)$$

Given that we are not interested in proving the optimality of \tilde{f} , the further constraint in (68) does not impair the value of our analysis. As a matter of fact, we will use \tilde{f} only to bound the false negative error exponent and show that, even if the $HT_{tr,a}^{lr}$ game is less favorable to the attacker than the $HT_{tr,b}^{lr}$ game, the two games have the same indistinguishability region.

By following the same flow of ideas used in Section VI, we consider the set of sequences for which the attacker is able to move y^n within the acceptance region $\Lambda_{tr}^{*,n}(P_{t_D^N})$, i.e.:

$$\tilde{\Gamma}_{tr,a}^n = \{(y^n, t_D^N, t_A^N) : \tilde{f}(y^n, t_A^N) \in \Lambda_{tr}^{*,n}(P_{t_D^N})\}. \quad (69)$$

Thanks to the additional constraint in (68), and by reasoning as in the proof of Property 1 in [19], it is easy to show that $\tilde{\Gamma}_{tr,a}^n$ is a union of triple of type classes, hence permitting us to redefine it in terms of types. Similarly to version b of the game, we find it useful to introduce the following definition:

$$\begin{aligned} \tilde{\Gamma}_{tr,a}^n(P_{t_D^N}, P_{t_A^N}) &= \{P_{y^n} \in \mathcal{P}_n : \forall y^n \in T(P_{y^n}), \\ &\quad (\tilde{f}(y^n, t_A^N), t_D^N) \in \Lambda_{tr}^{*,n}\}. \end{aligned} \quad (70)$$

By using the generalized function h_c instead of h , we can apply the above definition to any pair of pmf's. Specifically, given two pmf's Q and R , we define:

$$\begin{aligned} \tilde{\Gamma}_{tr,a}^n(Q, R) &= \{P_{y^n} \in \mathcal{P}_n : \forall y^n \in T(P_{y^n}), \\ &\quad (\tilde{f}(y^n, R), Q) \in \Lambda_{tr}^{*,n}\}. \end{aligned} \quad (71)$$

It is easy to see that:

$$\begin{aligned} \tilde{\Gamma}_{tr,a}^n(Q, R) &\subseteq \tilde{\Gamma}_{tr,a}^n(Q, Q) \\ \tilde{\Gamma}_{tr,a}^n(Q, Q) &= \Gamma_{tr,b}^n(Q), \end{aligned} \quad (72)$$

since when (and only when) $Q = R$, A performs its attack by using exactly the same acceptance region adopted by D, while in all the other cases he can rely only on an estimate based on its own training sequence. Paralleling the analysis of the $HT_{tr,b}^{lr}$ game, we introduce the set

$$\tilde{\Gamma}_{tr,a}^\infty(Q, R) = cl \left(\bigcup_n \tilde{\Gamma}_{tr,a}^n(Q, R) \right) \quad (73)$$

for which, thanks to (72), we have $\tilde{\Gamma}_{tr,a}^\infty(Q, R) \subseteq \tilde{\Gamma}_{tr,a}^\infty(Q, Q) = \Gamma_{tr,b}^\infty(Q)$.

We are now ready to prove our main result regarding the $HT_{tr,a}^{lr}$ game.

Theorem 6 (Asymptotic Payoff of the $HT_{ir,a}^{lr}$ Game): The error exponent of the payoff associated to the profile $(\Lambda_{ir}^{*,n}(P_{t_D}^N), \tilde{f}(\cdot, t_A^N))$ is lower (res. upper) bounded as follows

$$\tilde{\varepsilon}_{ir,a} \geq \min_{Q,R} \{c[\mathcal{D}(Q||P_X) + \mathcal{D}(R||P_X)] + \min_{P \in \tilde{\Gamma}_{ir,a}^\infty(Q,R)} \mathcal{D}(P||P_Y)\}, \quad (74)$$

$$\tilde{\varepsilon}_{ir,a} \leq \min_Q [2c \cdot \mathcal{D}(Q||P_X) + \min_{P \in \tilde{\Gamma}_{ir,a}^\infty(Q,Q)} \mathcal{D}(P||P_Y)]. \quad (75)$$

Proof: The proof is similar to the proof of Theorem 4, with the noticeable difference that now the lower and upper bounds are different hence preventing us to derive a precise expression for the error exponent. Let us start with the lower bound. By recalling the definition of the false negative error probability, for any n we can write:

$$\begin{aligned} P_{FN} &= \sum_{t_D^N} \sum_{t_A^N} P_X(t_D^N) P_X(t_A^N) \sum_{P \in \tilde{\Gamma}_{ir,a}^n(P_{t_D}^N, P_{t_A}^N)} P_Y(T(P)) \\ &= \sum_{Q \in \mathcal{P}_N} \sum_{R \in \mathcal{P}_N} P_X(T(Q)) P_X(T(R)) \sum_{P \in \tilde{\Gamma}_{ir,a}^n(Q,R)} P_Y(T(P)) \\ &\leq \sum_{Q \in \mathcal{P}_N} \sum_{R \in \mathcal{P}_N} P_X(T(Q)) P_X(T(R)) \\ &\quad \cdot (n+1)^{|\mathcal{X}|} 2^{-n \min_{P \in \tilde{\Gamma}_{ir,a}^n(Q,R)} \mathcal{D}(P||P_Y)} \\ &\leq \sum_{Q \in \mathcal{P}_N} P_X(T(Q)) (n+1)^{|\mathcal{X}|} (N+1)^{|\mathcal{X}|} \\ &\quad \cdot 2^{-n \min_{R \in \mathcal{P}_N} [c\mathcal{D}(R||P_X) + \min_{P \in \tilde{\Gamma}_{ir,a}^n(Q,R)} \mathcal{D}(P||P_Y)]} \\ &\leq (n+1)^{|\mathcal{X}|} (N+1)^{2|\mathcal{X}|} \\ &\quad \cdot 2^{-n \min_{Q,R} [c\mathcal{D}(Q||P_X) + c\mathcal{D}(R||P_X) + \min_{P \in \tilde{\Gamma}_{ir,a}^n(Q,R)} \mathcal{D}(P||P_Y)]}, \end{aligned} \quad (76)$$

where the use of the minimum instead of the infimum is justified by the compactness of the involved sets, and where in the last inequality we replaced the minimization over all Q and R in \mathcal{P}_N , with a minimization over the entire space of pmf's. By taking the logarithm of both sides and by letting n tend to infinity, the lower bound in (74) is proved.

We now turn the attention to the upper bound. To do so, let Q^* be the pmf achieving the minimum in (75). Due to the density of rational numbers within real numbers, we can find two sequences of pmf's Q_n and R_n that tend to Q^* when n tends to infinity, and such that $Q_n \in \mathcal{P}_n$, $R_n \in \mathcal{P}_n$, $\forall n$. By remembering that $N = nc$, we can say that the subsequences $Q_N = Q_{nc}$ and $R_N = R_{nc}$ also tend to Q^* when n (and hence N) tends to infinity. We can, then, use the subsequences Q_N and R_N to write the following chain of inequalities:

$$\begin{aligned} P_{FN} &= \sum_{Q \in \mathcal{P}_N} \sum_{R \in \mathcal{P}_N} P_X(T(Q)) P_X(T(R)) \sum_{P \in \tilde{\Gamma}_{ir,a}^n(Q,R)} P_Y(T(P)) \\ &\geq \sum_{Q,R \in \mathcal{P}_N} P_X(T(Q)) P_X(T(R)) \sum_{P \in \tilde{\Gamma}_{ir,a}^n(Q,R)} \frac{2^{-n\mathcal{D}(P||P_Y)}}{(n+1)^{|\mathcal{X}|}} \end{aligned}$$

$$\begin{aligned} &\stackrel{(a)}{\geq} \sum_{Q,R \in \mathcal{P}_N} P_X(T(Q)) P_X(T(R)) \frac{2^{-n \min_{P \in \tilde{\Gamma}_{ir,a}^n(Q,R)} \mathcal{D}(P||P_Y)}}{(n+1)^{|\mathcal{X}|}} \\ &\geq \sum_{Q,R \in \mathcal{P}_N} P_X(T(Q)) \frac{2^{-n[c\mathcal{D}(R||P_X) + \min_{P \in \tilde{\Gamma}_{ir,a}^n(Q,R)} \mathcal{D}(P||P_Y)]}}{(N+1)^{|\mathcal{X}|} (n+1)^{|\mathcal{X}|}} \\ &\stackrel{(b)}{\geq} \sum_{Q \in \mathcal{P}_N} P_X(T(Q)) \frac{2^{-n[c\mathcal{D}(R_N||P_X) + \min_{P \in \tilde{\Gamma}_{ir,a}^n(Q,R_N)} \mathcal{D}(P||P_Y)]}}{(N+1)^{|\mathcal{X}|} (n+1)^{|\mathcal{X}|}} \\ &\geq \sum_{Q \in \mathcal{P}_N} \frac{2^{-n[c\mathcal{D}(Q||P_X) + c\mathcal{D}(R_N||P_X) + \min_{P \in \tilde{\Gamma}_{ir,a}^n(Q,R_N)} \mathcal{D}(P||P_Y)]}}{(N+1)^{2|\mathcal{X}|} (n+1)^{|\mathcal{X}|}} \\ &\stackrel{(c)}{\geq} \frac{2^{-n[c\mathcal{D}(Q_N||P_X) + c\mathcal{D}(R_N||P_X) + \min_{P \in \tilde{\Gamma}_{ir,a}^n(Q_N,R_N)} \mathcal{D}(P||P_Y)]}}{(N+1)^{2|\mathcal{X}|} (n+1)^{|\mathcal{X}|}}, \end{aligned} \quad (77)$$

where inequalities (a), (b) and (c) have been obtained by replacing the summation with a single element of the sum (two elements of the sequences Q_N and R_N for (b) and (c)), and the others rely on a known lower bound on the probability of a type class ([24, Ch. 12]). By taking the logarithm of each side in (77), we can write:

$$\tilde{\varepsilon}_{ir,a} \leq c\mathcal{D}(Q_N||P_X) + c\mathcal{D}(R_N||P_X) + \min_{P \in \tilde{\Gamma}_{ir,a}^n(Q_N,R_N)} \mathcal{D}(P||P_Y) + \beta_n, \quad (78)$$

with $\beta_n = [2|\mathcal{X}| \log(N+1) + |\mathcal{X}| \log(n+1)]/n$ tending to 0 for $n \rightarrow \infty$. Let then P^* be defined as:

$$P^* = \arg \min_{P \in \tilde{\Gamma}_{ir,a}^\infty(Q^*,Q^*)} \mathcal{D}(P||P_Y). \quad (79)$$

By recalling that both Q_N and R_N tend to Q^* for increasing N , we can invoke Lemma 5 in Appendix C, to build a sequence P_n such that each term of the sequence belongs to $\tilde{\Gamma}_{ir,a}^n(Q_N, R_N)$ and $P_n \rightarrow P^*$, when $n \rightarrow \infty$. By recalling that

$$Q^* = \arg \min_Q [2c \cdot \mathcal{D}(Q||P_X) + \min_{P \in \tilde{\Gamma}_{ir,a}^\infty(Q,Q)} \mathcal{D}(P||P_Y)], \quad (80)$$

and by reasoning as in the proof of Theorem 4 (equations (51) and (52) and discussion therein), we can eventually prove the upper bound (75). ■

Theorem 6 has an important corollary.

Corollary 3 (Indistinguishability Region for $HT_{ir,a}^{lr}$): The false negative error exponent associated to the profile $(\Lambda_{ir}^{*,n}(P_{t_D}^N), \tilde{f}(\cdot, t_A^N))$ is equal to zero if and only if $P_Y \in \tilde{\Gamma}_{ir,a}^\infty(P_X, P_X) = \Gamma_{ir,b}^\infty(P_X)$, and hence the indistinguishability region of the $HT_{ir,a}^{lr}$ game is equal to that of the $HT_{ir,b}^{lr}$ game.

Proof: From the upper bound in Theorem 6, it follows that $\tilde{\varepsilon}_{ir,a} = 0$ if $P_Y \in \tilde{\Gamma}_{ir,a}^\infty(P_X, P_X)$, whereas from the lower bound we see that $\tilde{\varepsilon}_{ir,a} = 0$ implies that $P_Y \in \tilde{\Gamma}_{ir,a}^\infty(P_X, P_X)$. ■

Corollary 3 provides an interesting insight into the achievable performance of the $HT_{ir,a}^{lr}$ game. While, in general, version a of the game is less favorable to the attacker than version b , since in the latter case the attacker knows exactly

the acceptance region adopted by the defender, if the attacker adopts the strategy \tilde{f} , the indistinguishability regions of the two games are the same. Such a strategy, then, is optimal at least as far as the indistinguishability region is concerned. On the other side, there is no guarantee that the attacker can achieve the same payoff as for version b .

A. Training Sequences With Different Length

We conclude this section by discussing briefly the case in which the training sequences t_D^N and t_A^K have different lengths, i.e. $c \neq d$. To simplify the analysis we assume that c is known to the attacker, in this way A knows at least the exact form the h_c function used by D. We focus on the following attacking strategy: use the training sequence t_A^K to estimate $P_{t_D^N}$ and use the estimate to attack the sequence y^n . Specifically, the attacker may use the following estimate of $P_{t_D^N}$:

$$\begin{aligned} \tilde{P}_{t_D^N}(i) &= \frac{1}{N} \left[P_{t_A^K}(i) \cdot N \right] \quad \forall i = 1 \dots |\mathcal{X}| - 1, \\ \tilde{P}_{t_D^N}(|\mathcal{X}|) &= 1 - \sum_{i=1}^{|\mathcal{X}|-1} \tilde{P}_{t_D^N}(i), \end{aligned} \tag{81}$$

to implement the attacking function:

$$\tilde{f}(y^n, t_A^K) = \arg \min_{z^n: d(z^n, y^n) \leq nL} h_c(P_{z^n}, \tilde{P}_{t_D^N}). \tag{82}$$

With the above definitions, we can easily extend the analysis carried out for the case $c = d$ and obtain very similar results. Specifically, the upper bound in Theorem 6 can be rewritten as:

$$\tilde{\epsilon}_{tr,a} \leq \min_Q [(c + d) \cdot \mathcal{D}(Q||P_X) + \min_{P \in \tilde{\Gamma}_{tr,a}^\infty(Q,Q)} \mathcal{D}(P||P_Y)], \tag{83}$$

whose proof is practically identical to the proof of Theorem 6 and is omitted for sake of brevity. By observing that the performance achievable by the defender in version a of the game are at least as good as those achievable in version b , since in the latter case A knows exactly the acceptance region adopted by D and hence his attacks will surely be more effective, equation (83) allows us to conclude that the indistinguishability region is equal to that obtained for the case $c = d$.

