

Applications of Information Theory: Statistics and Statistical Mechanics

Khizar Qureshi *

Abstract

The method of optimizing entropy is used to (i) conduct Asymptotic Hypothesis Testing and (ii) determine the energy distribution for which Entropy is maximized. This paper focuses on two related applications of information theory: Statistics and Statistical Mechanics.

Introduction

Entropy is one measure of uncertainty within a system, and is often used to describe the disorder of sequences of quantized random variables. However, entropy can also be extended to methods within optimization, in which the disorder of a system of interest may be maximized or minimized. Such methods are prevalent within statistics, the physical sciences, and econometrics.

The concerns of a statistician observing a sequence of outcomes include the validity of an explanatory hypothesis, its degree of significance, and any assumptions underlying the statistical tests. While linear hypothesis testing is often sufficient, larger sequences exhibit large deviations in behavior that should receive separate treatment. Traditional linear hypothesis testing trivially assigns a constant multiple of an explanatory parameter β to an observation when forming a hypothesis ($H_0 = k\beta = 0$). Optimal entropy, in which disorder is locally minimized or maximized, can be used to construct *asymptotic*, non-linear hypothesis tests. Unlike linear hypothesis testing, error probability can be minimized.

Optimizing entropy extends to thermodynamic systems. The *Third Law of Thermodynamics* states that the entropy of a closed system, i.e. one in which no mass or energy is added or removed, must be bounded from below by zero. Achieving a non-entropic system is nearly impossible, except within a perfect crystal lattice. A more probable state is one for which the entropy of a system is maximized, and observations follow a Boltzmann distribution.

In our study of optimal entropy, we will use classical Statistics and Statistical Mechanics as a lens. We will demonstrate the concept of minimal entropy through two statistical tests: the univariate optimality test defined by *Stein's Lemma*, and a multivariate optimality test, as defined by the *Chernoff Bounds*. We will see that there exists a distribution, the Boltzmann distribution, that approaches maximum entropy as temperature goes to infinity. Finally, we will apply the concept of asymptotic hypothesis testing to Statistical Mechanics. In particular, we will test and observe the evolution of error probability with a growing sample size. We will also compare Q-function error probability with that of the Chernoff bound. The remainder of the paper is organized as follows. To better understand the atypicality of sequences, we will study the method of types. We will then learn how such sequences behave through the large deviation theory. The focus of the paper will then shift to hypothesis testing, in which we will develop tools for

*Prof. Peter Shor provided generous amounts of feedback
Department of Mathematics
Massachusetts Institute of Technology
18.424: Seminar in Information Theory

recognizing the asymptotic optimality of entropy. Illustrative examples of this will include Stein's Lemma and Chernoff bounds. For the interest of the physical sciences, we will rigorously derive the Boltzmann distribution, for which entropy is nearly maximized. Finally, we will converge the aforementioned topics through simulations of asymptotic statistical testing.

Hypothesis Testing

Statisticians are often concerned with not just observed data, but the several possible underlying explanations. A few examples include: testing for the effectiveness of a drug, determining whether or not a coin is biased, and the effect of gender on wage growth. We begin with a simple case in which we decide between two hypothesis, each of which is represented by an independent and identical distribution, or i.i.d. Let X_1, X_2, \dots, X_n be i.i.d. $\sim Q(x)$. For an observed outcome, we have two possible explanations:

- $H_1 : Q = P_1$
- $H_2 : Q = P_2$

We now define a general decision function, whose value reflects the acceptance and rejection of the above hypothesis. Namely, for general decision $g(x_1, x_2, \dots, x_n)$, $g(x) = i$ indicates that H_i is accepted. In the binary case, the set A over which $g(x) = i$ is complemented by the set A^c , over which $g(x) \neq i$.

Quite often, statisticians are concerned with accepting incorrect hypotheses and rejecting correct ones. Such occurrences, recognized as Type I/II errors, often occur when sequences exhibit atypicality and large deviating behavior (See appendix). Error probabilities are reflected through the decision function using weights α, β :

$$\begin{aligned}\alpha &= P(g(x) = 2 | H_1 \text{ true}) = P_1^n(A^c) \\ \beta &= P(g(x) = 1 | H_2 \text{ true}) = P_2^n(A)\end{aligned}\tag{1}$$

Notice that the general decision function takes on values contradicting those implied by the conditional hypothesis. The first implies that H_2 was accepted even though H_1 was true, and the second implies that H_1 was accepted even though H_2 was true. Type I (reject true) and Type II (accept false) errors similarly prove detrimental to experiments, and so we wish to minimize probabilities α and β . Minimizing α increases β , and minimizing β increases α . We will now explore methodology to minimize the overall probability of error by optimizing entropy as a weighted sum of α and β .

