

PRIVACY-AWARE COMMUNICATION OVER A WIRETAP CHANNEL WITH GENERATIVE NETWORKS

Ecenaz Erdemir, Pier Luigi Dragotti and Deniz Gündüz

Department of Electrical and Electronic Engineering, Imperial College London, UK

ABSTRACT

We study privacy-aware communication over a wiretap channel using end-to-end learning. Alice wants to transmit a source signal to Bob over a binary symmetric channel, while passive eavesdropper Eve tries to infer some sensitive attribute of Alice’s source based on its overheard signal. Since we usually do not have access to true distributions, we propose a data-driven approach using variational autoencoder (VAE)-based joint source channel coding (JSCC). We show through simulations with the colored MNIST dataset that our approach provides high reconstruction quality at the receiver while confusing the eavesdropper about the latent sensitive attribute, which consists of the color and thickness of the digits. Finally, we consider a parallel-channel scenario, and show that our approach arranges the information transmission such that the channels with higher noise levels at the eavesdropper carry the sensitive information, while the non-sensitive information is transmitted over more vulnerable channels.

Index Terms— Privacy-utility trade-off, wiretap channel, physical layer security, generative networks, variational autoencoders.

1. INTRODUCTION

Secrecy and privacy in data communication and data sharing systems have been extensively studied in the literature [1–12]. Although deep learning applications of data transmission have also been well investigated, deep learning in wireless communications and physical layer security has only recently become popular [13, 14]. The similarity between the communication systems and end-to-end learning motivates the use of autoencoder based neural network architectures, which simultaneously learn encoding and decoding [14, 15]. Recently, it has been shown that end-to-end approaches can also be utilized for physical layer secrecy [16–19]. In a wiretap channel setting, these techniques exploit the physical characteristics of the legitimate receiver’s channel over the eavesdropper’s, and allow communication with secrecy guarantees.

In this work, we consider a wiretap channel scenario in which Alice wants to deliver its source, S^m , to Bob over a noisy communication channel, while a passive eavesdropper Eve tries to infer a latent sensitive information T about S^m . For example, S^m may be an image or a video captured by Alice, while T may be the presence of a particular object or an activity within the scene. We assume binary symmetric

channels (BSCs) from Alice to both Bob and Eve. Our aim is to optimize the trade-off between the reconstruction distortion of source S^m at Bob and the privacy leakage of T to Eve, which is measured by the mutual information (MI) between the sensitive information and the noisy codewords observed by Eve. Note that, the wiretap channel model considered here is normally studied in the context of secure communications. Indeed, when $T = S^m$, our problem becomes a special case of the one studied in [20]. We, instead, call this “privacy-aware communications” since secrecy typically focuses on making the information leakage negligible, while privacy tolerates some leakage in return of utility [21]. Hence, we propose a privacy-utility trade-off (PUT) for communication over the wiretap channel by balancing the information leakage to the eavesdropper and the distortion at the legitimate receiver. We highlight that in the special case of identical channels to Bob and Eve, our problem also reduces to the well-known privacy funnel [7] with a noisy communication channel. In that scenario, Bob and Eve merge into a single receiver, to which we want to send S^m with the highest fidelity while hiding T . Therefore, our problem generalizes both the wiretap channel and the privacy funnel problems. Additionally, unlike in [20] and [7], we follow a data-driven approach by using an encoder-decoder pair, represented by a VAE network and a classifier which represents the eavesdropper.

Similar data-driven wiretap channel approaches have recently been proposed for Gaussian channels in [16–19]. However, [16, 18, 19] focus on channel coding, and [18, 19] enforce coding structure to the encoder, while we carry out end-to-end joint learning corresponding to a JSCC approach. In addition, unlike these works, we are interested in hiding an underlying sensitive information that is correlated with, but different from the original signal. The same problem is considered in [17] for an additive white Gaussian channel using a generative adversarial network (GAN), which minimizes the distortion of the reconstructed signal at the legitimate receiver while characterizing the privacy with a constraint on the likelihood of the sensitive information. On the other hand, we propose a PUT for a BSC wiretap channel using a VAE-based neural network architecture.

VAEs provide several advantages in this framework compared to standard autoencoders (AEs) [15]. They embed the input to a distribution rather than a point, and a random

channel input is sampled from the latent distribution rather than being generated by the encoder directly. Hence, VAEs are more aligned with the stochastic encoding approach employed in information theoretic derivation of the wiretap channel capacity [2, 20]. Additionally, VAEs provide significant control over how to model the latent distribution, since the encoder is designed as a generative network. This is difficult to achieve within the AE framework, and also allows a tractable calculation of the variational approximations of our cost function based on MI. Last but not least, it is challenging to optimize AEs for communication over discrete channels due to their non-differentiability, whereas sampling discrete codewords from a latent distribution is possible for VAEs.