We end this section by considering yet another version of the HT_{tr}^{lr} game. In certain cases, in fact, we may assume that D has the possibility to observe the output of the system under H_0 , and hence the output of the source X , for a longer time than A (see [37] for a multimedia forensics scenario in which such an assumption holds quite naturally). In our framework, we can model such a situation by assuming that the sequence t_A^K is a subsequence of t_D^N , leading to the following definition.

Definition 6: The $HT_{tr,sub}^{lr}(S_D, S_A, u)$ game is a zero-sum, strategic, game defined as the $HT_{tr,a}^{lr}$ game with the only difference that $t_A^K = (t_{D,l+1}, t_{D,l+2} \dots t_{D,l+K})$ with l and K known to D.

In some sense, we can say that this new version of the game is halfway between versions a and b . Like in version a , the attacker does not have a perfect knowledge of the training sequence used by the defender and hence he must resort to

an estimate of the true acceptance region. On the other hand, the situation is more favorable to the attacker with respect to version a with $d < c$, since now D knows at least part of the training samples used by D. Given that versions a and b of the game have the same indistinguishability region, we can conclude that the indistinguishability region of this latest version of the game will also be the same.

VIII. CONCLUDING REMARKS

The need to protect the cyberworld that surrounds us has spurred researchers to look for suitable countermeasures against the ever increasing number of attacks that every day are brought against the digital world we live in. In many cases, though, research has focused on specific security threats, each time by developing tailored solutions that can not be easily extended to other scenarios. It is no surprise, then, that similar solutions are re-invented several times, and that the same problems are faced with again and again by ignoring that satisfactory solutions have already been discovered in contiguous fields. Even worse, the lack of a unifying view does not permit to grasp the essence of the addressed problems and work out effective and general solutions to be applied with limited effort to different fields. Times are ripe to develop general tools and models that can be used to analyze very general classes of problems wherein the presence of an adversary aiming at system failure can not be neglected.

As a first attempt in this sense, we have introduced a framework to analyze the achievable performance of binary hypothesis testing in an adversarial setting, i.e. in the presence of an adversary with the explicit goal of degrading the performance of the test. We did so by casting the hypothesis testing problem into a game-theoretic framework. In this way, in fact, we have been able to define rigorously the goals and constraints of the two contenders, namely the analyst, a.k.a. the defender, and the adversary or attacker. More specifically, we introduced several versions of the hypothesis testing game, by paying attention to distinguish between hypothesis testing with known sources and hypothesis testing with training data. Then we focused on hypothesis testing with training data (the HT_{tr} game). From a more technical point of view, we derived the asymptotic equilibrium point of two different versions of the game, and analyzed the achievable payoff at the equilibrium. In addition to shedding a new light on the achievable performance of hypothesis testing in an adversarial environment, the analysis we carried out has the merit to clearly show the potentiality of the use of game-theoretic concepts coupled with tools typical of information theory and statistics.

Several directions for future research can be pointed out. The extension of our analysis to multiple hypothesis testing and classification is one of the most promising research directions, together with the extension to games characterized by more than two players. The analysis of situations in which the players do not have a perfect knowledge of the strategies available to the other players, or even the payoff function, by modeling the problem as a game with imperfect knowledge, is another interesting research direction. The investigation of

sequential games, in which the players move repeatedly each time by exploiting the result of the previous round of the game also offers many interesting hints for a fruitful research.

From a more focused perspective, it would be interesting to study specific instances of the HT_{tr} game, in which the sources belong to a specific class, e.g. binary sources. We predict that in this way simpler expressions of the equilibrium point, the achievable payoff and, most of all, the indistinguishability region could be obtained, thus permitting to get additional insights. Finally, we mention the opportunity of extending the analysis to the case of continuous sources. While the general ideas would remain the same, passing from discrete to continuous sources does not seem to be a trivial step, since our analysis relies heavily on the method of types, whose extension to continuous sources, though possible, comes with a number of additional difficulties.

APPENDIX

A. Proof of Lemma 1

We rewrite (24) by moving all the non zero terms to the left-hand side:

$$n\mathcal{D}(P_{x^n}||P_{r^{n+N}}) + N\mathcal{D}(P_{t^N}||P_{r^{n+N}}) - n\mathcal{D}(P_{x^n}||P_X) - N\mathcal{D}(P_{t^N}||P_X) \leq 0. \quad (A1)$$

By using the definition of empirical KL divergence stated in (1) and grouping the first term with the third and the second with the fourth, the left hand side of (A1) is equivalent to

$$n \sum_{a \in \mathcal{X}} P_{x^n}(a) \log \frac{P_X(a)}{P_{r^{n+N}}(a)} + N \sum_{a \in \mathcal{X}} P_{t^N}(a) \log \frac{P_X(a)}{P_{r^{n+N}}(a)}. \quad (A2)$$

Being r^{n+N} the concatenation of x^n and t^N , we argue that $nP_{x^n}(a) + NP_{t^N}(a) = (n+N)P_{r^{n+N}}(a) \forall a \in \mathcal{X}$, which permits to rewrite the sum in (A2) as follows:

$$(n+N) \sum_{a \in \mathcal{X}} P_{r^{n+N}}(a) \log \frac{P_X(a)}{P_{r^{n+N}}(a)} = -(n+N)\mathcal{D}(P_{r^{n+N}}||P_X). \quad (A3)$$

Hence, the proof of relation (A1) follows from the positivity of the divergence function, which equals zero if and only if $P_X = P_{r^{n+N}}$.