Stein's Lemma

We first fix either α or β , and manipulate the other to minimize the probability of error.

Theorem 1 (Stein's Lemma). *Let X_1, X_2, \dots, X_n be i.i.d. $\sim Q$. Further, let $D(P_1 P_2)$ represent the Kullback-Leibler distance, or relative entropy between the probability densities. Consider the hypothesis test between two alternatives $Q = P_1$ and $Q = P_2$ where $D(P_1 P_2) < \infty$. Let $A_n \subseteq \mathcal{H}^n$ be an acceptance region for hypothesis 1. Let the probabilities of error be*

$$\begin{aligned}\alpha_n &= P_1^n(A_n^c) \\ \beta_n &= P_2^n(A_n)\end{aligned}\tag{2}$$

and for $0 < \epsilon < \frac{1}{2}$, define

$$\beta_n^\epsilon = \min_{A_n \subseteq X^n} \beta_n\tag{3}$$

Then,

$$\lim_{\epsilon \rightarrow 0} \lim_{n \rightarrow \infty} \frac{1}{n} \log \beta_n^\epsilon = -D(P_1 P_2) \quad (4)$$

Proof. See Appendix □

Thus, no sequence of sets B_n has an exponent better than $D(P_1 P_2)$. But the sequence A_n achieves the exponent $D(P_1 P_2)$. Thus A_n is asymptotically optimal, and the best error exponent is $D(P_1 P_2)$.

Chernoff Bound

Thus far, α and β have been treated separately. The approach underlying Stein's Lemma was to set one error probability to be infinitesimally small, and measure the effect on the resulting probability. We saw that setting $\alpha \leq \epsilon$ achieved $\beta_n = 2^{-nD}$. However, the distribution of error amongst α and β may be highly asymmetrical, in which case univariate optimization may not suffice. We now explore methodology for a bivariate optimization.

An alternative approach is to minimize the weighted sum of α and β . The resulting error exponent is known as the *Chernoff Information*. Consider a distribution of i.i.d. random variables: X_1, X_2, \dots, X_n representative of the decision function. We assign P_1 to Q with probability π_1 and P_2 to Q with probability π_2 . Upholding the definition of α and β , the overall probability of error is

$$P_\epsilon^n = \pi_1 \alpha_n + \pi_2 \beta_n \quad (5)$$

Theorem 2 (Chernoff). *The best achievable exponent in the Bayesian probability of error is D^* , where*

$$D^* = \text{text} \lim_{n \rightarrow \infty} \min_{A_n \subseteq \mathcal{H}^n} -\frac{1}{n} \log P_\epsilon^n = D(P_{\lambda^*} P_1) = D(P_{\lambda^*} P_2) \quad (6)$$

where

$$P_\lambda = \frac{P_1^\lambda(x) P_2^{1-\lambda}(x)}{\sum_{a \in \mathcal{X}} P_1^\lambda(a) P_2^{1-\lambda}(a)} \quad (7)$$

and λ^* the value of λ such that

$$D(P_{\lambda^*} P_1) = D(P_{\lambda^*} P_2). \quad (8)$$

Proof. See Appendix □

Physical Chemistry

Claude Shannon first proposed that the uncertainty due to possible errors in a message could be encapsulated by

$$U(W) = \log W \quad (9)$$

where W is the number of possible ways (state space) of encoding random information. Intuitively, the uncertainty increases with increasing W , and is zero if $W=1$. The concept of entropy provides a deep-rooted link between information theory and statistical mechanics. The state with the least information available, or greatest entropy occurs when the set of all states are equiprobable. This is also the state with maximum uncertainty. An information theoretic perspective dictates that explicit knowledge of various probabilities associated with the system constitutes greater information. Similarly, the thermodynamics of a system of isolated particles indicate that entropy is directly correlated with expected energy level. Below is a molecular orbital diagram that illustrates the possible energy states, all of which depend on the position an electron occupies.