We apply our approach to privacy-aware image transmission and show that while the receiver can reconstruct high quality images, the eavesdropper is confused about the sensitive information. We also consider a parallel-channel case in which Bob and Eve might experience different noise levels over each channel. We show that our end-to-end approach judiciously adjusts its transmission to exploit the more secure channels to transmit the sensitive information.

2. PROBLEM STATEMENT

We consider a communication scenario in which a user wants to reliably transmit data from one point to another over a noisy communication channel, while a passive eavesdropper tries to infer a latent sensitive information through its noisy observation of the transmitted signal. Fig. 1 illustrates the communication problem via a simple example. Alice wants to reveal her data $S^m \in \mathcal{S}$, e.g., images of the applicants for a certain job position, to Bob over a noisy channel. Eve eavesdrops through her own channel and receives a noisy version of the transmitted signal by Alice. Eve's goal is to extract Alice's sensitive information $T \in \mathcal{T}$, e.g., ethic or socioeconomic background of the applicants, which is correlated with S^m but not explicitly observed by any of the involved parties. Alice's goal, on the other hand, is to encode the source such that it can be reconstructed by Bob with high fidelity, while the sensitive information T cannot be accurately detected by Eve. The source is encoded into codewords $X^n \in \{x_1, \dots, x_n\}$, where $X_i \in \mathcal{X} = \{0, 1\}$, by a stochastic encoding function $f_{enc}(S^m) = X^n$ represented by a conditional distribution $P(X^n|S^m)$. We consider a BSC characterized by the joint conditional distribution $P(Y_B^n, Y_E^n|X^n)$, $Y_{B,i}, Y_{E,i} \in \mathcal{X}$. The noisy codeword received by Bob is decoded as $f_{dec}(Y_B^n) = \hat{S}^m$, and Eve receives its own noisy observation Y_E^n .

We model the joint distribution of $T, S^m, X^n, Y_B^n, \hat{S}^m$, i.e., the random variables (r.v.'s) for the sensitive information, source signal, transmitted codeword, noisy codeword received by Bob, and his reconstruction, respectively, using the following graphical model $T - S^m \rightarrow X^n \rightarrow Y_B^n \rightarrow \hat{S}^m$ as:

$$P(T, S^m, X^n, Y_B^n, \hat{S}^m) = P(T, S^m)P(X^n|S^m)P(Y_B^n, Y_E^n|X^n)P(\hat{S}^m|Y_B^n). \quad (1)$$

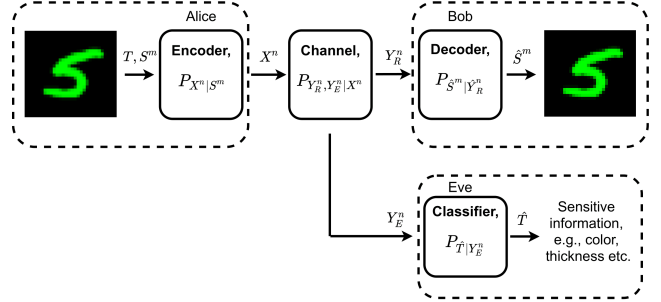


Fig. 1: Communication system with wiretap channel.

The two BSCs independently flip each bit in the transmitted codeword with crossover probabilities ϵ_B and ϵ_E at Bob's and Eve's channels, respectively. Hence, the joint probability of the channel can be decomposed as follows:

$$P(Y_B^n|X^n) = \prod_{i=1}^n \epsilon_B^{x_i \oplus y_{B,i}} (1 - \epsilon_B)^{x_i \oplus y_{B,i} \oplus 1}, \quad (2)$$

$$P(Y_E^n|X^n) = \prod_{i=1}^n \epsilon_E^{x_i \oplus y_{E,i}} (1 - \epsilon_E)^{x_i \oplus y_{E,i} \oplus 1}, \quad (3)$$

where \oplus represents the exclusive OR operation, and $x_i, y_{B,i}$ and $y_{E,i}$ are the i^{th} bits of X^n, Y_B^n and Y_E^n , respectively.

