Is Semantic Communication Secure? A Tale of Multi-Domain Adversarial Attacks

Yalin E. Sagduyu, Tugba Erpek, Sennur Ulukus, and Aylin Yener

For semantic communications, transmitter-receiver functionalities are modeled as an autoencoder, followed by a task classifier that evaluates the meaning of the conveyed information. This approach transfers compressed feature vectors reliably with a small number of channel uses while keeping the semantic loss low. Multi-domain security vulnerabilities of using deep neural networks (DNNs) for semantic communications are identified. Based on adversarial machine learning, test-time (targeted and non-targeted) adversarial attacks on these DNNs are introduced.

ABSTRACT

Semantic communication seeks to transfer information from a source while conveying a desired meaning to its destination. We model the transmitter-receiver functionalities as an autoencoder, followed by a task classifier that evaluates the meaning of the conveyed information. The autoencoder consists of an encoder at the transmitter that jointly models source coding, channel coding, and modulation, and a decoder at the receiver that jointly models demodulation, channel decoding, and source decoding. By augmenting the reconstruction loss with a semantic loss, this encoder-decoder pair is interactively trained with the semantic task classifier. This approach transfers compressed feature vectors reliably with a small number of channel uses while keeping the semantic loss low. We identify the multi-domain security vulnerabilities of using deep neural networks (DNNs) for semantic communications. Based on adversarial machine learning, we introduce test-time (targeted and non-targeted) adversarial attacks on these DNNs. As a computer vision attack, small perturbations are injected into the images at the input of the transmitter's encoder. As a wireless attack, small perturbation signals are transmitted to interfere with the input of the receiver's decoder. By launching these attacks individually or jointly (as a multi-domain attack), we show that it is possible to change the semantics of the transferred information (with larger impact than conventional jamming) and highlight the need of defense methods for the safe adoption of semantic communications.

INTRODUCTION

Conventional communication pursues the goal of reliable transfer of messages in terms of symbols (or bits) without a special focus on how the semantics of information pertinent to these messages is preserved. Semantic communications seeks to change this paradigm by preserving the semantics in recovered messages beyond conventional reliability measures. As an example, consider a surveillance system of edge devices equipped with cameras. Each edge device takes images and needs to transfer them over a wireless channel to a security center. The semantics of the transferred information is of paramount importance. For example, if the security center needs to classify images with respect to the intruders detected, the goal would be not only to reconstruct the images reliably at the security center, but also to preserve their semantics, namely minimize the semantic loss with respect to the errors in image classification over wireless links to detect intruders at the security center.

Semantic communications aims to reliably communicate the meanings of messages through a channel by minimizing the semantic error [1] to best preserve the meaning of recovered messages. Semantic communications is envisioned to serve different applications such as text [2], *speech/audio* [3, 4], *image* [5] and *video* [6] communications. Semantic communications has been studied in terms of information-theoretical foundations [7] and networking aspects [8]. In addition, *task-oriented communications* has been formulated to utilize the semantics of information via its significance relative to the goal of information transfer when performing an underlying task [9].

()

In conventional communications, the transmitter and receiver functionalities are typically designed as separate modules such as source coding, channel coding, and modulation at the transmitter and demodulation, channel decoding and source decoding at the receiver. The goal of conventional communications is to reconstruct the transmitter's data samples (messages) at the receiver by minimizing the symbol/bit error rate or a signal distortion metric such as the mean squared error (MSE). The joint design of communication functionalities is ultimately needed to recover the semantic information in addition to the transfer of messages themselves.

Semantic communications can be set up in a deep learning-driven end-to-end communication framework by training an autoencoder that consists of an encoder at the transmitter and a decoder at the receiver. The encoder is modeled as a deep neural network (DNN) for joint operations of source coding, channel coding, and modulation at the transmitter. Then, a second DNN is used to model the *decoder* for joint operations of demodulation, channel decoding, and source decoding at the receiver. The input of the encoder is of general data type (e.g., an image). The encoder and the decoder are separated by a wireless channel such that the output of the encoder is a modulated signal transmitted over a wireless channel and the received signal becomes the input to the decoder. Then, the decoder output

Digital Object Identifier: 10.1109/MCOM.006.2200878 Yalin E. Sagduyu and Tugba Erpek are with Virginia Tech, USA; Sennur Ulukus is with the University of Maryland, USA; Aylin Yener is with The Ohio State University, USA.