Figure 1: The figure above is a molecular orbital diagram. When two atoms interact, and possibly share their valence (outermost) electrons, the electrons must occupy a particular orbital. Once occupying an orbital, the atoms are able to form a bond. Here, we are not concerned with the type of the bond, but rather, paired occupancy. The empty orbitals are *indistinguishable*, and will be occupied with equal probability.

Entropy may also be observed in macroscopic states. The *second law of thermodynamics* states that in equilibrium, changes in entropy are proportional to changes in system heat per unit temperature. We have $dW = \frac{dQ_{sys}}{T}$. We can better understand this through an illustration. Consider a system of gas particles that may be expanded or compressed. We can study the system under various entropic regimes. The diagram below illustrates how available work decreases (increases) for gaseous expansion (compression) under bivariate states of pressure and volume.

Consider a perfectly structured crystal lattice structure, in which the positions of each contributing molecule is fixed. If we observe such a system, depart, and return after n periods, the position of each molecule within the crystal will have remained the same *almost surely*. If the particles did not displace, then they also carried zero kinetic energy, which is representative of a zero temperature system. This near certainty of a thermodynamic system is an example of an optimization in which entropy is minimized. If instead, the system consisted of a fair coin toss, with no extra information, entropy would be maximized.

The Boltzmann Distribution

The distribution that maximizes the state space for a fixed energy level is the Boltzmann distribution. We will now derive such a distribution, and show that it uniquely maximizes entropy on each energy level. Consider a crystal containing N particles, each of which has available energy levels, ϵ_n . The state space, W , is the number of ways the total energy $E = \sum_n N_n \epsilon_n$ can be distributed amongst the particles in each energy level, across all energy levels, $N = \sum_n N_n$. The expected number of particles in each energy level, N_n , is the product of the probability that a particle is at an energy level, P_n , and the total number of particles, N . The only consideration for such a distribution is the number of particles in each energy levels, not necessarily the amount of energy allocated to each particle. Energy is conserved amongst states and across energy levels. While there exist several ways of assigning the number of particles in

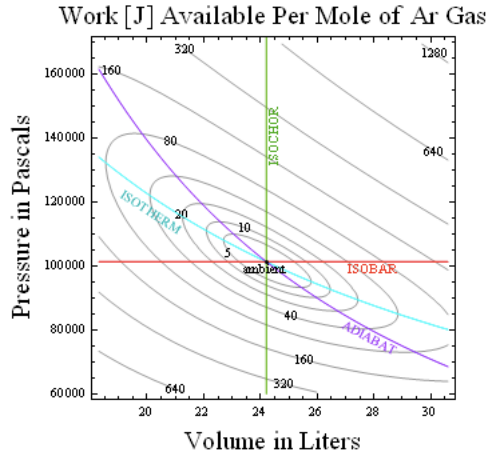


Figure 2: The figure above is a pressure-volume diagram for a system of Argon gas particles. Expanding or compressing a gas requires energy, the extent of which depends on the state of the system. Notice that for an adiabatic system ($dQ_{sys} = 0$), compressing the gas requires the least relative energy. That is, when the change in entropy is minimized, a system can be most naturally expanded/compressed. To reach minimum work available, we move down the gradient of steepest descent until entropy is globally minimized.

each energy level ϵ_n , we wish to find the state with the distribution achievable in the most number of ways for fixed energy levels. Our first constraint is that the total state space W is the sum of individual states occupied, W_i for all possible distributions

$$W = \sum W_i \tag{10}$$

Amongst all possible distributions of particles, there exists one that can be achieved in more ways than any other. A distribution that approaches maximum entropy for fixed energy levels is the Boltzmann distribution.

$$\log W \cong \log W_B \tag{11}$$

To find the most probable distribution that maximizes W , we first note that each particle in the crystal can be distinguished from the others because it occupies a defined position in space. Therefore, such a setting allows us to number the particles $1, 2, \dots, N$. We assume a large N , to maintain consistency with typical non-deficient states. A particular microstate of the crystal will place particle 1 in energy level ϵ_i , particle 2 in energy level ϵ_j , and so on. Initially, we seek the number of W_i microstates in a distribution for which there are N_1 particles in ϵ_1 , N_2 particles in ϵ_2 , and so on. We choose, at random, particles from the crystal, N_i , and assign them to energy levels, ϵ_i . The number of ways this can be done is equal to the number of different orders in which the particles can be chosen from the crystal. The first particle can be chosen from a group of N . With $N - 1$ particles remaining, the second can be chosen in $N - 1$ ways. We see that the number of ways for selecting the first two particles is $N(N - 1)$. Following this procedure, we then we see that the number of ways for selecting the N particles is $N(N - 1)(N - 2)(N - 3) \dots (3)(2)(1)$, or $N!$.

