ARL-MR-0991 • Dec 2018

US Army Research Laboratory

Optimization Techniques for Network Security, Distributed Agents, and RF Sensor Coexistence

by Michael J Weisman, Anthony F Martone, Gunjan Verma, and Robert J Drost

NOTICES

Disclaimers

The findings in this report are not to be construed as an official Department of the Army position unless so designated by other authorized documents.

Citation of manufacturer's or trade names does not constitute an official endorsement or approval of the use thereof.

Destroy this report when it is no longer needed. Do not return it to the originator.

ARL-MR-0991 ● DEC 2018

US Army Research Laboratory

Optimization Techniques for Network Security, Distributed Agents, and RF Sensor Coexistence

by Michael J Weisman, Gunjan Verma, and Robert J Drost
Computational and Information Sciences Directorate, ARL

Anthony F Martone
Sensors and Electron Devices Directorate, ARL

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.

1. REPORT DATE *(DD-MM-YYYY)*	2. REPORT TYPE	3. DATES COVERED (From - To)
Dec 2018	Memorandum Report	October 20–September 2018

4. TITLE AND SUBTITLE
Optimization Techniques for Network Security, Distributed Agents, and RF Sensor Coexistence

5a. CONTRACT NUMBER

5b. GRANT NUMBER

5c. PROGRAM ELEMENT NUMBER

6. AUTHOR(S)
Michael J Weisman, Anthony F Martone, Gunjan Verma, and Robert J Drost

5d. PROJECT NUMBER

5e. TASK NUMBER

5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
US Army Research Laboratory
ATTN: RDRL-CIN-D
Adelphi, MD 20783-1138

8. PERFORMING ORGANIZATION REPORT NUMBER
ARL-MR-0991

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)

10. SPONSOR/MONITOR'S ACRONYM(S)

11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION/AVAILABILITY STATEMENT
Approved for public release; distribution is unlimited.

13. SUPPLEMENTARY NOTES
<michael.j.weisman2.civ@mail.mil>

14. ABSTRACT
A primary objective of the US Army Research Laboratory (ARL) is to bring together the scientific and military communities through collaboration on research that will directly benefit the Warfighter. A wide array of projects that ARL supports involves optimization of a noise-corrupted loss function over a potentially high-dimensional parameter space. In this report, we detail three areas of research that ARL is involved in that may benefit from stochastic optimization: adversarial machine learning, distributed agents, and spectrum sensing for radar.

15. SUBJECT TERMS
adversarial machine learning, stochastic optimization, distributed agents, spectrum sensing multioptimization

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			Michael J Weisman
Unclassified	Unclassified	Unclassified	UU	22	19b. TELEPHONE NUMBER (Include area code) 301-394-1237

Contents

List of Figures

1. Introduction

One objective of the US Army Research Laboratory (ARL) is to bring together the scientific and military communities through collaboration on research that will directly benefit the Warfighter. A wide array of projects that ARL supports involves optimization of a noise-corrupted loss function over a potentially high-dimensional parameter space. In this report, we detail three areas of research that ARL is involved in that may benefit from stochastic optimization.

The first research area, adversarial machine learning, involves deep neural networks, where the number of parameters is vast but evaluation of the loss function (forward-propagation) is relatively cheap. Here, we seek to solve a stochastic optimization problem subject to a minimum norm constraint. In the second research area, distributed agents, the number of parameters is of modest size (of order 100 for typical problems) but the challenge is that access to the full parameter vector is limited in that any given agent in the system can only control a small subset of the parameters. Also, the global loss function to be optimized is an aggregate (e.g., sum) of many local loss functions (e.g., pairwise functions across all pairs of nodes) and each agent can only observe some of the local loss functions. As a result, computing the global loss function itself requires some internode coordination. As we would like to minimize such coordination, we seek methods that minimize the number of loss function evaluations.

Finally, in the third research area, spectrum sensing for radar, we are confronted with a constrained multiobjective function that we wish to optimize. The multiobjective function seeks to maximize two conflicting objective functions, namely the signal-to-interference-plus-noise ratio (SINR) and the bandwidth. A particular challenge is that each constituent function of the overall objective is measured with possibly different levels of noise. As such, we are more likely to encounter higher-power RF interference as the bandwidth increases. The goal then becomes that of identifying the best trade-off available for the measured spectrum, where a sub-band is estimated (bandwidth and center frequency) that optimizes both conditions.