In hindsight, relation (A1) derives from the property that the empirical probability distribution $P_{r^{n+N}}$ maximizes the probability that a source outputs the concatenation of x^n and t^N , i.e., $P_X(r^{n+N}) \leq P_{r^{n+N}}(r^{n+N}) \forall P_X$. To show this, from (A3) we write:

$$\sum_{a \in \mathcal{X}} N_{r^{n+N}}(a) \log \frac{P_X(a)}{P_{r^{n+N}}(a)} \leq 0. \quad (A4)$$

Exploiting the properties of the logarithm, relation (A4) is equivalent to the following

$$\log \prod_{a \in \mathcal{X}} P_X(a)^{N_{r^{n+N}}(a)} \leq \log \prod_{a \in \mathcal{X}} P_{r^{n+N}}(a)^{N_{r^{n+N}}(a)}, \quad (A5)$$

which implies

$$P_X(r^{n+N}) \leq \prod_{a \in \mathcal{X}} P_{r^{n+N}}(a)^{N_{r^{n+N}}(a)} = P_{r^{n+N}}(r^{n+N}). \quad (A6)$$

B. Topology of $\Gamma_{tr,b}^\infty(Q)$

Many of the proofs in the body of the paper relies on the following lemma.

Lemma 4: Let $\{Q_{N(n)}\}$ be a sequence of pmf's such that $Q_{N(n)} \rightarrow Q$ when $n \rightarrow \infty$. Then, for any pmf P^* in $\Gamma_{tr,b}^\infty(Q)$ a sequence P_n exists such that $P_n \in \Gamma_{tr,b}^n(Q_{N(n)})$ for each n and $P_n \rightarrow P^*$.

Proof: To prove the lemma, we will show that for any $\varepsilon > 0$ and n large enough, we can find a pmf $P_n \in \Gamma_{tr,b}^n(Q_{N(n)})$ such that $d(P_n, P^*) < \varepsilon$, where d is a proper distance measure between pmf's. Specifically, we will do so by assuming that $P^* \in \cup_n \Gamma_{tr,b}^n(Q)$. If this is not the case, in fact, by the definition of $\Gamma_{tr,b}^\infty(Q)$, it is possible to find a pmf $P' \in \cup_n \Gamma_{tr,b}^n(Q)$ that is arbitrarily close to P^* , and then a pmf in $\Gamma_{tr,b}^n(Q_{N(n)})$ that is arbitrarily close to P' and hence to P^* . Let then P^* belong to $\Gamma_{tr,b}^m(Q)$ for some m . This means that for any sequence $y^m \in T(P^*)$ a mapping f exists that transforms y^m into a sequence z^m such that

$$h_c(P_{z^m}, Q) = \lambda - \delta_m - \delta \quad (A7)$$

with $\delta_m = |\mathcal{X}|[\log(m+1)(N(m)+1)]/m$ and where δ is a strictly positive quantity. Due to the density of $\cup_n \mathcal{P}_n$ in set of all pmf's, we can find a sequence of pmf's P_n that tends to P^* when n tends to infinity and for which $P_n \in \mathcal{P}_n$ for each n . Let y^n be a sequence in $T(P_n)$. Let $n_f(i \rightarrow j)$ indicate the number of times that the mapping f transforms the i -th symbol of the alphabet into the j -th one. By starting from the mapping f , for each n we build a mapping $v_{(n)}$ for which $n_{v_{(n)}}(i \rightarrow j) = \lfloor n_f(i \rightarrow j) \cdot n/m \rfloor$. When n increases the type of $z^n = v_{(n)}(y^n)$ will approach that of $z^m = f(y^m)$ for any y^m in P^* . By remembering that $Q_{N(n)} \rightarrow Q$ when n tends to infinity, and by exploiting the continuity of the h function, we can write:

$$\begin{aligned} h_c(P_{z^n}, Q_{N(n)}) &= h_c(P_{z^n}, Q) + \beta'_n \\ &= h_c(P_{z^n}, Q) + \beta'_n + \beta''_n \\ &= \lambda - \delta_m - \delta + \beta'_n + \beta''_n, \end{aligned} \quad (A8)$$

where β'_n and β''_n tend to zero when $n \rightarrow \infty$. Given that δ is a fixed and strictly positive number, we conclude that when n is large, P_{z^n} belongs to $\Lambda_{tr}^n(Q_{N(n)})$ and hence P_n belongs to $\Gamma_{tr,b}^n(Q_{N(n)})$, thus completing the proof. ■

C. Topology of $\tilde{\Gamma}_{tr,a}^\infty(Q, Q)$

The analysis of version a of the hypothesis testing game needs that lemma 4 is generalized as follows.