We have over counted the ways of achieving a given distribution, and have assumed that all states are distinguishable. Consider the placement of the first two particles into energy level ϵ_1 . It makes no difference whether the first particle is placed into ϵ_1 prior to or following particle 2. That is, the states

are *indistinguishable*. This relaxes the strictness on order, and so permutations are ignored. Thus, the state space, W_i , for a given distribution, is $N!$ divided by the product of all $N!$

$$W_i = \frac{N!}{\prod_n N_n!} \quad (12)$$

To find the distribution that maximizes W_i , we note Stirling's approximation for $\log N!$

$$\log N! \approx N \log N - N \quad (13)$$

Finding the maximum of W_i is equivalent to finding the maximum of $\log W_i$, so we combine equations (13) and (14), and re-arrange as follows

$$\log W_i \approx \log N! - \sum_n \log N_n! = N \log N - \sum_n N_n \log N_n = \left(\sum_n N_n \right) \log \sum_n N_n - \sum_n N_n \log N_n \quad (14)$$

However, the set of particles is conserved. Moreover, the net energy within the system is conserved. This provides the following two constraints

$$\begin{aligned} \sum_j N_j &= N \\ \sum_j \epsilon_j N_j &= E \end{aligned} \quad (15)$$

We now use Lagrange's method of undetermined multipliers. When W_i is maximized, its differential \log must be zero

$$d \log W_i = \sum_j \frac{\partial \log W_i}{\partial N_j} dN_j = 0 \quad (16)$$

We multiply the constraints on particle count and energy by constants α, β , and then take the differential to obtain

$$\begin{aligned} \alpha \sum_j dN_j &= 0 \\ \beta \sum_j \epsilon_j dN_j &= 0 \end{aligned} \quad (17)$$

Subtracting these two constraints from the log-entropy, we obtain

$$\sum_j \left(\frac{\partial \log W_i}{\partial N_j} - \alpha - \beta \epsilon_j \right) dN_j = 0 \quad (18)$$

Through use of log-properties and algebraic manipulation (See Appendix), the expression above is reduced to

$$\log N_j = \log N - \alpha - \beta \epsilon_j \quad (19)$$

which, after exponentiating both sides is

$$N_j = N e^{-\alpha} e^{-\beta \epsilon_j} \quad (20)$$

The significance of this result is that it shows the occupancy of an energy state ϵ_j is proportional to $e^{-\beta \epsilon_j}$.

Key Result

To account for energy states, thermodynamicists often make use of temperature, an intrinsic quantity. Temperature is equivalent to the average kinetic energy of a system of particles. Because this varies across systems, we normalize. For a temperature T , and the Boltzmann constant, k_B ,

$$\beta = \frac{1}{k_B T}. \quad (21)$$

Inducting on the one particle case, in which, $e^\alpha = \sum_j e^{\frac{-\epsilon_j}{k_B T}}$, and combining the preceding three expressions, providing the desired result

$$P_n = \frac{N_n}{N} = \frac{e^z}{\sum_N e^z} \quad (22)$$

where $z = \frac{-\epsilon_j}{k_B T}$

We have now found the Boltzmann distribution, for which entropy is maximized. The Boltzmann probability above expresses the fraction of particles placed in each quantum state n to maximize entropy W_i of the distribution over each energy level, ϵ_j .

Simulation: Asymptotic Hypothesis Testing

Thus far, we have studied various methods of statistical testing, highlighting the importance of asymptotic tests such as the Chernoff Information Bound. We have also (briefly) explored Statistical Mechanics, in which we show that entropy is maximized for a Boltzmann distribution. We now demonstrate the importance of our learnings through a representative example.

Robust methods of signal interpretation allow for communication, and involve the separation of *signal* and *noise*. A simple signal will follow a Gaussian distribution, for which entropy is maximized. A hypothesis consists of assigning observations as either signals, or as noise. Such hypotheses carry error probabilities, and should be studied with both linear testing, as well as asymptotic testing.