2. Adversarial Machine Learning

2.1 Recent Work

Machine-learning algorithms are being operated not only in benign settings, but also in the presence of malicious attacks.[1] ARL's Computational and Information Sciences Directorate is actively engaged in adversarial machine learning research. For example, in recent work with Penn State,[2] we found that inherent vulnerabilities in traditional machine learning classifiers, such as deep neural networks, can be exploited by an adversary during training and/or at the classification stage. For handwritten digits from the MNIST dataset, an adversary can, with minimal perturbation, cause a deep neural network to misclassify the digit as any other digit. Additionally, of concern to the security of autonomous vehicles, it has been shown that it is possible to fool a deep neural network into misclassifying road signs. For example, a stop sign can be forced to classify as a yield sign.

2.2 Adversarial Attacks on Deep Neural Networks

We first present equations of a simplified multilayer perceptron (a deep neural network may have many more layers). We then define the crafting of an adversarial attack as the solution to an optimization problem, which we denote as P. Presently, we solve P using heuristics. Our interest is twofold:

1. To solve P more efficiently, making use of a smaller number of iterations or evaluations
2. To develop more principled or effective methods and heuristics that achieve even better solutions to P

2.2.1 Multilayer Perceptron

By way of example, we consider a three-layer perceptron made up of an input layer, one hidden layer, and an output layer. Generalizing to a larger number of layers is straightforward. The components of the input feature vector are given by $\mathbf{X} = (x_1, x_2, \ldots, x_n)$. Each neuron h_j in the hidden layer has an activation function ϕ. For example, ϕ is often taken as the sigmoid function $\phi : \; x \mapsto 1/\left(1 + \exp\{-x\}\right)$. The hidden layers each compute

$$h_j(\mathbf{X}) \;=\; \phi\left(\sum_{k=1}^{n} w_{j,k}\, x_k \;+\; w_{j,0}\right),$$

where the $w_{j,k}$ are the weights and the $w_{j,0}$ are bias terms, each computed during training. At the final layer, the output of the network is

$$\mathbf{F}(\mathbf{X}) \;=\; \phi\left(\sum_{j=1}^{m} \nu_j h_j(\mathbf{X}) + \nu_0\right),$$

where m is the number of neurons in the hidden layer and the ν_j and ν_0 are similarly weights and a bias term, respectively.

2.2.2 Setting Up the Adversarial Attack Problem

In Papernot et al.,[3] a general recipe for an adversarial attack is given along with a more specific attack, known as a saliency attack, aimed at fooling deep neural networks. The adversary starts with a legitimate input sample $\mathbf{X}$ whose true classification is $F(\mathbf{X}) = \mathbf{Y}$. The aim of the adversary is to craft an adversarial sample $\mathbf{X}^*$ that will go unnoticed (i.e., be interpreted as $\mathbf{X}$) by a human overseeing the process but that will be misclassified as $\mathbf{Y}^*$. Formally, we write the problem as

$$\arg\min_{\delta_{\mathbf{X}}} ||\delta_{\mathbf{X}}|| \text{ s.t. } \mathbf{F}(\mathbf{X} + \delta_{\mathbf{X}}) = \mathbf{Y}^*. \tag{1}$$

Note that finding the exact solution to this problem may be not be feasible due to the generally complex form of $\mathbf{F}$. Existing approaches rely on approximations; we detail one such approximation called the Jacobian saliency map (JSM).[3] The JSM is constructed and $\mathbf{X}$ is iteratively modified by choosing the pixel or pixels at the maximum point in the saliency map. The saliency map is defined to be

$$S(\mathbf{X}, t)[i] = \begin{cases} -\alpha\beta, & \text{whenever } \alpha > 0 \text{ and } \beta < 0, \\ 0, & \text{otherwise,} \end{cases}$$

where $\alpha = \dfrac{\partial \mathbf{F}_t(\mathbf{X})}{\partial \mathbf{X}_i}$ and $\beta = \displaystyle\sum_{j \neq t} \frac{\partial \mathbf{F}_j(\mathbf{X})}{\partial \mathbf{X}_i}$.

Intuitively, this method looks to modify those components of $\mathbf{X}$ that would most rapidly increase the probability of the target class (i.e., the adversary's target class) and decrease the probability of the non-target classes.

