One-Shot Information Hiding

Yanxiao Liu and Cheuk Ting Li

The Chinese University of Hong Kong

Overview

We consider the information hiding problem (Moulin and O'Sullivan (2003)):

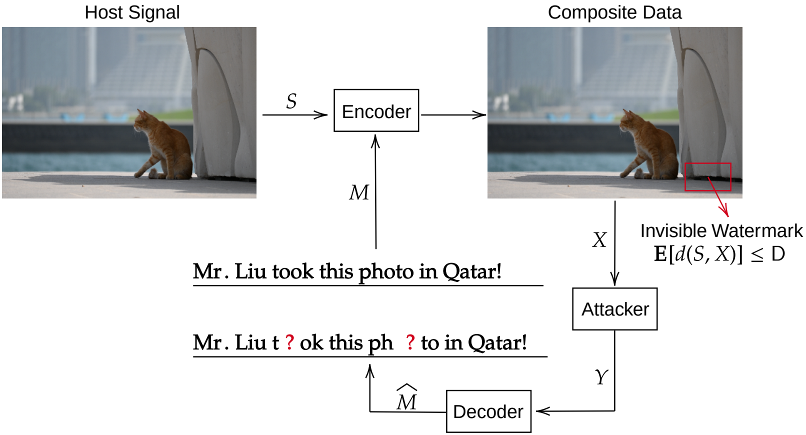

Moulin, P., & O'Sullivan, J. A. (2003). Information-theoretic analysis of information hiding. IEEE Transactions on information theory, 49(3), 563-593.

Background: Information Hiding

Information Hiding and Applications

- Wide range of applications:
  1. Watermarking: protect personal identification contained in messages;
  2. Fingerprinting: identify a unique user even if users collude;
  3. Steganography and cryptography;
- Copyright protection in modern scenarios:
  1. Machine learning tools (Midjourney, ChatGPT, etc) can be possibly trained on public data without obtaining permissions from the authors.
  2. Watermarking tools are important, e.g., Ji et al. (2024).
  3. A better understanding on the fundamental limits of information hiding could help design practical watermarking techniques.

Moulin, P., & O'Sullivan, J. A. (2003). Information-theoretic analysis of information hiding. IEEE Transactions on information theory, 49(3), 563-593.

Ji, Z., Hu, Q., Zheng, Y., Xiang, L., & Wang, X. (2024). A Principled Approach to Natural Language Watermarking. In ACM Multimedia 2024.

Overview

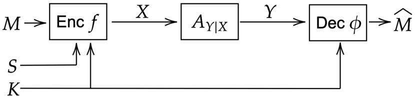

Overview

1. One-shot analysis of the information hiding problem.
2. Game-theoretic formulation:
   1. Team A: an encoder (information hider) and a decoder, trying to embed a message into a host signal and reconstruct it;
   2. Team B: an attacker (noisy channel) trying to remove the hidden information.
3. **Our Contributions**: one-shot achievability results that:
   1. apply to any host distribution, and any class of attack channels;
   2. do not assume the decoder know the attacker's choice, similar to Somekh-Baruch and Merhav (2004);
   3. recover the asymptotic hiding capacity by Moulin and O'Sullivan (2003), hence provide a simple alternative proof.

Merhav, N. (2004). On the capacity game of public watermarking systems. IEEE Transactions on Information Theory, 50(3), 511-524.

Moulin, P., & O'Sullivan, J. A. (2003). Information-theoretic analysis of information hiding. IEEE Transactions on information theory, 49(3), 563-593.

Background: One-Shot Information Theory

One-Shot Information Theory

Each source and channel is only used **once**, i.e., $n = 1$ (discussed in Feinstein (1954); Shannon (1957); Yassaee et al. (2013); Li and Anantharam (2021)).

1. Sources and channels are **arbitrary**: no need to be memoryless or ergodic.
2. Goal: obtain one-shot results that can recover existing (first-order and second-order) **asymptotic** results when applied to memoryless sources and channels and also **finite blocklength** results like Polyanskiy et al. (2010) and Kostina and Verdú (2012).

Feinstein, A. (1954). A new basic theorem of information theory.

Shannon, C. E. (1957). Certain results in coding theory for noisy channels. Information and control, 1(1), 6-25.