We formulate the optimization problem as

$$\begin{aligned} \min_{f_{enc}, f_{dec}} \quad & \mathbb{E}[d(S^m, \hat{S}^m)] - I(S^m; Y_B^n) + \lambda I(T; Y_E^n) \\ \text{s.t.} \quad & T, S^m \rightarrow X^n \rightarrow Y_B^n \rightarrow \hat{S}^m, \end{aligned} \quad (4)$$

where λ is the tuning parameter for the privacy level. Here, in addition to the reconstruction distortion between S^m and \hat{S}^m , measured by $d(\cdot, \cdot)$, we also maximize the MI between the user's data S^m and the noisy codewords observed by Bob, i.e., $I(S^m; Y_B^n)$, for improved utility. While minimizing the distortion $\mathbb{E}[d(S^m, \hat{S}^m)]$ improves pixel-wise data reconstruction quality, we have observed in our simulations that maximizing the MI between the source signal and Bob's channel output enhances the information flow and helps with capturing the high level features at the receiver side.

Exact calculation of the MI is difficult when the data distribution is not known. Hence, we approximate $I(S^m; Y_B^n)$ and $I(T; Y_E^n)$ via their variational representations [22]. Due to the intractability of the true posteriors $P(S^m|Y_B^n)$ and $P(T|Y_E^n)$, we use their amortized variational approximations $f_{enc}(Y_B^n) = P(\hat{S}^m|Y_B^n)$ and $f_{eve}(Y_E^n) = P(\hat{T}|Y_E^n)$, respectively. Here, we assume that the eavesdropper tries to predict the sensitive information T as $f_{eve}(Y_E^n) = \hat{T}$. We can write $I(S^m; Y_B^n)$ as follows:

$$I(S^m; Y_B^n) = H(S^m) - H(S^m|Y_B^n) \quad (5)$$

$$= H(S^m) + D(P(S^m|Y_B^n) || f_{dec}(Y_B^n)) + \mathbb{E}[\log f_{dec}(Y_B^n)] \quad (6)$$

$$\geq H(S^m) + \max_{f_{dec}} \mathbb{E}[\log f_{dec}(Y_B^n)], \quad (7)$$




Fig. 2: PUT curve of our privacy-aware JSCC mechanism for $\epsilon_B = 0.1$ and $\epsilon_E = \{0, 0.2, 0.3\}$.

where $D(\cdot||\cdot)$ denotes the KL divergence, $H(S^m)$ is constant, (5) follows from the definition of MI, (6) holds for any distribution $f_{dec}(Y_B^n)$ over S^m given the values in Y_B^n . Finally, (7) follows from the fact that maximum is attained when the decoder is optimum, i.e., $f_{dec}(Y_B^n) = P(S^m|Y_B^n)$. Likewise, the information leakage to the eavesdropper becomes

$$I(T; Y_E^n) = H(T) - H(T|Y_E^n) \quad (8)$$

$$= H(T) + D(P(T|Y_E^n)||f_{eve}(Y_E^n)) + \mathbb{E}[\log f_{eve}(Y_E^n)] \quad (9)$$

$$\geq H(T) + \max_{f_{eve}} \mathbb{E}[\log f_{eve}(Y_E^n)], \quad (10)$$

where $H(T)$ is a constant term, (8), (9) and (10) follow similarly to (5), (6) and (7), respectively. Here, (7) is attained when the decoder is optimum since we maximize $I(S^m; Y_B^n)$ in our objective. However, (10) is not attained even if the classifier representing the eavesdropper is optimum, because we minimize $I(T; Y_E^n)$ in the objective. This is due to intractability of representing $I(T; Y_E^n)$ with an upper-bound [23]. On the other hand, our numerical results indicate that although we do not optimize exact bounds for MI terms, in practice our model still learns an effective PUT.

2.1. PARALLEL-CHANNEL SCENARIO

In this section, we assume the codewords are transmitted over parallel channels with different noise levels, e.g., due to OFDM. Our setting represents the scenario in which the transmitter divides the total available bandwidth into non-overlapping bands carrying separate portions of the data. Each of the parallel bands face a different noise level for both the receiver and the eavesdropper, i.e., $(\epsilon_{B_i}, \epsilon_{E_i})$. For instance, in a three-channel scenario with equal bandwidths $n/3$, crossover probabilities $\epsilon_B = \{\epsilon_{B_1}, \epsilon_{B_2}, \epsilon_{B_3}\}$ and $\epsilon_E = \{\epsilon_{E_1}, \epsilon_{E_2}, \epsilon_{E_3}\}$, channel probabilities can be written as