Example: Binary Detection

The classical binary detection problem involves the reception of finite-length signals realized as a random process $r[n], n = 1, 2, \dots$, [3]. The signal can be attributed to either Gaussian white noise $n[i]$ or a deterministic signal $s[i]$. Basic studies involve the interpretation of the signal-to-noise ratio, a measure of quality. Consider a binary detection problem:

$$r = s_i + n, i \in \{1, 2\} \quad (23)$$

- Detections are composed of signals and noise
- n : N -dimensional noise vector
- i.i.d. Gaussian random variables and $\sim (0, 1)$
- $s_1 = (m, m, \dots, m)$
- $s_2 = (0, 0, \dots, 0)$

Evaluate the error probability for both $N = 1$ and $N = 4$ when $m \in \{1, 2, 3, 4, 5, 6\}$.

Traditionally, error probability is evaluated through the Q-function, which represents the probability that a normal random variable will obtain a value larger than x standard deviations above the mean. It can also be thought of as the "tail" probability of the standard normal distribution, and is useful for linear hypothesis testing.

$$Q(x) = \frac{1}{\sqrt{2\pi}} \int_x^\infty \exp\left(-\frac{t^2}{2}\right) dt = \frac{1}{2} \operatorname{erfc}\left(\frac{x}{\sqrt{2}}\right) \quad (24)$$

Given a signal-to-noise Ratio, the Q-function can be used to determine the error probability

$$P^e = Q\left(\frac{\rho}{\sigma}\right) \text{ where } \frac{\rho^2}{\sigma^2} = \frac{Nm^2}{4} \quad (25)$$

We also know that any error probability is bounded from above by the Chernoff Information bound. For the Q-function,

$$Q(x) \leq \exp\left(-\frac{x^2}{2}\right) \quad (26)$$

And so,

$$P^e = Q\left(\frac{\sqrt{Nm}}{2}\right) \leq \exp\left(-\frac{m^2 N}{8}\right) \quad (27)$$

The figure below illustrates the growth of error in both forms of testing.

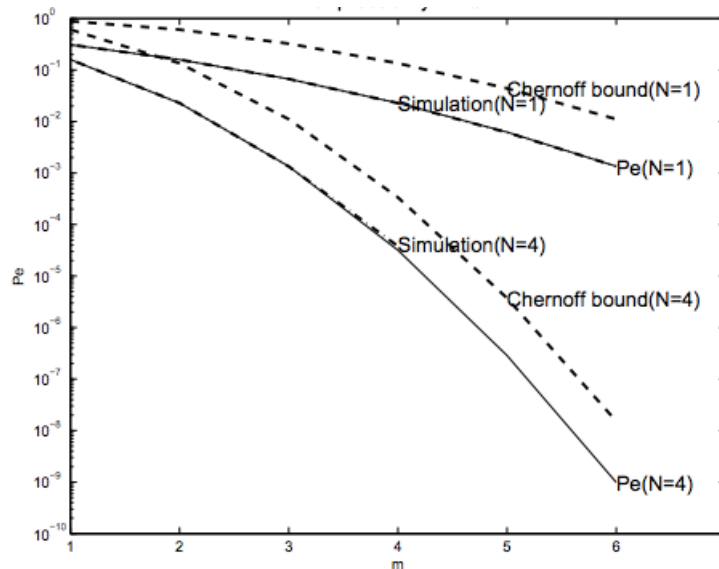


Figure 3: The figure above shows the evolution of error probability with an increasing signal length (m). Error probability decreases as: (i) the number of bits in the signal increases, and (ii) the number of elements in the noise vector increases.

Concluding Remarks

Optimizing entropy demonstrates the applicability of information theory beyond computing. Asymptotic testing captures error probability in atypical sequences, and a Boltzmann distribution of particles approaches maximum entropy, as temperature goes to infinity.

Appendix

Motivation for asymptotic testing arises from atypicality and large deviations in sequences. We *briefly* review this, and encourage the ambitious reader to study further.

The Method of Types

The Asymptotic Equipartition Property formalizes that although there exist several possible outcomes of a stochastic process, there exists a set from which sequences are *typical*, or most frequently observed. The centric approach underlying the AEP involves defining an almost sure convergence in probability between the expectation of a sequence to its entropy. Similarly, the Method of Types defines strong bounds on the number of sequences of a particular distribution, as well as the probability of each such sequence being observed.