Yassaee, M. H., Aref, M. R., & Gohari, A. (2013, July). Non-asymptotic output statistics of random binning and its applications. In 2013 IEEE International Symposium on Information Theory (pp. 1849-1853). IEEE.

Li, C. T., & Anantharam, V. (2021). A unified framework for one-shot achievability via the Poisson matching lemma. IEEE Transactions on Information Theory, 67(5), 2624-2651.

Polyanskiy, Y., Poor, H. V., & Verdú, S. (2010). Channel coding rate in the finite blocklength regime. IEEE Transactions on Information Theory, 56(5), 2307-2359.

Kostina, V., & Verdú, S. (2012). Fixed-length lossy compression in the finite blocklength regime. IEEE Transactions on Information Theory, 58(6), 3309-3338.

One-Shot Information Hiding

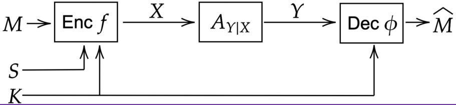

Problem Formulation

- A message M is uniformly chosen from the set $[1 : \mathsf{L}]$;
- Common randomness $K \in \mathcal{K}$ available to the encoder-decoder team;
- **Encoder**: hides M into a host signal $S \in \mathcal{S}$ and $X = f(S, K, M)$.
  1. S, K can be correlated: $S, K \sim P_{S,K}$.
  2. X should be close to S and $d_1(S, X)$ is small with $d_1 : \mathcal{S} \times \mathcal{X} \to [0, \infty)$.
- **Attacker**: chooses an attack channel $A_{Y|X} \in \mathcal{A}$ and tries to destroy M.
  1. We assume the attacker knows the distributions (but not the values) of S, M, K, and the code that the encoder-decoder team uses.
- **Decoder**: observes Y, K and recovers M by $\hat{M} = \phi(K, Y)$
  1. Decoder should be uninformed of the attacker's strategy in one-shot case.
  2. We bound the following worst case failure probability:

$$P_e := \sup_{A_{Y|X} \in \mathcal{A}} \mathbf{P}\Big(d_1(S, X) > \mathrm{D}_1 \ \text{ OR } \ M \neq \hat{M}\Big).$$

One-Shot Information Hiding

Comparison with literature

- We drop several assumptions in Moulin and O'Sullivan (2003):
  1. We do not assume the attack channels are memoryless.
  2. Similar to Somekh-Baruch and Merhav (2004), we do not assume the decoder completely knows the attacker.
  3. In Somekh-Baruch and Merhav (2004):
     1. an asymptotic achievable rate is expressed as the limit of a sequence of single-letter expressions;
     2. K is a shared key of unlimited size independent of M, S that can be chosen as a part of the coding scheme, but our K is a given side information and cannot be changed.
- Our code accounts for all possible attack channels, thus can be regarded as a combination of Gelfand-Pinsker coding and compound channel.
- Techniques in literature (usually based on Gelfand-Pinsker coding) are inapplicable in the one-shot case.

Moulin, P., & O'Sullivan, J. A. (2003). Information-theoretic analysis of information hiding. IEEE Transactions on information theory

Somekh-Baruch, A., & Merhav, N. (2004). On the capacity game of public watermarking systems. IEEE Transactions on Information Theory, 50(3), 511-524.

Technique: Poisson Functional Representation

Poisson Functional Representation (PFR)

- Fix a distribution $\bar{P}$ over $\mathcal{U}$ and a Poisson process $(T_i)_{i=1,2,\ldots}$ of rate 1.
- Let $(\bar{U}_i)_{i=1,2,\ldots}$ be an independent i.i.d. sequence with distribution $\bar{P}$.
- The "marked" Poisson process $(\bar{U}_i, T_i)_i$ supports a "query operation" given by the PFR, where one inputs P, and gets a sample $\tilde{U}_P \sim P$.
- Poisson Functional Representation:

$$\tilde{U}_P := \bar{U}_K$$

  where $K := \arg\min_i T_i \cdot \left(\frac{dP}{d\bar{P}}(\bar{U}_i)\right)^{-1}$.
- Various applications: minimax learning (Li et al. (2020)), neural network compression (Lei et al. (2022)), differential privacy (Liu et al. (2024)) etc.