(namely, the reconstructed samples) is given as input to a *semantic task classifier* (the third DNN) that aims to verify the semantics of reconstructed samples (e.g., presence or absence of the intruder in the surveillance scenario described above). If the accuracy of this semantic task classifier is high, then we can say that the semantics of the information is preserved with high fidelity.

The autoenconder is trained by accounting for channel effects as well as preserving the semantics of information. To that end, semantic communications extends autoencoder communications, where the encoder encompasses channel coding and modulation operations, the decoder encompasses demodulation and channel decoding operations, and the sole goal is the same as conventional communications, namely reconstructing messages (in form of symbols) at the receiver [10]. On the other hand, the autoencoder for semantic communications incorporates source coding and source decoding, and reconstructs the input data samples such as images. More importantly, this autoencoder is trained by a custom loss function that augments the reconstruction loss (e.g., the MSE between the input image at the transmitter and the reconstructed image at the receiver) with the semantic loss that is represented by the penalty of violating the constraint that the loss of the semantic task classifier (that is designed to capture the semantics of the recovered information) exceeds a target threshold.

The training of the autoencoder and the semantic task classifier can be separated or combined. The former way is to consider a fixed task classifier that is trained offline (such as in [11]). However, if this classifier is trained with clean input data without taking the channel effects and the corresponding reconstruction losses into account, it cannot achieve high accuracy especially in the low signal-to-noise ratio (SNR) regime. On the other hand, the autoencoder's output is not known in advance before training it with respect to channel effects, so it cannot be readily used as the input to the semantic task classifier for offline training purposes.

To that end, we consider *interactive training* of the autoencoder and the semantic task classifier over multiple rounds. In each round, the autoencoder for semantic communications is trained first and then its output is used to build the training data that is leveraged to train the semantic task classifier. Along with the reconstruction loss (the MSE loss), the loss of this classifier is then used in the custom loss function of the autoencoder for the next round. This process is repeated over multiple rounds while ingesting new training and validation data samples in each round. This training process seeks to improve the fidelity of both the autoencoder and the semantic task classifier.

As deep learning becomes a core part of semantic communications, there is an increasing concern about the vulnerability of the underlying DNNs to *adversarial effects*. In our case, three DNNs are utilized, an encoder at the transmitter and a decoder and a classifier at the receiver. Smart adversaries may leverage emerging machine learning techniques to exploit vulnerabilities and tamper with the learning functionalities of all these DNNs embedded in semantic communications. Learning in the presence of adversaries has been studied under *adversarial* machine learning for various data domains such as computer vision and natural language processing (NLP). Due to the shared and open nature of wireless medium, wireless applications are highly susceptible to adversaries such as eavesdroppers and jammers that can further observe and manipulate the training and test (inference) time operations of machine learning used for wireless applications [12].

In test time, an adversarial (evasion) attack can add a small perturbation to the input samples of a victim DNN and fool it into making wrong decisions. The complex decision of the DNN makes it highly sensitive to even small variations in the input samples. Both the encoder at the transmitter and the decoder at the receiver take inputs that can be manipulated by the adversaries. The input of the encoder is a signal of general type. If it is an image such as in the surveillance scenario discussed earlier, the adversary can position a small deceptive object in front of the camera and this is captured by the camera as a small perturbation in this computer vision attack. The input of the decoder is the wireless signal received over the channel. This signal can be manipulated by the adversary that transmits a perturbation signal over the air. This way, the received signal includes the perturbation signal superimposed with the transmitted signal and receiver noise. Adversarial attacks on wireless signals have been considered for signal classifications [13] and autoencoder communications [14]. Semantic noise has been considered in [15] that causes misleading between the intended semantic symbols and received ones for image classification tasks. In both computer vision and wireless attacks, the adversarial perturbations are determined as solutions to an optimization problem to minimize the power of perturbation signal subject to the condition that the DNN's decision becomes incorrect.