2.3 Optimization Opportunities

Papernot has identified a number of instances where the intelligent use of optimization could help to improve adversarial machine learning techniques (i.e., to create more effective attacks and/or defenses).[4]

- White-box attacks (where the adversary not only has access to the type of classifier being implemented by the defense, but also has access to either the data, the parameters of the classifier, or both) can be improved if we can find better heuristics (relative to, for example, the JSM described previously) for solving Eq. 1. The JSM is a greedy algorithm, but other authors have used Limited Memory Broyden–Fletcher–Goldfarb–Shanno or Adam optimization algorithms. Fundamentally, these are all approximations to solve Eq. 1. The key question is, can we do better? That is, can we 1) find perturbations of even *smaller* norm and/or 2) find such perturbations *faster*?
- With black-box attacks (where the classification method and its parameters are hidden from the adversary), which is perhaps the more realistic adversarial setting, we still wish to find vulnerabilities in the machine learning model $\mathbf{F}$. We are only able to query the model $\mathbf{F}$ and cannot know its internal structure/parameters. In such a situation, the adversary does not know $\mathbf{F}$ but still wishes to solve Eq. 1. A typical approach is for the adversary to build a surrogate $\hat{\mathbf{F}}$ to $\mathbf{F}$ based on querying the model (i.e., based on input/output queries of $\mathbf{F}$). Can we use optimization to either 1) minimize the number of such queries and/or 2) discover the optimal input locations $\mathbf{X}$ where the queries should be made?
- Papernot also identified Bayesian optimization as a potential method to reduce risk and maximize learning associated with each query made to the black-box model.[4]

3. Stochastic Optimization with Distributed Agents

3.1 Application Background

Due to increasingly complex, congested, and contested tactical communication environments, the Army has a keen interest in exploring alternative communication modalities. Augmenting Army networks with alternative communication capabilities has the potential to increase the resilience of the network, allowing the Army

to maintain connectivity even when conventional systems are inappropriate or ineffective. One such alternative modality involves the use of deep UV light (wavelengths of 200–300 nm) for communication. At these wavelengths, increased atmospheric scattering of transmitted (signal) radiation and the increased atmospheric absorption of solar (noise) radiation enables the use of extremely sensitive photon-counting receivers (i.e., photomultiplier tubes) to establish novel non-line-of-sight optical links through the atmospheric scattering channel. The unique features of this modality and channel have sparked significant research effort into UV communications (UVC) in recent years.[5]

Much of the existing literature has focused on point-to-point links and, in particular, understanding and modeling the point-to-point communication channel. Of the proposed modeling approaches, the most general involves the Monte Carlo simulation of photon propagation through the atmosphere.[6] Combining path loss estimates from such a model with a Poisson counting noise model allows one to explore a variety of point-to-point system design issues (e.g., He et al.[7]). Experimental measurement systems have been used for validation of this prior research (e.g., Liao et al.[8]).

Lacking from the existing literature is a deep understanding of the multiuser UVC scenario, such as a network of UVC systems operating in the same vicinity. To enhance throughput in a multiuser environment, the atmospheric channel can be spatially multiplexed so that different regions of the sky are predominately utilized by different communication links. Using spatial multiplexing, the optimization of the pointing direction (i.e., the azimuth and elevation angles) of each transmitter and receiver becomes crucial to proper operation, since even a single transmitter pointing suboptimally (perhaps in a greedy attempt to optimize the particular link that it is trying to establish) has the potential to inject an unacceptable level of noise into many of the other network links, resulting in their disconnection and, possibly, complete network failure. Of interest, for example, is both the ability to perform offline optimization based on existing (Monte Carlo) channel models in order to gain insight and further theoretical UVC networking research, as well as online algorithms that could enable the real-time adaptation of UVC nodes based on channel measurements.

The complexity of optimizing the pointing directions in a UVC network is a key challenge, particularly as one considers increasingly large networks. Even for small

networks, the Monte Carlo evaluation of the channel model for each configuration of parameters that is examined during an optimization routine can be prohibitive. For example, the application of conventional optimization approaches to a simple "network" of two transceiver nodes attempting to form a full-duplex link proves computationally challenging.[9] One approach to amelioriating this issue involves combining channel model precomputation and interpolation with an analytical model for range dependency,[10] but this approach is approximate with little guarantee as to the quality of the approximation.