Li, C. T., & El Gamal, A. (2018). Strong functional representation lemma and applications to coding theorems. IEEE Transactions on Information Theory.

Li, C. T., Wu, X., Özgür, A., & El Gamal, A. (2020). Minimax learning for distributed inference. IEEE Transactions on Information Theory, 66(12), 7929-7938.

Lei, E., Hassani, H., & Bidokhti, S. S. (2022). Neural estimation of the rate-distortion function with applications to operational source coding. Journal on Selected Areas in Information Theory.

Liu, Y., Chen, W. N., Özgür, A., & Li, C. T. (2024). Universal Exact Compression of Differentially Private Mechanisms. Advances in Neural Information Processing Systems, 37.

Technique: Poisson Matching Lemma

Poisson Matching Lemma

- In communication settings, e.g., Li and Anantharam (2021); Liu and Li (2024),
  1. Encoder queries the process using the prior distribution of the signal to obtain the codeword;
  2. Decoder queries using the posterior distribution of the signal given the noisy observation to obtain the reconstruction.
- One wishes to bound the error probability (the probability of mismatch between the Poisson functional representations applied on different distributions).
- **Poisson matching lemma**: For two distributions $P, Q \ll \bar{P}$, almost surely, we have:

$$\mathbf{P}\left[\tilde{U}_Q \neq \tilde{U}_P \middle| \tilde{U}_P\right] \leq 1 - \left(1 + \frac{dP}{dQ}(\tilde{U}_P)\right)^{-1}.$$

Li, Cheuk Ting, and Venkat Anantharam. "A unified framework for one-shot achievability via the Poisson matching lemma." IEEE Transactions on Information Theory 67, no. 5 (2021): 2624-2651.

Liu, Y., & Li, C. T. (2024). One-shot coding over general noisy networks. In 2024 IEEE International Symposium on Information Theory (ISIT). IEEE.

Technique: ϵ-covering number

ϵ-covering number

- Since the encoder-decoder team accounts for all possible attack channels in $\mathcal{A}$, it suffers a penalty depending on the "size" of $\mathcal{A}$.
- Though the cardinality of $\mathcal{A}$ could be infinite, we can often find a finite subset $\tilde{A}$ such that every $A \in \mathcal{A}$ is close enough to some $\tilde{A} \subseteq \mathcal{A}$.
- This notion of size is captured by the ϵ-covering number, also appeared in Moulin and O'Sullivan (2003); Blackwell et al. (1959).

Definition

Given a set of channels $\mathcal{A}$ from $\mathcal{X}$ to $\mathcal{Y}$, its **ϵ-covering number** is defined as

$$N_\epsilon(\mathcal{A}) := \min\left\{|\tilde{\mathcal{A}}| : \tilde{\mathcal{A}} \subseteq \mathcal{A},\ \sup_{A \in \mathcal{A}} \min_{\tilde{A} \in \tilde{\mathcal{A}}} \sup_{x \in \mathcal{X}} \|A_{Y|X}(\cdot|x) - \tilde{A}_{Y|X}(\cdot|x)\|_{\mathrm{TV}} \le \epsilon\right\},$$

where $\|A_{Y|X}(\cdot|x) - \tilde{A}_{Y|X}(\cdot|x)\|_{\mathrm{TV}} \in [0,1]$ denotes the TV distance between $A_{Y|X}(\cdot|x)$ and $\tilde{A}_{Y|X}(\cdot|x)$.

Moulin, P., & O'Sullivan, J. A. (2003). Information-theoretic analysis of information hiding. IEEE Transactions on information theory, 49(3), 563-593.

Blackwell, D., Breiman, L., & Thomasian, A. J. (1959). The capacity of a class of channels. The Annals of Mathematical Statistics, 1229-1241.