Lemma 5: Let $\{Q_{N(n)}\}$ and $\{R_{N(n)}\}$ be two sequences of pmf's such that $Q_{N(n)} \rightarrow Q$ and $R_{N(n)} \rightarrow Q$ as $n \rightarrow \infty$. Then, for any pmf P^* in $\tilde{\Gamma}_{tr,a}^\infty(Q, Q)$ a sequence $P_n \in \tilde{\Gamma}_{tr,a}^n(Q_{N(n)}, R_{N(n)})$ exists such that $P_n \rightarrow P^*$.

Proof: Given the definition of $\tilde{\Gamma}_{tr,a}^\infty(Q, Q)$ as the closure of $\cup_n \tilde{\Gamma}_{tr,a}^n(Q, Q)$, we proceed as in Lemma 4, and limit our proof to the special case in which $P^* \in \cup_n \tilde{\Gamma}_{tr,a}^n(Q, Q)$. Let then P^* belong to $\tilde{\Gamma}_{tr,a}^m(Q, Q)$ for some m . This means that for any sequence $y^m \in T(P^*)$ the function

$$\tilde{f}(y^m, Q) = \arg \min_{z^m: d(z^m, y^m) \leq mL} h_c(P_{z^m}, Q) \quad (A9)$$

maps y^m into a sequence z^m such that

$$h_c(P_{z^m}, Q) = \lambda - \delta_m - \delta, \quad (\text{A10})$$

with $\delta_m = |\mathcal{X}|[\log(m+1)(N(m)+1)]/m$ and where δ is a strictly positive quantity. Due to the density of rational numbers in the real line, we can find a sequence of pmf's P_n that tends to P^* when n tends to infinity and for which $P_n \in \mathcal{P}_n$ for each n . Let u^n be a sequence in $T(P_n)$. By reasoning as in Lemma 4, we can define a, distortion-limited, mapping v for which $n_v(i \rightarrow j) = \lfloor n_f(i \rightarrow j) \cdot n/m \rfloor$. As we have shown in Lemma 4, when n is large, v brings the sequence u^n into a sequence $w^n \in \Lambda_{tr}^{*,n}$, so that, for any $Q_{N(n)} \rightarrow Q$ and n large enough, we have:

$$h_c(P_{w^n}, Q_{N(n)}) = \lambda - \delta_m - \delta + \beta'_n, \quad (\text{A11})$$

with β'_n approaching zero when n increases. The above property, however, is not enough to ensure that $P_n \in \tilde{\Gamma}_{tr,a}^n(Q_{N(n)}, R_{N(n)})$, since for this to be the case the minimization (A9) has to be carried by using $R_{N(n)}$ instead of $Q_{N(n)}$, and so there is no guarantee that the optimum mapping will be v (or another mapping better than that). Despite this observation, we can exploit the fact that $R_{N(n)}$ tends to Q when n tends to infinity, to show that indeed, for large n , $P_n \in \tilde{\Gamma}_{tr,a}^n(Q_{N(n)}, R_{N(n)})$. Given $u^n \in P_n$, let r^n be defined as follows:

$$r^n = \arg \min_{r^n: d(r^n, u^n) \leq nL} h_c(P_{r^n}, R_{N(n)}). \quad (\text{A12})$$

Since both $Q_{N(n)}$ and $R_{N(n)}$ tend to Q , when n increases they will get arbitrarily close to each other. By exploiting the continuity of the h_c function, we can write the following chain of inequalities:

$$\begin{aligned} h_c(P_{r^n}, Q_{N(n)}) &\stackrel{(a)}{\leq} h_c(P_{r^n}, R_{N(n)}) + \beta''_n \\ &\stackrel{(b)}{\leq} h_c(P_{w^n}, R_{N(n)}) + \beta''_n \\ &\stackrel{(c)}{\leq} h_c(P_{w^n}, Q_{N(n)}) + \beta''_n + \beta'''_n \\ &\stackrel{(d)}{=} \lambda - \delta_m - \delta + \beta'_n + \beta''_n + \beta'''_n \\ &\stackrel{(e)}{\leq} \lambda - \delta_m - \delta + \beta'_n + \beta''_n + \beta'''_n, \end{aligned} \quad (\text{A13})$$

where β'_n, β''_n and β'''_n can be made arbitrarily small by increasing n , and where (a) and (c) derive from the continuity of the h_c function and from the fact that $Q_{N(n)}$ and $R_{N(n)}$ tend to the same limit, (b) is due to the fact that r^n is the solution of the minimization in (A12), (d) derives from (A11), and (e) is due to the fact that for a fixed m when n increases $\delta_n = |\mathcal{X}|[\log(n+1)(N(n)+1)]/n \leq |\mathcal{X}|[\log(m+1)(N(m)+1)]/m = \delta_m$.