$$P(Y_B^n | X^n) = \prod_{i=1}^{\frac{n}{3}} \epsilon_{B_1}^{x_i \oplus y_{B,i}} (1 - \epsilon_{B_1})^{x_i \oplus y_{B,i} \oplus 1} \quad (11)$$

$$\times \prod_{j=\frac{n}{3}+1}^{\frac{2n}{3}} \epsilon_{B_2}^{x_j \oplus y_{B,j}} (1 - \epsilon_{B_2})^{x_j \oplus y_{B,j} \oplus 1} \prod_{k=\frac{2n}{3}+1}^n \epsilon_{B_3}^{x_k \oplus y_{B,k}} (1 - \epsilon_{B_3})^{x_k \oplus y_{B,k} \oplus 1}$$

for the receiver, and as follows for the eavesdropper:

$$P(Y_E^n | X^n) = \prod_{i=1}^{\frac{n}{3}} \epsilon_{E_1}^{x_i \oplus y_{E,i}} (1 - \epsilon_{E_1})^{x_i \oplus y_{E,i} \oplus 1} \quad (12)$$

$$\times \prod_{j=\frac{n}{3}+1}^{\frac{2n}{3}} \epsilon_{E_2}^{x_j \oplus y_{E,j}} (1 - \epsilon_{E_2})^{x_j \oplus y_{E,j} \oplus 1} \prod_{k=\frac{2n}{3}+1}^n \epsilon_{E_3}^{x_k \oplus y_{E,k}} (1 - \epsilon_{E_3})^{x_k \oplus y_{E,k} \oplus 1}.$$

We solve (4) using the channel probabilities (11) and (12). We want our solution for (4) to control the transmission through the channels such that the sensitive information T is transmitted over the channels in which Eve experiences high noise, while the rest of the source is transmitted over the channels Bob experiences low noise, independent of Eve's channel. We numerically verify that the proposed VAE-based encoder indeed satisfies these expectations.

3. SIMULATION RESULTS

We consider the wiretap channel in Fig. 1, where the encoder and decoder at Alice and Bob are represented by a VAE, while Eve employs a classifier. For the encoder-decoder pair, we employed the network structure "NECST" proposed in [24]. We designed our privacy aware JSCC network by incorporating our classifier based eavesdropper in NECST. We used colored MNIST handwritten digits as S^m for $m = 32 \times 32$ pixels, and color and thickness of the digits as the sensitive information $T \in \mathcal{T} = \{(R, 0), (R, 1), (R, 2), (G, 0), (G, 1), (G, 2), (B, 0), (B, 1), (B, 2)\}$, where R, G and B denote red, green and blue colors, while 0, 1 and 2 represent thin, medium and thick digits, respectively. We set the total channel bandwidth to $n = 200$ bits. We present MI using mutual information neural estimator (MINE) [23], which is a realistic approximation of MI. Due to its high computational complexity, we only used [23] for calculating the information leakage in the test time.

3.1. Single Channel

We first consider a single channel scenario. In Fig. 2, information leakage $I(T; Y_E^n)$ and Eve's classification accuracy are shown with respect to ℓ_2 distortion per image. Dotted, dashed and straight lines represent the cases with $\epsilon_E = 0$, $\epsilon_E = 0.2$ and $\epsilon_E = 0.3$, respectively, while we have $\epsilon_B = 0.1$ for all cases. Data points are taken at $\lambda = \{0, 5, 10, 20\}$. Fig. 2 shows that the information leakage about the sensitive information decreases as the image distortion increases, which is expected due to the PUT. This is the case even when there is no noise in Eve's channel, i.e., $\epsilon_E = 0$, since minimizing the MI between T and Y_E^n provides a certain level of privacy.




Fig. 3: Original images and their reconstructions by Bob and Eve from top to bottom, respectively, for $\epsilon_B = 0.1$, $\epsilon_E = 0.3$.

Moreover, noisier eavesdropper channel leaks less information at the same level of distortion. Similar trend can be seen for Eve’s accuracy. In Fig. 2, we also observe that a MI gap as small as 0.06 corresponds to 20% accuracy gap between $\epsilon_E = 0.2$ and $\epsilon_E = 0.3$ cases.