One-Shot Achievability

Theorem

Fix any $P_{U,X|S,K}$ and channel $\hat{A}_{Y|X}$. For any $\epsilon \geq 0$, there exists an information hiding scheme satisfying

$$P_e \leq N_\epsilon(\mathcal{A}) \cdot \sup_{A_{Y|X}\in\mathcal{A}} \mathbf{E}_{Y|X\sim A_{Y|X}}\left[1 - \mathbf{1}\{d_1(S,X) \leq \mathrm{D}_1\} \cdot \left(1 + \mathrm{L}2^{-\hat{\iota}(U;Y|K)+\iota(U;S|K)}\right)^{-1}\right] + \epsilon,$$

where we assume $(S,K,U,X,Y) \sim P_{S,K}P_{U,X|S,K}A_{Y|X}$ in the expectation, and $\hat{\iota}(U;Y|K)$ is the information density computed by the joint distribution $P_{S,K}P_{U,X|S,K}\hat{A}_{Y|X}$ (instead of $A_{Y|X}$), assuming that $\iota(U;S|K), \hat{\iota}(U;Y|K)$ are almost surely finite for every $A_{Y|X} \in \mathcal{A}$.

One-Shot Achievability: Proof

Proof (part 1/3)

- We design the decoder assuming the attack channel is fixed to $\hat{A}_{Y|X}$, and hope it works for every attack channel $A_{Y|X}$.
- Codebook: $\mathcal{C} := ((\bar{U}_i, \bar{M}_i), T_i)_i$ where $(T_i)_i$ is a Poisson process, $\bar{U}_i \overset{iid}{\sim} P_U$, and $\bar{M}_i \overset{iid}{\sim} P_M = \mathrm{Unif}[1:\mathrm{L}]$. It will be fixed later.
- Encoding and decoding utilize the Poisson Functional Representation:
  1. Encoder: $U = \tilde{U}_{P_{U|S,K}(\cdot|S,K)\times\delta_M}$ and sends $X|(S,K,U) \sim P_{X|S,K,U}$;
  2. Decoder: observes Y, K and outputs $\hat{M} = \tilde{M}_{\hat{P}_{U|Y,K}(\cdot|Y,K)\times P_M}$, where $\hat{P}_{U|Y,K}$ is computed by the joint distribution $P_{S,K}P_{U,X|S,K}\hat{A}_{Y|X}$.

One-Shot Achievability: Proof

Proof (part 2/3)

When the attack channel is $A_{Y|X} \in \mathcal{A}$, the error probability is

$$
\begin{aligned}
P_e(A) &:= 1 - \mathbf{P}_{Y|X\sim A_{Y|X}}\Big(d_1(S,X) \le \mathrm{D}_1 \ \text{ AND } \ M = \hat{M}\Big) \\
&= \mathbf{E}\Big[1 - \mathbf{1}\{d_1(S,X) \le \mathrm{D}_1\} \cdot \mathbf{P}\Big(M = \hat{M} | M, S, U, Y, K\Big)\Big] \\
&\le \mathbf{E}\Big[1 - \mathbf{1}\{d_1(S,X) \le \mathrm{D}_1\} \cdot \mathbf{P}\Big((U,M) = (\tilde{U}, \tilde{M})_{\hat{P}_{U|Y,K}(\cdot|Y,K)\times P_M} | M, S, U, Y, K\Big)\Big] \\
&\overset{(a)}{\le} \mathbf{E}\left[1 - \mathbf{1}\{d_1(S,X) \le \mathrm{D}_1\} \cdot \left(1 + \frac{\mathrm{d}P_{U|S,K}(\cdot|S,K) \times \delta_M}{\mathrm{d}\hat{P}_{U|Y,K}(\cdot|Y,K) \times P_M}(U,M)\right)^{-1}\right] \\
&= \mathbf{E}\left[1 - \mathbf{1}\{d_1(S,X) \le \mathrm{D}_1\}\left(1 + \mathrm{L}2^{-\hat{\iota}(U;Y|K)+\iota(U;S|K)}\right)^{-1}\right] \\
&\le \sup_{A_{Y|X}\in\mathcal{A}} \mathbf{E}_{Y|X\sim A_{Y|X}}\left[1 - \mathbf{1}\{d_1(S,X) \le \mathrm{D}_1\} \cdot \left(1 + \mathrm{L}2^{-\hat{\iota}(U;Y|K)+\iota(U;S|K)}\right)^{-1}\right] \\
&=: \overline{P_e},
\end{aligned}
$$

where (a) is by the Poisson matching lemma.