The adversary can launch either a *non-targeted attack* (where the adversary seeks to change the semantics of recovered information to any other incorrect meaning) or a *targeted attack* (where the adversary seeks to change the semantics of recovered information to a specific incorrect meaning). All these attacks are very effective even when a small perturbation is used, and significantly outperform conventional attacks such as jamming, where a Gaussian signal is transmitted as the perturbation.

We also present a *multi-domain attack*, where these attacks can be launched either separately or together by adding perturbations to the input data samples (e.g., images) as well as to the channel data (with over-the-air transmissions). We show that these attacks are very effective individually or better when combined in a multi-domain attack. Combining computer vision and wireless attacks (each using only a small perturbation added to the input image or the wireless signal) substantially reduces the performance of semantic communications beyond what a single-domain attack can individually achieve.

The rest of the article is organized as follows. The next section describes the deep learning-driven autoencoder-based semantic communications and evaluates its performance. We then present the multi-domain adversarial attack vectors against semantic communications and highlight

As deep learning becomes a core part of semantic communications, there is an increasing concern about the vulnerability of the underlying DNNs to adversarial effects. Combining computer vision and wireless attacks (each using only a small perturbation added to the input image or the wireless signal) substantially reduces the performance of semantic communications beyond what a single-domain attack can individually achieve.



FIGURE 1. System model for semantic communications.

the major loss in preserving the semantic information beyond the reconstruction loss. The final section concludes the article.

Semantic Communications with End-to-end Deep Learning

We consider deep learning-enabled semantic communications shown in Fig. 1. The transmitter-receiver functionalities are designed as an autoencoder. The encoder that is trained as a DNN at the transmitter takes the data samples (e.g., images) as the input and jointly performs source coding, channel coding and modulation operations. The output of the encoder is the modulated signals that are transmitted over the air in multiple channel uses. The size of latent space at the output of the encoder corresponds to the number of channel uses conceptually. The decoder that is trained as another DNN at the receiver takes the received signals as the input and jointly performs demodulation, channel decoding and source decoding operations. The output of the decoder is the reconstructed data samples.

The encoder and decoder DNNs are trained jointly. Autoencoder communications has been considered for joint training of channel coding and modulation at the transmitter and demodulation and channel decoding at the receiver to minimize the categorical cross-entropy (CCE) loss for symbol recovery [10]. In our setting, the input data is a general signal such as an image instead of symbols. Therefore, source coding and source decoding are incorporated in the encoder and decoder, respectively. Then, the goal is extended to reconstruct the input signal at the receiver. To construct data samples at the receiver, a distortion loss can be minimized such as the MSE.

In semantic communications, the goal is not only to reconstruct the signals at the receiver but also preserve the meaning conveyed by the reconstructed data samples by minimizing a custom loss that incorporates both the reconstruction loss and semantic loss for the semantic classifier. The accuracy of the semantic classifier is indirectly reflected in the training of the autoencoder. Overall, the autoencoder and the semantic classifier are jointly trained that goes beyond the traditional training process of the encoder-decoder pair. On the other hand, autoencoder communications would reconstruct the transmitter inputs by minimizing the reconstruction loss, whereas task-oriented communications would classify transmitter inputs by minimizing the classification loss.

For numerical results, we use the MNIST dataset of handwritten digits as the input data and the digit classifier as the semantic task classifier such that the goal of semantic communications is to ensure that the reconstructed signals can be still reliably recognized with respect to its digit labels. The custom loss for semantic communications is the MSE loss of reconstructed signals augmented with the penalty of violating the condition that the target (CCE) loss of the semantic task classifier exceeds a certain threshold, namely, it is the MSE loss plus a weight times the gap of classifier loss from a threshold. For numerical results, we set the weight as 0.2 and the threshold as the loss of classifier with clean inputs in the absence of channel effects. If the autoencoder was trained by reconstruction or semantic loss only, it could not reliably preserve the semantic information or reconstruct images, respectively.