For illustration purposes, we trained an additional decoder on the noisy bits received by Eve (Y_E^n) with the same structure as Bob’s decoder. Fig. 3 depicts the original images, reconstructed images by Bob and Eve, respectively, from top to bottom. We can see that in the absence of privacy ($\lambda = 0$), both Bob and Eve can reconstruct the images rather accurately, while, thanks to the employed JSCC approach, Bob’s better channel allows it to have better fidelity. On the other hand, when privacy is imposed ($\lambda = 20$), we can see that Eve cannot recover neither the colour nor the thickness information. On the other hand, we can see that this information is available to Bob; and hence, it has been successfully hidden from Eve while being available in the transmitted signal.

3.2. Parallel Channels

Next, we consider a parallel-channel scenario, where the signal is transmitted over multiple channels with different noise levels. We use 4 parallel channels each with a bandwidth of $n/4 = 50$ bits. Error probability pairs for Bob’s and Eve’s channels are set as $(\epsilon_B, \epsilon_E) = \{\text{Ch1} : (0.1, 0.1), \text{Ch2} : (0.001, 0.2), \text{Ch3} : (0.2, 0.001), \text{Ch4} : (0.001, 0.001)\}$. Table 1 shows the information leakage, Eve’s classification accuracy on T , and separately on the sensitive attributes *color* and *thickness* for each channel. *Accuracy*, *Color* and *Accuracy*, *Thickness*

Table 1: Information leakage and Eve’s classification accuracy for the sensitive r.v. \hat{T} and individual sensitive attributes at each channel for $\lambda = 10$

Channels	Ch1	Ch2	Ch3	Ch4
$I(T; Y_E^n)$	1.0836	1.5703	0.0689	0.6411
Accuracy, T	13.65%	31.85%	16.15%	19.6%
Accuracy, Color	34.5%	62.35%	41.2%	45.25%
Accuracy, Thickness	35.75%	48.65%	38.05%	38.55%




Fig. 4: Original images and reconstructions by Bob, Eve, Bob’s individual channels (Ch1-4), and Eve’s channels (Ch1-4), respectively, from top to bottom, for $\lambda = 10$.

Thickness are calculated as the success of the classifications for only the *color* and only the *thickness*, respectively. Our privacy-aware generative network obtains the PUT by minimizing the information leakage of the sensitive attributes and the distortion. This leads to smaller information leakage at the best quality channel of Eve, i.e., Ch2, and larger at the worst one, i.e., Ch3. Eve’s classification accuracy of T , and individual color and thickness attributes, are the highest for Ch2 and lowest for Ch1. We observed that Ch1 accuracy is low because the classifier is confused between *blue* and *green*, as well as the *medium* and *thick*, but still has high accuracy for *red* and *thin* attributes. On the other hand, Ch3 has low accuracy for all the attributes. This leads to the difference between the leakage and accuracy results for Ch1 and Ch3.

In Fig. 4, we show the original and reconstructed images by Bob, Eve, Ch1 to Ch4 of Bob, and Ch1 to Ch4 of Eve, respectively, from top to bottom. First three rows show similar results with the single channel case, i.e., Eve is confused about the color and thickness of the digits, while Bob can reconstruct at high quality. Moreover, Ch1 and Ch3 do not have meaningful reconstructions for either Bob or Eve. This is because Eve faces less noise in these channels, which might lead to larger leakage. Hence, our network minimizes the information flow through these channels. Ch2, on the other hand, carries more information than Ch4 since it can better hide the sensitive attributes from Eve while maximizing the information transmission for Bob.

4. CONCLUSION

We proposed a VAE-based privacy-aware communication scheme over a wireless wiretap channel. In our simulation results, we showed that our end-to-end learning approach provides minimally distorted source transmission with maximum channel capacity while minimizing the information leakage about sensitive information to an eavesdropper. We also showed that our approach balances the information flow in a parallel-channel scenario such that the PUT is obtained according to the receiver’s and eavesdropper’s channel noises.