One-Shot Achievability: Proof

Proof (part 3/3)

- The only common randomness between the encoder and the decoder is K, which we cannot control. We need to fix the codebook.
- Let $\tilde{\mathcal{A}} \subseteq \mathcal{A}$ attain the minimum in $N_\epsilon(\mathcal{A})$ and write $P_e(A) = \mathbf{E}_{\mathcal{C}}[P_e(A, \mathcal{C})]$.
- For any $A \in \mathcal{A}$, let $\tilde{A} \in \tilde{\mathcal{A}}$ satisfy $\sup_{x \in \mathcal{X}} \|A_{Y|X}(\cdot|x) - \tilde{A}_{Y|X}(\cdot|x)\|_{\mathrm{TV}} \le \epsilon$.
- By $|P_e(A, c) - P_e(\tilde{A}, c)| \le \epsilon$,

$$\mathbf{E}_{\mathcal{C}}\left[\sup_{A \in \mathcal{A}} P_e(A, \mathcal{C})\right] \le \mathbf{E}_{\mathcal{C}}\left[\sum_{\tilde{A} \in \tilde{\mathcal{A}}} P_e(\tilde{A}, \mathcal{C}) + \epsilon\right] = \sum_{\tilde{A} \in \tilde{\mathcal{A}}} P_e(\tilde{A}) + \epsilon \le |\tilde{\mathcal{A}}| \cdot \overline{P_e} + \epsilon,$$

and complete the proof by the existence of a codebook c such that $\sup_{A \in \mathcal{A}} P_e(A, c) \le |\tilde{\mathcal{A}}| \cdot \overline{P_e} + \epsilon$.

Remark

- When $K = \emptyset$, $d_1(s, x) = 0$, and $\mathcal{A} = \{A_{Y|X}\}$ is a singleton set, taking $\hat{A}_{Y|X} = A_{Y|X}$, our theorem reduces to the one-shot Gelfand-Pinsker coding result in Li and Anantharam (2021).

Recovering the Asymptotic Result

Proposition: simple bound on the ϵ-covering number

If $\mathcal{X}$ and $\mathcal{Y}$ are discrete and finite, then

$$N_\epsilon(\mathcal{A}) \le \left(\frac{1}{2\epsilon} + \frac{|\mathcal{Y}|+1}{2}\right)^{|\mathcal{X}|\cdot|\mathcal{Y}|}. \tag{1}$$

Proof of the Proposition

- Write $d(A,\tilde{A}) := \sup_{x\in\mathcal{X}} \|A_{Y|X}(\cdot|x) - \tilde{A}_{Y|X}(\cdot|x)\|_{\mathrm{TV}}$.
- Start with $\tilde{\mathcal{A}} = \emptyset$, add $A \in \mathcal{A}$ not currently covered by $\tilde{\mathcal{A}}$ to $\tilde{\mathcal{A}}$ one by one. The $(\epsilon/2)$-balls $\{A :\ d(A,\tilde{A}) \le \epsilon/2\}$ must be disjoint for $\tilde{A} \in \tilde{\mathcal{A}}$.
- Treat $A_{Y|X}$ as a transition probability matrix $A \in \mathbb{R}^{|\mathcal{Y}|\times|\mathcal{X}|}$.
- The volume of $\{A \in \mathbb{R}^{|\mathcal{Y}|\times|\mathcal{X}|} :\ d(A,\tilde{A}) \le \epsilon/2\}$ is $b := ((2\epsilon)^{|\mathcal{Y}|}/(|\mathcal{Y}|!))^{|\mathcal{X}|}$. They are subsets of $\{A \in \mathbb{R}^{|\mathcal{Y}|\times|\mathcal{X}|} : \min_{x,y} A_{y,x} \ge -\epsilon,\ \max_x \sum_y A_{y,x} \le 1+\epsilon\}$, which has a volume $B := ((1+(|\mathcal{Y}|+1)\epsilon)^{|\mathcal{Y}|}/(|\mathcal{Y}|!))^{|\mathcal{X}|}$.
- Hence $|\tilde{\mathcal{A}}|$ is bounded by $\frac{B}{b}$, giving (1).