Large Deviation Theory

Recall that *type* of a sequence $x_i^n \in A^n$ is representative of its empirical distribution $\hat{P} = \hat{P}_{x_i^n}$ where:

$$\hat{P}(a) = \frac{|\{i : x_i = a\}|}{n}, a \in A. \quad (28)$$

A distribution P on A is called an *n-type* if it is the type of some $x_i^n \in A^n$. The set of all $x_i^n \in A^n$ of type P is called the *type class of the n-type P* and is denoted by \mathcal{T}_P^n .

Lemma 3. *The number of possible n-types is*

$$\binom{n + |A| - 1}{|A| - 1} \quad (29)$$

Proof.

$$\mathcal{T}(P) = \{x \in \mathcal{X}^n : P_x = P\} \quad (30)$$

The combinatoric cardinality of $\mathcal{T}(P)$ provides the result □

Lemma 4. *For any n-type P,*

$$\binom{n + |A| - 1}{|A| - 1}^{-1} 2^{nH(P)} \leq |\mathcal{T}_P^n| \leq 2^{nH(P)} \quad (31)$$

Proof. First, we prove the upper bound using $P(\mathcal{T}(P)) \leq 1$.

$$1 \geq P^n(\mathcal{T}(P)) = \sum_{x \in \mathcal{T}(P)} P^n(x) = \sum_{x \in \mathcal{T}(P)} 2^{-nH(P)} = T(P)2^{-nH(P)} \quad (32)$$

Consequently, $T(P) \leq 2^{nH(P)}$. For the lower bound, using the fact that $T(P)$ has the highest probability amongst all type classes in \mathcal{P} , we can bound the ratio of probabilities

$$\frac{P^n(T(P))}{P^n(T(\hat{P}))} = \frac{T(P) \prod_{a \in \mathcal{X}} P(a)^{nP(a)}}{T(\hat{P}) \prod_{a \in \mathcal{X}} P(a)^{n\hat{P}(a)}} = \prod_{a \in \mathcal{X}} \frac{(n\hat{P}(a))!}{(nP(a))!} P(a)^{n(P(a) - \hat{P}(a))} \quad (33)$$

Using the identity $\frac{m!}{n!} \geq n^{m-n}$, we see

$$\frac{P^n(T(P))}{P^n(T(\hat{P}))} \geq \prod_{a \in \mathcal{X}} n^{n(P(a) - \hat{P}(a))} = n^{n(1-1)} = 1 \quad (34)$$

So $P^n(T(P)) \geq P^n(T(\hat{P}))$. The lower bound can now be found as

$$\begin{aligned} 1 &= \sum_{Q \in \mathcal{P}_n} P^n(T(Q)) \leq \sum_{Q \in \mathcal{P}_n} P^n(T(Q)) \\ &= \max_Q P^n(T(Q)) \\ &= \sum_{Q \in \mathcal{P}_n} P^n(T(Q)) \\ &\leq (n+1)^{\mathcal{X}} P^n(T(P)) = (n+1)^{\mathcal{X}} \sum_{x \in T(P)} P^n(x) \\ &= (n+1)^{\mathcal{X}} \sum_{x \in T(P)} 2^{-nH(P)} \\ &= (n+1)^{\mathcal{X}} T(P) 2^{-nH(P)} \end{aligned} \quad (35)$$

□

To connect the theory of types with general probability theory, we must develop a sense of relative entropy. For any distribution P on A , let P^n denote the distribution of n independent drawings from P , that is, $P^n(x_1^n) = \prod_{i=1}^n P(x_i)$, $x_1^n \in A^n$.

Lemma 5. For any distribution P on A and any n -type Q

$$\begin{aligned} \frac{P^n(x_1^n)}{Q^n(x_1^n)} &= 2^{-nD(QP)}, \text{ if } x_1^n \in \mathcal{T}_Q^n \\ \binom{n+|A|-1}{|A|-1}^{-1} 2^{-nD(QP)} &\leq P(\mathcal{T}_P^n) \leq 2^{-nD(QP)} \end{aligned} \quad (36)$$

Proof. For probability $P \in \mathcal{P}_n$, distribution Q , the probability of type class $T(P)$ under Q^n is $2^{-nD(PQ)}$. We see

$$\begin{aligned} Q^n(T(P)) &= \sum_{x \in T(P)} Q^n(x) \\ &= \sum_{x \in T(P)} 2^{-n(D(PQ)+H(P))} \\ &= T(P) 2^{-n(D(PQ)+H(P))} \end{aligned} \quad (37)$$