Another key optimization challenge is that whether an optimization algorithm is designed for use with channel estimates from models or measurements, the resulting estimates may be quite noisy. In the case where a channel model is employed, this noise may be well modeled as Gaussian (since an extremely large number of simulated photons are usually considered), while estimates based on measurements may require the consideration of Poisson counting noise. In either case, there exists a potentially important tradeoff between the frequency of estimates available to an optimization routine (as dictated by the number of simulated photons in a Monte Carlo model or the length of time used to measure the communication channel) and the magnitude of the noise in the estimates.

The objective function itself may have a complicated form that captures some notion of global network performance. Each node may have direct access only to channel measurements associated with (signal and/or interference) links at that node, and it might not be possible to isolate interference signals (i.e., it might be possible to measure only the cumulative interference from all network nodes to a given node for a particular configuration). Sharing of channel measurements might enable the global computation of the network performance objective function, but this incurs overhead that should be accounted for in a real-time adaptation algorithm. However, despite the desire to limit communication overhead, the fact that, as previously noted, a single suboptimally pointed transmitter can have a drastic effect on the overall network performance implies a delicate balance between the low-complexity distributed action of individual nodes and their joint cooperation in global optimization. Another distributed aspect of the optimization problem is that each node may have direct control and perfect knowledge only of the pointing directions associated with the links that are formed to or from that node. Knowledge of the configuration of the rest of the parameter space again could be shared at a cost of communication overhead that should be accounted for.

3.2 Optimization Framework

Consider the case where there are N spatially distributed agents in a system. Each agent i ($i = 1,\ 2, \ldots,\ N$) has direct control only over a subset of parameters and may only be able to observe a part of the overall loss function. In what follows, we make this more concrete.

Let $\boldsymbol{\theta}_i$ be a vector of length $M_i \geq 1$; agent i can only directly control $\boldsymbol{\theta}_i$. Let $\boldsymbol{\theta} = [\boldsymbol{\theta}_1^T\ \boldsymbol{\theta}_2^T \ldots \boldsymbol{\theta}_N^T]^T$. Our goal is to minimize a global loss function $L(\boldsymbol{\theta}) = L(\boldsymbol{\theta}_1, \ldots, \boldsymbol{\theta}_N)$, where the loss function is, for example, a sum of many component loss functions $L(\boldsymbol{\theta}) = \sum_j L_j(\boldsymbol{\theta})$. Importantly, any given node may only be able to observe some of the L_j; in other words, some of the constituent loss functions may be unknown to any given node. In the process of minimizing $L(\boldsymbol{\theta})$, we may seek to minimize the frequency (equivalently number) of parameter updates made by each agent. Also, by sharing measurements of $L_j(\boldsymbol{\theta})$ among nodes, a global estimate of $L(\boldsymbol{\theta})$ could be constructed, but such sharing 1) incurs overhead cost that one might wish to account for and 2) might be prohibitive when considering distributed optimization algorithms.

Some particular wrinkles to the general problem that may be relevant include the following:

- The optimization could proceed "one agent at a time". That is, each agent in turn and according to some fixed (e.g., round-robin) or adaptive schedule updates its parameter set, with a noisy measurement of L taken after each individual update. It is quite possible that some agents will have a much greater influence over L than others (e.g., a node in the center of a communications network might affect L greatly, while one on the periphery might have less influence).
- The noise process could be non-Gaussian (e.g., Poisson) noise.
- The loss function L may exhibit threshold effects that result in null measurements. For example, we may be unable to measure $L(\boldsymbol{\theta})$ for some configurations of $\boldsymbol{\theta}$ when the signal-to-noise ratio (SNR) is extremely degraded.
- The actual goal may not be to minimize L but to reduce L below some prespecified threshold (e.g., to ensure some minimum SNR).

- L could actually be time-varying (i.e., the communications channel could change due to weather conditions or due to the mobility of nodes). In this case, understanding how well an optimization can track the changing channel provides insight into how fast one can allow the channel to change (e.g., how quickly nodes can move spatially) without breaking the network. Also, given a particular rate of change for the channel, this insight could inform choices in trading off optimization step size with error at convergence in order to ensure a sufficient rate of adaptation.