Recovering the Asymptotic Result

Recovering the Asymptotic Result

- Moulin and O'Sullivan (2003): S, K, X, Y are finite and discrete, and $A_{Y|X}$ must be memoryless and is subject to a distortion constraint.
- Consider sequences $S^n = (S_1, \ldots, S_n)$, K^n, X^n, Y^n where $(S_i, K_i) \overset{iid}{\sim} P_{S,K}$.
- For a input distribution P_X, the class of attackers $\mathcal{A}_n = \mathcal{A}_n(P_X)$ is

$$\mathcal{A}_n(P_X) := \left\{A^n_{Y|X} : A_{Y|X} \in \mathcal{A}(P_X)\right\},$$

$$\mathcal{A}(P_X) := \left\{A_{Y|X} : \mathbf{E}_{(X,Y)\sim P_X A_{Y|X}}[d_2(X, Y)] \le \mathrm{D}_2\right\},$$

  where d_2 is a distortion measure and $A^n_{Y|X}$ is memoryless.
- The asymptotic hiding capacity given in Moulin and O'Sullivan (2003) is

$$C = \max_{P_{U,X|S,K}} \min_{A_{Y|X}: \mathbf{E}[d_2(X,Y)] \le \mathrm{D}_2} \Big(I(U; Y|K) - I(U; S|K)\Big).$$

  where the maximum is over $P_{U,X|S,K}$ with $\mathbf{E}[d_1(S, X)] \le \mathrm{D}_1$.

Recovering the Asymptotic Result

Recovering the Asymptotic Result

- Let $P_{U,X|S,K}$ achieve the maximum subject to $\mathbf{E}[d_1(S,X)] \le \mathrm{D}_1'$, $\mathrm{D}_1' < \mathrm{D}_1$.
- $\hat{A}_{Y|X}$ is the minimizer of R-D function $\min_{A_{Y|X}: \mathbf{E}[d_2(X,Y)] \le \mathrm{D}_2} I(U;Y|K)$.
- We show a rate $R < \hat{I}(U;Y|K) - I(U;S|K)$ is achievable:
  1. For any $A_{Y|X}$ with $\mathbf{E}[d_2(X,Y)] \le \mathrm{D}_2$, let $A^{\lambda}_{Y|X} := (1-\lambda)\hat{A}_{Y|X} + \lambda A_{Y|X}$.
  2. For $I_\lambda(U;Y|K)$ (from $Y|X \sim A^{\lambda}_{Y|X}$), check
  $$\frac{\mathrm{d}}{\mathrm{d}\lambda} I_\lambda(U;Y|K)\Big|_{\lambda=0} = \mathbf{E}_{Y|X\sim A_{Y|X}}[\hat{\imath}(U;Y|K)] - \hat{I}(U;Y|K),$$
  which is nonnegative due to the optimality of $\hat{A}$.
  3. For i.i.d. sequences $(S^n, K^n, U^n, X^n, Y^n) \sim P^n_{S,K} P^n_{U,X|S,K} A^n_{Y|X}$ and $\mathrm{L} = \lfloor 2^{nR} \rfloor$, by the law of large numbers, as $n \to \infty$, exponentially
  $$\mathrm{L} 2^{-\hat{\imath}(U^n;Y^n|K^n)+\imath(U^n;S^n|K^n)} \le 2^{nR - \sum_{i=1}^{n}(\hat{\imath}(U_i;Y_i|K_i) - \imath(U_i;S_i|K_i))} \to 0.$$
  4. To bound the $N_\epsilon(\mathcal{A}_n(P_X))$: we construct a ϵ-cover of $\mathcal{A}_n(P_X)$ using an (ϵ/n)-cover of $\mathcal{A}(P_X)$, hence $N_\epsilon(\mathcal{A}_n(P_X)) \le N_{\epsilon/n}(\mathcal{A}(P_X)) = O((n/\epsilon)^{|\mathcal{X}|\cdot|\mathcal{Y}|})$, which grows much slower than the exponential decrease of the expectation
  5. Take $\epsilon = 1/n$, $P_e \to 0$ as $n \to \infty$. Taking $\mathrm{D}_1' \to \mathrm{D}_1$ completes the proof.