First, we can assume that the semantic task classifier is pretrained using the CCE loss for the clean data (without channel impairments) as in [11]. However, it takes the reconstructed signals as the input and is sensitive to the noise in the input data due to channel effects. This data mismatch in training and test times decreases the accuracy of the semantic classifier and consequently decreases the performance of the autoencoder that makes use of the CCE loss of this classifier as part of its custom loss function. To mitigate this issue, we pursue a multi-round interactive training process as follows. In each round, the autoencoder is retrained by using both the MSE and the CCE loss of the semantic task classifier that was trained in the previous round. Then, the semantic task classifier is retrained using the reconstructed samples collected at the output of the decoder in test time of the current round. Starting with pretraining of the semantic task classifier with clean data, we repeat this process in multiple rounds.

()

The DNN architectures of encoder, decoder and semantic task classifier network are provided in Fig. 2. Feedforward neural networks are used in each DNN. The encoder-decoder pair does not need to be symmetric in general. Hyperparameter optimization can be performed for encoder, decoder and semantic classifier models (such as number of layers and neurons). The MNIST dataset of handwritten digit images (each of 28×28 grayscale pixels) is used as the input data (60K for training and 10K for testing). The wireless transmissions are carried out over an additive white Gaussian noise (AWGN) channel. The accuracy of semantic task classifier for both cases of fixed pretraining and interactive retraining over multiple

52

IEEE Communications Magazine • November 2023

rounds is shown in Fig. 3 as a function of the SNR (when the number of channel uses is varied). This interactive retraining helps preserve the semantic information and the semantic task classifier accuracy increases as the SNR and the number of channel uses increase, while the reconstruction loss remains small. For example, when the number of channel uses is 40, the reconstruction loss is 0.026, 0.021, 0.019, 0.017, and 0.016 when the SNR is 0dB, 3dB, 5dB, 8dB, and 10dB, respectively. The baseline without retraining achieves a similar MSE, whereas semantic communications further preserves the semantics of information transfer.

Adversarial Attacks on Semantic Communications

In test time, adversarial (evasion) attacks seek to manipulate the test data input to the adversary's model (e.g., by adding a small perturbation) such that it cannot make a reliable decision for these samples. The effect of this attack is measured in terms of the model accuracy for the manipulated test input samples (the lower this accuracy drops, the more effective the attack becomes). The perturbation is selected by minimizing the perturbation power subject to the conditions that an error occurs in the decision of the victim model, and the perturbation power remains upper bounded by a threshold. Since solving this optimization problem is difficult, Fast Gradient Method (FGM) can be applied by linearizing the loss function and using the gradient of the loss function when crafting the perturbation. Fast Gradient Sign Method (FGSM) takes the sign of the gradient to design the perturbation. Other attack methods include Basic Iterative Method (BIM), Projected Gradient Descent (PGD), Momentum Iterative Method, DeepFool, and Carlini Wagner (C&W).

The adversary can launch *targeted* and *non-targeted attacks*. The targeted attacks seek to cause errors in the DNN outputs only for samples from a specific set of non-target labels (classes) to other target labels (by minimizing the loss function of the victim DNN with respect to the target labels). On the other hand, the non-targeted attacks seek to cause errors for samples from all labels (by maximizing the loss function of the victim DNN for all samples under the non-targeted attack).

The adversary can launch adversarial attacks on semantic communications in two different ways. First, the adversary adds a small perturbation to the input sample (namely, the image in our case) and manipulate the semantic meaning of messages (namely, the reconstructed image is classified to a wrong digit label) although the reconstruction loss remains small. Second, the adversary adds a small perturbation to the input of the decoder (namely, the received wireless signal) at the receiver (potentially with an overthe-air transmission). These adversarial attacks in different data domains are illustrated in Fig. 4.