Replacing $T(P)$ with the result from Lemma 3, we see the result. □

Corollary 6. Let \hat{P}_n denote the empirical distribution (type) of a random sample of size n drawn from P . Then

$$P \left(D \left(\hat{P}_n P \right) \geq \delta \right) \leq \binom{n + |A| - 1}{|A| - 1} 2^{-n\delta}, \forall \delta > 0 \quad (38)$$

Proof. Given an $\epsilon > 0$, we can define a typical set \mathcal{T}_Q^ϵ of sequences for the distribution Q as $\mathcal{T}_Q^\epsilon = \{x^n : D(P_{x^n} Q) \leq \epsilon\}$. Then the probability of an atypical sequence is

$$\begin{aligned} 1 - Q^n(\mathcal{T}_Q^\epsilon) &= \sum_{P: D(PQ) \geq \epsilon} Q^n(T(P)) \\ &\leq \sum_{P: D(PQ) \geq \epsilon} 2^{-nD(PQ)} \\ &\leq \sum_{P: D(PQ) \geq \epsilon} 2^{-n\delta} \\ &\leq \binom{n + |A| - 1}{|A| - 1} 2^{-n\delta}, \forall \delta \geq 0 \end{aligned} \quad (39)$$

□

Theorem 7. *Sanov's Theorem* Let Π be a set of distributions on A whose closure is equal to the closure of its interior. Then for the empirical distribution of a sample from a strictly positive distribution P on A ,

$$-\frac{1}{n} \log P \left(\hat{P} \in \Pi \right) \rightarrow D(\Pi P) \quad (40)$$

Proof. Sanov's Theorem Let \mathcal{P}_n be the set of possible n -types and let $\Pi_n = \Pi \cap \mathcal{P}_n$. The previous lemma implies that

$$\begin{aligned} \text{Prob} \left(\hat{P}_n \in \Pi_n \right) &= P^n \left(\cup_{Q \in \Pi_n} \mathcal{T}_Q^n \right) \text{ is upper bounded by} \\ &\binom{n + |A| - 1}{|A| - 1} 2^{-nD(\Pi_n P)} \text{ and lower bounded by} \\ &\binom{n + |A| - 1}{|A| - 1}^{-1} 2^{-nD(\Pi_n P)} \end{aligned} \quad (41)$$

Since $D(QP)$ is continuous in Q , the hypothesis on Π implies that $D(\Pi_n P)$ is arbitrarily close to $D(\Pi P)$ if n is large. □

Proofs

Stein's Lemma. To prove the theorem, we construct a sequence of acceptance regions $A_n \subseteq X^n$ such that $A_n < \epsilon$ and $\beta_n = 2^{-nD(P_1 P_2)}$. We then show that no other sequence of tests has an asymptotically better exponent.

First, we define

$$A_n = \left\{ x \in X^n : 2^{+n(D(P_1 P_2) - \delta)} \leq \frac{P_1 x}{P_2 x} \leq 2^{+n(D(P_1 P_2) + \delta)} \right\} \quad (42)$$

Then, we have the following properties:

1. $P_1^n(A_n) \rightarrow 1$. This follows from:

$$P_1^n(A_n) = P_1^n \left(\frac{1}{n} \sum_{i=1}^n \log \frac{P_1(X_i)}{P_2(X_i)} \in (D(P_1 P_2) - \delta_1 D(P_1 P_2) + \delta) \right) \quad (43)$$

by the strong LLN, since $D(P_1 P_2) = E_{P_1} \left(\log \frac{P_1(X)}{P_2(X)} \right)$. Hence, for sufficiently large n , $A_n < \epsilon$.

2. $P_2^n(A_n) \leq 2^{+n(D(P_1 P_2) - \delta)}$. Using the definition of A_n , we have

$$\begin{aligned} P_2^n(A_n) &= \sum_{A_n} P_2(x) \leq \sum_{A_n} P_1(x) 2^{+n(D(P_1 P_2) - \delta)} \\ &= 2^{+n(D(P_1 P_2) - \delta)} \sum_{A_n} P_1(x) \\ &= 2^{+n(D(P_1 P_2) - \delta)} (1 - \alpha_n). \end{aligned} \quad (44)$$