Equation (A13) proves that $P_n \in \tilde{\Gamma}_{tr,a}^n(Q_{N(n)}, R_{N(n)})$, thus completing the proof of the lemma. ■

REFERENCES

- [1] M. Barni and F. Pérez-González, "Coping with the enemy: Advances in adversary-aware signal processing," in *Proc. IEEE ICASSP*, Vancouver, BC, Canada, May 2013, pp. 8682–8686.
- [2] E. Delp, N. Memon, and M. Wu, "Special issue on digital forensics," *IEEE Signal Process. Mag.*, vol. 26, no. 2, pp. 38–48, Mar. 2009.
- [3] N. Dalvi, P. Domingos, P. Mausam, S. Sanjay, and D. Verma, "Adversarial classification," in *Proc. 10th ACM SIGKDD Int. Conf. Knowl. Discovery Data Mining*, 2004, pp. 99–108.
- [4] A. K. Jain, A. Ross, and U. Uludag, "Biometric template security: Challenges and solutions," in *Proc. EUISPCO*, 2005, pp. 469–472.
- [5] M. Martínez-Díaz, J. Fierrez-Aguilar, F. Alonso-Fernández, J. Ortega-García, and J. A. Sigüenza, "Hill-climbing and brute-force attacks on biometric systems: A case study in match-on-card fingerprint verification," in *Proc. 40th Annu. IEEE Int. Carnahan Conf. Security Technol.*, Oct. 2006, pp. 151–159.
- [6] I. Cox, M. Miller, and J. Bloom, *Digital Watermarking*. San Mateo, CA, USA: Morgan Kaufmann, 2002.
- [7] J. Fridrich, *Steganography in Digital Media: Principles, Algorithms, and Applications*. Cambridge, U.K.: Cambridge Univ. Press, 2009.
- [8] K. M. C. Tan, K. S. Killourhy, and R. A. Maxion, "Undermining an anomaly-based intrusion detection system using common exploits," in *Proc. RAID*, Zürich, Switzerland, Oct. 2002, pp. 54–73.
- [9] Y. Yang, Y. L. Sun, S. Kay, and Q. Yang, "Defending online reputation systems against collaborative unfair raters through signal modeling and trust," in *Proc. 24th ACM Symp. Appl. Comput.*, Honolulu, HI, USA, Mar. 2009.
- [10] W. Wang, H. Li, Y. Sun, and Z. Han, "Securing collaborative spectrum sensing against untrustworthy secondary users in cognitive radio networks," *EURASIP J. Adv. Signal Process.*, vol. 2010, pp. 695750-1–695750-4, Nov. 2009.
- [11] A. S. Rawat, P. Anand, H. Chen, and P. K. Varshney, "Collaborative spectrum sensing in the presence of Byzantine attacks in cognitive radio networks," *IEEE Trans. Signal Process.*, vol. 59, no. 2, pp. 774–786, Feb. 2011.
- [12] R. Böhme and M. Kirchner, "Counter-forensics: Attacking image forensics," in *Digital Image Forensics*, H. T. Sencar and N. Memon, Eds. Berlin, Germany: Springer-Verlag, 2012.
- [13] C. Cachin, "An information-theoretic model for steganography," in *Proc. 2nd Int. Workshop IH*, vol. 1525, 1998, pp. 306–318.
- [14] M. C. Stamm, W. S. Lin, and K. J. R. Liu, "Forensics vs. anti-forensics: A decision and game theoretic framework," in *Proc. IEEE ICASSP*, Kyoto, Japan, Mar. 2012, pp. 1749–1752.
- [15] A. S. Cohen and A. Lapidoth, "The Gaussian watermarking game," *IEEE Trans. Inf. Theory*, vol. 48, no. 6, pp. 1639–1667, Jun. 2002.
- [16] P. Moulin and M. K. Mihcak, "The parallel-Gaussian watermarking game," *IEEE Trans. Inf. Theory*, vol. 50, no. 2, pp. 272–289, Feb. 2004.
- [17] A. Somekh-Baruch and N. Merhav, "On the capacity game of public watermarking systems," *IEEE Trans. Inf. Theory*, vol. 50, no. 3, pp. 511–524, Mar. 2004.
- [18] N. Merhav and E. Sabbag, "Optimal watermark embedding and detection strategies under limited detection resources," *IEEE Trans. Inf. Theory*, vol. 54, no. 1, pp. 255–274, Jan. 2008.
- [19] M. Barni and B. Tondi, "The source identification game: An information-theoretic perspective," *IEEE Trans. Inf. Forensics Security*, vol. 8, no. 3, pp. 450–463, Mar. 2013.
- [20] D. Bernheim, "Rationalizable strategic behavior," *Econometrica*, vol. 52, no. 4, pp. 1007–1028, 1984.
- [21] D. Pearce, "Rationalizable strategic behavior and the problem of perfection," *Econometrica*, vol. 52, no. 4, pp. 1029–1050, 1984.
- [22] M. Barni and B. Tondi, "Optimum forensic and counter-forensic strategies for source identification with training data," in *Proc. IEEE Int. WIFS*, Tenerife, Spain, Dec. 2012, pp. 199–204.
- [23] M. Gutman, "Asymptotically optimal classification for multiple tests with empirically observed statistics," *IEEE Trans. Inf. Theory*, vol. 35, no. 2, pp. 401–408, Mar. 1989.
- [24] T. M. Cover and J. A. Thomas, *Elements of Information Theory*. New York, NY, USA: Wiley, 1991.
- [25] I. Csiszár and P. Shields, *Information Theory and Statistics: A Tutorial*. Delft, The Netherlands: Now Publishers, 2004.
- [26] I. Csiszár and J. Körner, *Information Theory: Coding Theorems for Discrete Memoryless Systems*, 2nd ed. Cambridge, U.K.: Cambridge Univ. Press, 2011.
- [27] J. Nash, "Equilibrium points in n-person games," *Proc. Nat. Acad. Sci.*, vol. 36, no. 1, pp. 48–49, 1950.
- [28] M. J. Osborne and A. Rubinstein, *A Course in Game Theory*. Cambridge, MA, USA: MIT Press, 1994.
- [29] Y. C. Chen, N. Van Long, and X. Luo, "Iterated strict dominance in general games," *Games Econ. Behavior*, vol. 61, no. 2, pp. 299–315, Nov. 2007.
- [30] S. M. Kay, *Fundamentals of Statistical Signal Processing: Detection Theory*, vol. 2. Englewood Cliffs, NJ, USA: Prentice-Hall, 1998.