First, we consider the adversarial attack on the transmitter (encoder) input. The adversarial perturbation is generated by the FGSM by taking the gradient with respect to the concatenation of the autoencoder and the semantic task classifier. Therefore, the attack depends on the DNN models used. The perturbation is computed as the gradient weighted with the *perturbation strength* and added to the image sample before inputting it to



FIGURE 2. The DNN architectures of the autoencoder and the semantic task classifier.



FIGURE 3. Semantic classification accuracy with and without retraining over time.

the encoder. This corresponds to a *computer vision attack*. The attack success rate is defined as the average error probability of the semantic task classifier under the non-targeted attack and as the probability that the semantic task classifier classifies a non-target label as a target label under the targeted attack. In this article, we consider a white-box attack where the adversary knows the input (either the image or the wireless signal corresponding to the encoded input). To relax this assumption, a universal adversarial perturbation (UAP) can be generated by generating adversarial perturbations for different inputs and corresponding signals, and then reducing the dimension (e.g., via principal component analysis) to a common perturbation [13, 14].

The success rate of the adversarial attack (at 5dB SNR and with 40 channel uses) is shown in Fig. 5 as a function of the perturbation strength. For the targeted attack, we consider two cases; namely, averaged over non-target labels and target labels. We compute the attack success rate for each pair of non-target and target labels. In the former case, we find the best attack success rate over all target labels for a given non-target label, and then average this best attack performance over all target labels. In the latter case, we find the best attack success rate over all non-target labels for a given target label, and then average this best attack performance over all non-target labels. The non-targeted attack is easier as it needs to flip any label to any other label, ۲

()



FIGURE 4. Adversarial attacks on semantic communications.



FIGURE 5. Effects of adversarial attacks from the computer vision domain on semantic communications.



FIGURE 6. Effects of the adversarial attacks from the wireless domain domain and the multi-domain adversarial attacks on semantic communications.

so it achieves high attack performance. Targeted attacks are also effective on the average and likely flip the labels from a specific non-target label to another one or to a specific target-label.

The reconstruction loss caused by each attack is found very close to each other suggesting that adversarial attack poses a more serious threat to the fidelity of semantic information than the information transfer itself. For example, the non-targeted attack with a small perturbation strength such as 0.3 reduces the classifier accuracy to 0.11, although the reconstruction error remains small, namely the MSE is 0.09 (while the MSE is increased, it is still significantly low compared to the range of MSE values possible for poor reconstruction of inputs). This suggests that the adversarial attack can effectively degrade the semantics of conveyed information even when it is reconstructed with a small distortion. For the MNIST data, this means that the reconstructed images at the receiver may look similar to the input images but they cannot be reliably classified to its digit labels. Second, we consider the adversarial attack on the receiver (decoder) input. The adversarial perturbation is generated by FGSM by taking the gradient with respect to the concatenation of the decoder of the autoencoder system and the semantic task classifier. The perturbation is computed as the gradient weighted with the perturbation-to-noise ratio (PNR) and added over the air to the wireless transmission. This corresponds to a wireless attack. The success rate of non-targeted attacks on semantic communications (at 5dB SNR and with 40 channel uses) is shown in Fig. 6. The performance is evaluated as a function of the PNR and compared with the case when Gaussian noise is used as the perturbation such as in conventional jamming attacks. The adversarial perturbation added over the transmitted signals is very effective in reducing the classifier accuracy even when the PNR is low, whereas Gaussian noise is not effective as a perturbation unless the perturbation noise power is much higher than the receiver noise.

۲

Figure 6 also shows the performance under the multi-domain attack combining wireless and computer vision attacks, where a perturbation is added to the receiver input over the air (a wireless attack) and another perturbation (of strength 0.1) is added to the transmitter input image (a computer vision attack). This attack is the most effective one and quickly reduces the classifier even the adversary uses small perturbations of low strength and PNR. The addition of perturbations to the input images does not lose its effect over the wireless transmission when another perturbation is added to the wireless signal. On the contrary, the introduction of the computer vision attack amplifies the effect of the overall adversarial attack and thus further reduces the PNR needed by the wireless adversarial attack to reduce the semantic accuracy below a target threshold.