5. REFERENCES

- [1] C. E. Shannon, "Communication theory of secrecy systems," *The Bell system technical journal*, vol. 28, no. 4, pp. 656–715, 1949.
- [2] A. D. Wyner, "The wire-tap channel," *Bell system technical journal*, vol. 54, no. 8, pp. 1355–1387, 1975.
- [3] M. Bloch and J. Barros, *Physical-Layer Security: From Information Theory to Security Engineering*, 01 2011.
- [4] F. du Pin Calmon and N. Fawaz, "Privacy against statistical inference," in *Allerton Conf. on Communication, Control, and Comp.*, 2012, pp. 1401–1408.
- [5] E. Erdemir, D. Gündüz, and P. L. Dragotti, "Smart meter privacy," in *Privacy in Dynamical Systems*, Farhad Farokhi, Ed. Springer Singapore, first edition, 2020.
- [6] A. Zamani, T. Oechtering, and M. Skoglund, "A design framework for epsilon-private data disclosure," *ArXiv*, vol. abs/2009.01704, 2020.
- [7] A. Makhdoumi, S. Salamatian, N. Fawaz, and M. Médard, "From the information bottleneck to the privacy funnel," in *2014 IEEE Information Theory Workshop (ITW 2014)*, 2014, pp. 501–505.
- [8] E. Erdemir, P. L. Dragotti, and D. Gündüz, "Privacy-aware time-series data sharing with deep reinforcement learning," *IEEE Transactions on Information Forensics and Security*, vol. 16, pp. 389–401, 2021.
- [9] B. Rassouli and D. Gündüz, "On perfect privacy," *IEEE Journal on Selected Areas in Information Theory*, pp. 1–1, 2021.
- [10] B. Rassouli and D. Gündüz, "Optimal utility-privacy trade-off with total variation distance as a privacy measure," *IEEE Transactions on Information Forensics and Security*, vol. 15, pp. 594–603, 2020.
- [11] E. Erdemir, P. L. Dragotti, and D. Gündüz, "Active privacy-utility trade-off against a hypothesis testing adversary," *CoRR*, vol. abs/2102.08308, 2021.
- [12] M. Z. Hameed, A. György, and D. Gündüz, "The best defense is a good offense: Adversarial attacks to avoid modulation detection," *IEEE Trans. Info. Forensics and Security*, vol. 16, 2021.
- [13] D. Gündüz, P. de Kerret, N. D. Sidiropoulos, D. Gesbert, C. R. Murthy, and M. van der Schaar, "Machine learning in the air," *IEEE Journal on Selected Areas in Communications*, vol. 37, no. 10, pp. 2184–2199, 2019.
- [14] T. O’Shea and J. Hoydis, "An introduction to deep learning for the physical layer," *IEEE Tran. on Cognitive Comms. and Networking*, vol. 3, no. 4, pp. 563–575, 2017.
- [15] E. Bourtsoulatze, D. B. Kurka, and D. Gündüz, "Deep joint source-channel coding for wireless image transmission," *IEEE Transactions on Cognitive Communications and Networking*, vol. 5, no. 3, pp. 567–579, 2019.
- [16] K.-L. Besser, P.-H. Lin, C. R. Janda, and E. A. Jorswieck, "Wiretap code design by neural network autoencoders," *IEEE Transactions on Information Forensics and Security*, vol. 15, pp. 3374–3386, 2020.
- [17] T. Marchioro, N. Laurenti, and D. Gündüz, "Adversarial networks for secure wireless communications," in *IEEE Int’l Conf. on Acoustics, Speech and Signal Proc. (ICASSP)*, 2020, pp. 8748–8752.
- [18] R. Fritschek, R. F. Schaefer, and G. Wunder, "Deep learning for the Gaussian wiretap channel," in *ICC 2019-2019 IEEE International Conference on Communications (ICC)*. IEEE, 2019, pp. 1–6.
- [19] R. Fritschek, R. F. Schaefer, and G. Wunder, "Deep learning based wiretap coding via mutual information estimation," New York, NY, USA, 2020, WiseML ’20, p. 74–79, Association for Computing Machinery.
- [20] N. Merhav, "Shannon’s secrecy system with informed receivers and its application to systematic coding for wiretapped channels," *IEEE Transactions on Information Theory*, vol. 54, no. 6, pp. 2723–2734, 2008.
- [21] M. Bloch, O. Günlü, A. Yener, F. Oggier, H. V. Poor, L. Sankar, and R. F. Schaefer, "An overview of information-theoretic security and privacy: Metrics, limits and applications," *IEEE Journal on Selected Areas in Information Theory*, vol. 2, no. 1, pp. 5–22, 2021.
- [22] D. Barber and F. Agakov, "The IM algorithm: a variational approach to information maximization," *Advances in Neural Information Processing Systems*, vol. 16, no. 320, pp. 201, 2004.
- [23] B. Poole, S. Ozair, A. Van Den Oord, A. Alemi, and G. Tucker, "On variational bounds of mutual information," in *International Conference on Machine Learning*. PMLR, 2019, pp. 5171–5180.
- [24] K. Choi, K. Tatwawadi, A. Grover, T. Weissman, and S. Ermon, "Neural joint source-channel coding," in *International Conference on Machine Learning*. PMLR, 2019, pp. 1182–1192.