- [31] M. Chen, J. Fridrich, M. Goljan, and J. Lukas, "Determining image origin and integrity using sensor noise," *IEEE Trans. Inf. Forensics Security*, vol. 3, no. 1, pp. 74–90, Mar. 2008.
- [32] A. Piva, "An overview on image forensics," *ISRN Signal Process.*, vol. 2013, Article id. 496701, pp. 1–22, 2013, doi: 10.1155/2013/496701.
- [33] M. C. Stamm and K. J. R. Liu, "Blind forensics of contrast enhancement in digital images," in *Proc. 15th IEEE ICIP*, San Diego, CA, USA, Oct. 2008, pp. 3112–3115.
- [34] I. Csiszar, "The method of types [information theory]," *IEEE Trans. Inf. Theory*, vol. 44, no. 6, pp. 2505–2523, Oct. 1998.
- [35] I. Csiszar and P. C. Shields, "Redundancy rates for renewal and other processes," *IEEE Trans. Inf. Theory*, vol. 42, no. 6, pp. 2065–2072, Nov. 1996.
- [36] M. Kendall and S. Stuart, *The Advanced Theory of Statistics*, vol. 2, 4th ed. New York, NY, USA: MacMillan, 1979.
- [37] M. Goljan, J. Fridrich, and M. Chen, "Sensor noise camera identification: Countering counter-forensics," *Proc. SPIE*, vol. 7541, p. 75410S, Jan. 2010.

Mauro Barni (M'92–SM'06–F'12) graduated in Electronic Engineering at the University of Florence in 1991. He received the PhD in Informatics and Telecommunications in October 1995. He has carried out his research activity for over 20 years first at the Department of Electronics and Telecommunication of the University of Florence, then at the Department of Information Engineering and Mathematics of the University of Siena where he works as associate Professor. During the last decade he has been studying the application of image processing techniques to copyright protection and authentication of multimedia, and the possibility of processing signals that has been previously encrypted without decrypting them (digital watermarking, multimedia forensics, signal processing in the encrypted domain). Lately he has been working on theoretical and practical aspects of adversarial signal processing.

He is author/co-author of about 270 papers published in international journals and conference proceedings, and holds four patents in the field of digital watermarking and image authentication. He is coauthor of the book *Watermarking Systems Engineering: Enabling Digital Assets Security and other Applications* (Dekker Inc., 2004). He participated to several National and European research projects on diverse topics, including computer vision, multimedia signal processing, remote sensing, digital watermarking, IPR protection. Currently he is involved in the REWIND project focusing on theoretical and practical aspects of Multimedia Forensics (FET program). He was the funding editor of the *EURASIP Journal on Information Security*. He is currently part of the editorial board of the IEEE SIGNAL PROCESSING MAGAZINE. Prof. Barni has been the chairman of the IEEE Information Forensic and Security Technical Committee (IFS-TC) from 2010 to 2011. He is a member of EURASIP. He was appointed DL of the IEEE SPS for the years 2013–2014.

Benedetta Tondi graduated (*cum laude*) in Electronics and Communications Engineering at the University of Siena, Siena, Italy, in 2012, with a thesis on Adversary-aware source identification in the area of multimedia forensics. She is currently a member of the Visual Information Processing and Protection (VIPP) group in the Department of Information Engineering and Mathematical Sciences, of the University of Siena, where she is pursuing her PhD. Her research interest focuses on the application of information theory and game theory concepts to the adversarial signal processing. Benedetta Tondi is a Student member of the IEEE society.