There are various conventional schemes such as channel hopping, spread spectrum, and adding artificial noise to protect communications from attacks such as jamming. We assume that the adversary operates on the same frequency-time block as the

communications system that pursues bandwidth efficiency (such as envisioned for NextG systems) compared to spread spectrum. In that case, adding artificial noise will reduce the PNR that will slightly reduce the attack success (the adversarial attack can be successful even when the perturbation is below the noise floor as shown in Fig. 6) at the expense of reducing the classifier accuracy since the SNR will also decrease as shown in Fig. 3.

In addition to these test-time attacks considered, it is possible to attack the DNNs in training time such as poisoning (causative) attacks (that manipulate the training data in terms of features and labels), whereas test-time and training-time attacks can be combined in backdoor (Trojan) attacks, where the adversary adds triggers for a target label in training data and later activates them in test time. With the open-source development such as O-RAN for next-generation communications, adversarial machine learning poses a serious threat to the use of deep learning such as we discussed for semantic communications. Therefore, it is essential to characterize this emerging attack surface and develop defense mechanisms.

CONCLUSION

We formulate an autoencoder-based semantic communications system enabled by deep learning to transfer information from a source to its destination while preserving the semantics of information in addition to reliability objectives. The transmitter and receiver functionalities are represented as an encoder-decoder pair that is trained with a custom loss function that combines the reconstruction loss with a semantic loss that is captured by the loss of a subsequent semantic task classifier. By accounting for channel effects, the DNNs for the autoencoder and the semantic task classifier are interactively trained. However, the use of the DNNs makes semantic communications vulnerable to adversarial attacks that seek to manipulate the DNN inputs. These adversarial attacks can be launched in different domains such as a computer vision attack that injects a perturbation to the input image at the transmitter and a wireless attack that transmits a perturbation signal that is received by the decoder at the receiver as superimposed with the transmitted signal. We show that these attacks are very effective individually and even more when combined to reduce the semantic communications performance, namely they lead to a major semantic loss such that the attempt to recover information cannot preserve the semantics.

Future research directions include:

- Extension to other datasets (both image and other data modalities such as text and speech)
- Extension to other machine learning tasks to capture semantic information
- · Extension to other channel properties (e.g., Rayleigh channel)
- Study of hyperparameter optimization for encoder, decoder and semantic classifier models (e.g., number of layers and layer sizes)
- Study of UAPs to generalize adversarial attacks on semantic communications
- Study of other adversarial machine learning attacks (e.g., poisoning and backdoor attacks) on semantic communications
- Study of defense schemes (e.g., adversarial training, randomized smoothing, and certi-

fied defense) against adversarial attacks on semantic communications.