Summary

Our Contributions

- A **one-shot** analysis for the information hiding problem that:
  1. applies to arbitrary channels (not necessarily memoryless or ergodic), any host distribution and any class of attackers;
  2. without assuming the decoder knows the attack channel;
  3. provides a simple alternative proof of the asymptotic hiding capacity.

Future Directions

- Generalized family of Gelfand-Pinsker problems by Moulin and Wang (2007).
- Stegotext reconstruction (also recover encoded data X) like Grover et al. (2015) or Xu et al. (2023).
- Better design of AI-based watermarking tools (e.g., Ji et al. (2024) designed watermarking tools following Moulin and O'Sullivan (2003)).

Moulin, P., & Wang, Y. (2007). Capacity and random-coding exponents for channel coding with side information. IEEE Transactions on Information Theory, 53(4), 1326-1347.

Grover, P., Wagner, A. B., & Sahai, A. (2015). Information embedding and the triple role of control. IEEE Transactions on Information Theory, 61(4), 1539-1549.

Xu, Y., Lu, J., Guang, X., & Xu, W. (2023). Information Embedding With Stegotext Reconstruction. IEEE Transactions on Information Forensics and Security, 19, 1415-1428.

Ji, Z., Hu, Q., Zheng, Y., Xiang, L., & Wang, X. (2024). A Principled Approach to Natural Language Watermarking. In ACM Multimedia 2024.

Acknowledgement

This work was partially supported by two grants from the Research Grants Council of the Hong Kong Special Administrative Region, China [Project No.s: CUHK 24205621 (ECS), CUHK 14209823 (GRF)].

References

Blackwell, D., Breiman, L., Thomasian, A., et al. (1959). The capacity of a class of channels. *The Annals of Mathematical Statistics*, 30(4):1229–1241.

Feinstein, A. (1954). A new basic theorem of information theory. *IRE Trans. Inf. Theory*, (4):2–22.

Kostina, V. and Verdú, S. (2012). Fixed-length lossy compression in the finite blocklength regime. *IEEE Trans. Inf. Theory*, 58(6):3309–3338.

Lei, E., Hassani, H., and Bidokhti, S. S. (2022). Neural estimation of the rate-distortion function with applications to operational source coding. *IEEE Journal on Selected Areas in Information Theory*, 3(4):674–686.

Li, C. T. and Anantharam, V. (2021). A unified framework for one-shot achievability via the poisson matching lemma. *IEEE Transactions on Information Theory*, 67(5):2624–2651.

Li, C. T., Wu, X., Ozgur, A., and El Gamal, A. (2020). Minimax learning for distributed inference. *IEEE Transactions on Information Theory*, 66(12):7929–7938.

Liu, Y., Chen, W.-N., Özgür, A., and Li, C. T. (2024). Universal exact compression of differentially private mechanisms. *Advances in Neural Information Processing Systems*.

Liu, Y. and Li, C. T. (2024). One-shot coding over general noisy networks. In *2024 IEEE International Symposium on Information Theory (ISIT)*, pages 3124–3129.

Moulin, P. and O'Sullivan, J. A. (2003). Information-theoretic analysis of information hiding. *IEEE Transactions on information theory*, 49(3):563–593.

Moulin, P. and Wang, Y. (2007). Capacity and random-coding exponents for channel coding with side information. *IEEE Transactions on Information Theory*, 53(4):1326–1347.

Polyanskiy, Y., Poor, H. V., and Verdú, S. (2010). Channel coding rate in the finite blocklength regime. *IEEE Trans. Inf. Theory*, 56(5):2307–2359.

Shannon, C. E. (1957). Certain results in coding theory for noisy channels. *Information and control*, 1(1):6–25.

Somekh-Baruch, A. and Merhav, N. (2004). On the capacity game of public watermarking systems. *IEEE Transactions on Information Theory*, 50(3):511–524.

Xu, Y., Lu, J., Guang, X., and Xu, W. (2023). Information embedding with stegotext reconstruction. *IEEE Transactions on Information Forensics and Security*, 19:1415–1428.

Yassaee, M. H., Aref, M. R., and Gohari, A. (2013). Non-asymptotic output statistics of random binning and its applications. In *2013 IEEE International Symposium on Information Theory*, pages 1849–1853. IEEE.