Similarly, $P_2^n(A_n) \geq 2^{+n(D(P_1 P_2) + \delta)} (1 - \alpha_n)$. And so, $\frac{1}{n} \log \beta_n \leq -D(P_1 P_2) + \delta + \frac{\log(1 - \alpha_n)}{n}$, and $\frac{1}{n} \log \beta_n \geq -D(P_1 P_2) - \delta + \frac{\log(1 - \alpha_n)}{n}$, hence $\lim_{n \rightarrow \infty} \frac{1}{n} \log \beta_n = -D(P_1 P_2)$. □

Chernoff Bound. The optimum hypothesis test is a likelihood ratio test, which follows the form:

$$D(P_{X^n} P_2) - D(P_{X^n} P_1) > T \quad (45)$$

The test divides the probability simplex into regions corresponding to hypothesis 1 and hypothesis 2, respectively. This is illustrated below: Let A be the set of types associated with hypothesis 1. From the




Figure 4: The figure above shows the probability simplex and the Chernoff bound. Notice that for error probabilities P_1 and P_2 , there exists an optimal error probability P_λ . This is determined as a weighted argmin.

preceding discussions, it follows that the closest point in the set A^c to P_1 is on the boundary of A , and is of the form given by [8]. Then, it is clear that P_λ is the distribution in A that is closest to P_2 . It is also the distribution in A^c that is closest to P_1 . By Sanov's theorem, we can calculate the associated probabilities of error:

$$\begin{aligned} \alpha_n &= P_1^n(A^c) = 2^{-nD(P_\lambda^* P_1)} \\ \beta_n &= P_2^n(A) = 2^{-nD(P_\lambda^* P_2)} \end{aligned} \quad (46)$$

In the Bayesian case, the overall probability of error is the weighted sum of the two probabilities of error,

$$P_e = \pi_1 2^{-nD(P_\lambda^* P_1)} + \pi_2 2^{-nD(P_\lambda^* P_2)} = 2^{-n \min\{-D(P_\lambda^* P_1), -D(P_\lambda^* P_2)\}} \quad (47)$$

since the exponential rate is determined by the worst exponent. Since $D(P_\lambda P_1)$ increases with λ and $D(P_\lambda P_2)$ decreases with λ , the maximum value of the minimum of $\{D(P_\lambda P_1), D(P_\lambda P_2)\}$ is attained when they are equal. We choose λ so that

$$D(P_\lambda P_1) = D(P_\lambda P_2) = C(P_1, P_2) \tag{48}$$

Thus $C(P_1, P_2)$ is the highest achievable exponent for the probability of error, and is called the Chernoff information.

The closest point in the set A^c to P_1 is on the boundary of A , and is of the form given by $[\cdot]$. Then from the previous discussion, it is clear that P_λ is the distribution in A that is closest to P_2 ; it is also the distribution in A^c that is closest to P_1 . By Sanov's theorem, we can calculate the associated probabilities of error:

$$\begin{aligned} \alpha_n &= P_1^n(A^c) = 2^{-nD(P_\lambda^* P_1)} \\ \beta_n &= P_2^n(A^c) = 2^{-nD(P_\lambda^* P_2)} \end{aligned} \tag{49}$$

In the Bayesian case, the overall probability of error is the weighted sum of the individual two probabilities of error,

$$P_e = \pi_1 2^{-nD(P_\lambda P_1)} + \pi_2 2^{-nD(P_\lambda P_2)} = 2^{-n \min(D(P_\lambda P_1), D(P_\lambda P_2))} \tag{50}$$

since the exponential rate is determined by the worst exponent. Since $D(P_\lambda P_1) = D(P_\lambda P_1)$ increases with λ and $D(P_\lambda P_2)$ decreases with λ , the maximum value of the minimum of $D(P_\lambda P_1), D(P_\lambda P_2)$ is attained when they are equal.

This is illustrated below:




Figure 5: The figure above shows the relative entropy for each error probability as a function of λ .

We choose λ so that

$$D(P_\lambda P_1) = D(P_\lambda P_2) = C(P_1, P_2) \tag{51}$$

Thus $C(P_1, P_2)$ is the highest achievable exponent for the probability of error, and is called the Chernoff information. □

References

- [1] T.M. Cover, J.A. Thomas, *Elements of Information Theory*, Wiley, 2nd edition, 2006.
- [2] D. Eisenberg, D. Crothers, *Physical Chemistry with Applications to the Life Sciences*, Benjamin Cummings, 3rd edition, 1979.
- [3] P. Gopych, *Sensitivity and Bias within the Binary Signal Detection Theory, BSDT*, Information Theories & Applications, Vol. 11, 2004.