REFERENCES

- [1] B. Guler and A. Yener, "Semantic Index Assignment," Proc. IEEE Int'l. Conf. Pervasive Computing and Commun. Workshops, Budapest, Hungary, 2014, pp. 431-36.
- [2] H. Xie et al., "Deep Learning Enabled Semantic Communication Systems," IEEE Trans. Signal Processing, vol. 69, 2021, pp. 2663-75.
- Z. Weng and Z. Qin, "Semantic Communication Systems for Speech Transmission," IEEE JSAC, vol. 39, no. 8, 2021, pp. 2434-44.
- [4] H. Tong et al., "Federated Learning Based Audio Semantic Communication over Wireless Networks," Proc. IEEE Clobal Commun. Conf., Madrid, Spain, 2021, pp. 1-6.
- Z. Qin et al., "Semantic Communications: Principles and Challenges," arXiv preprint, arXiv:2201.01389, 2021
- [6] P. Jiang et al., "Wireless Semantic Communications for Video Conferencing" IEEE JSAC, vol. 41, no. 1, Jan. 2023, pp. 230–44.
- [7] D. Gündüz et al., "Beyond Transmitting Bits: Context, Semantics, and Task-Oriented Communications," IEEE JSAC, vol. 41, no. 1, Jan. 2023, pp. 5-41.
- [8] E. Uysal et al., "Semantic Communications in Networked Systems," IEEE Network, vol. 36, no. 4, 2022, pp. 233-40.
- [9] J. Shao, Y. Mao, and J. Zhang, "Learning Task-Oriented Com munication for Edge Inference: An Information Bottleneck Approach," IEEE JSAC, vol. 40, no. 1, 2021, pp. 197–211.
 [10] T. J. O'Shea and J. Hoydis, "An Introduction to Deep Learn-ing for the Physical Layer," IEEE Trans. Cognitive Commun.
- Networking, vol. 3, no. 4, 2017, pp. 563-75
- [11] H. Zhang et al., "Deep Learning-Enabled Semantic Communication Systems With Task-Unaware Transmitter and Dynamic Data," *IEEE JSAC*, vol. 41, no. 1, 2023, pp. 170–85. [12] D. Adesina *et al.*, "Adversarial Machine Learning in Wireless
- Communications Using RF Data: A Review," IEEE Commun. Surveys & Tutorials, vol. 25, no. 1, 2023, pp. 77–100.
- [13] B. Kim et al., "Channel-Aware Adversarial Attacks against Deep Learning-Based Wireless Signal Classifiers," IEEE Trans.
- Wireless Commun., vol. 21, no. 6, June 2022, pp. 3868–80.
 M. Sadeghi and E. G. Larsson, "Physical Adversarial Attacks against End-to-End Autoencoder Communication Systems," IEEE Commun. Letters, vol. 23, no. 5, 2019, pp. 847-50.
- [15] Q. Hu et al., "Robust Semantic Communications against Semantic Noise," Proc. IEEE Vehicular Technology Conf., London, United Kingdom, 2022, pp. 1-6.

BIOGRAPHIES

YALIN E. SAGDUYU [S'02, M'08, SM'15] received his Ph.D. degree in Electrical and Computer Engineering from University of Mary-land, College Park. He is a Research Professor at Virginia Tech and a Visiting Research Professor at University of Maryland. Previously, he was the Director of Networks and Security at Intelligent Automation, Inc. He is an Editor of IEEE Trans. Communications and IEEE Trans. Cognitive Communications and Networking. He received an IEEE HST Best Paper Award.

TUGBA ERPEK is a Research Associate Professor at the Intelligent Systems Division of the Virginia Tech National Security Institute. She received her Ph.D. degree in Electrical and Computer Engineering from Virginia Tech. Prior to Virginia Tech, she was a Lead Scientist and Network Communications Technical Area Lead at the Intelligent Automation, Inc. and a Senior Communications Systems Engineer at the Shared Spectrum Company. Her research interests are in wireless communications, networks, security, machine learning, 5G and beyond.

SENNUR ULUKUS [F'16] is the Anthony Ephremides Professor in Information Sciences and Systems in the Department of Electrical and Computer Engineering at the University of Maryland, College Park. She received her Ph.D. from Rutgers University. She received the 2003 IEEE Marconi Prize Paper Award in Wireless Communications, 2019 IEEE Communications Society Best Tutorial Paper Award, 2020 IEEE Communications Society WICE Outstanding Achievement Award, and TCGCC Distinguished Technical Achievement Recognition Award.

AYLIN YENER [F'15] is Roy and Lois Chope Chaired Professor at The Ohio State University. She received the 2014 IEEE Marconi Paper Award, 2018 IÉEE Communications Society WICE Outstanding Achievement Award, 2019 IEEE Communications Society Best Tutorial Paper Award, 2020 IEEE Communication Theory Technical Achievement Award. She is Director-elect for IEEE Division IX. Previously, she was President of the IEEE Information Theory Society. Presently, she is Editor-in-Chief of the IEEE Trans. Green Communications and Networking

()