4. Spectrum-Sensing Multiobjective Optimization

Access to the electromagnetic spectrum is an ever-growing challenge for radar. Future radar will be required to mitigate RF interference from other RF sources, relocate to new frequency bands while maintaining quality of service, and share frequency bands with other RF systems. The spectrum-sensing multiobjective optimization (SS-MO) technique was recently investigated as a possible solution to these challenges.[11–16] SS-MO first generates a power spectrum based on a passive sensing of the electromagnetic environment (EME). This power spectrum estimates the noise and radio frequency interference in-band to the radar. Multiobjective optimization is then used to process the power spectrum to determine a sub-band (within the operating band of the radar) that jointly maximizes two conflicting objective functions: the radar range resolution and the SINR. Preliminary results indicate that the proposed technique has the capability to significantly increase SINR and maintain range resolution requirements.[11]

In this development we assume that the radar attempts to access the wideband frequency band B. Ideally, the radar would use the full band to enhance range resolution; however, access to this band could degrade the SINR if B is occupied by other RF systems. Alternatively, the radar can implement the SS-MO technique. A block diagram of the SS-MO technique is illustrated in Fig. 1. A passive sensing of the EME is first used at the start of the radar coherent processing interval (CPI). A power spectrum is then estimated and denoted as $\Theta = \{\theta_1, \theta_2, \ldots, \theta_N\}$ of size N for frequencies $F = \{f_1, f_2, \ldots, f_N\}$ with frequency resolution Δ_r. The SINR objective function is next formed to estimate the interference and noise in all possible

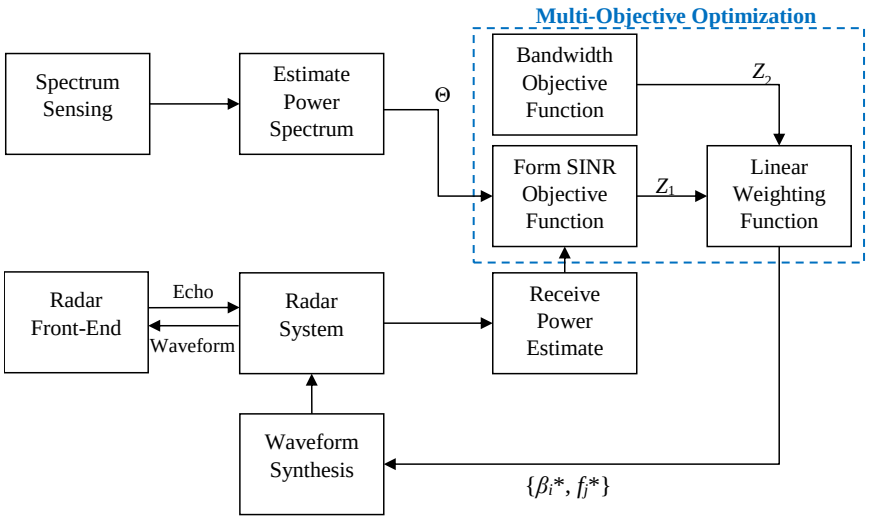

Fig. 1 SS-MO technique for radar. SS-MO requires an empirical measure of the spectral environment. Multiobjective optimization is used to find a sub-band for radar operation.

sub-band combinations and is calculated recursively as

$$\Gamma(\beta_i, f_j) = \begin{cases} \theta_j, & \text{if } i = 1 \text{ and } j = 1, 2, \ldots, N \\ \Gamma(\beta_1, f_j) + \Gamma(\beta_1, f_{1+j}), & \text{if } i = 2 \text{ and } j = 1, 2, \ldots, N-1 \\ \Gamma(\beta_{i-1}, f_j) + \Gamma(\beta_1, f_{i+j-1}), & \text{if } i = 3, 4, \ldots, N \text{ and} \\ & \qquad j = 1, 2, \ldots, N-i+1, \end{cases} \tag{2}$$

where $f_j \in F$ is the start frequency (the band edge), $\beta_i = i\Delta_r$ is the bandwidth of the ith sub-band, $i = \{1, 2, \ldots, N\}$, $j = \{1, 2, \ldots, N\}$, and there exists a total of $\bar{N} = \sum_{k=1}^{N}(N - k + 1) = (N^2 + N)/2$ possible sub-band combinations.

Figure 2 is used to illustrate the formation of Eq. 2 for $N = 5$ with $\{\theta_1, \theta_2, \ldots, \theta_5\}$, $\{f_1, f_2, \ldots, f_5\}$, and $\bar{N} = 15$ (the sub-bands available for evaluation). The start frequency (band edge) is shown below each cell and the bandwidth (β_i) is shown on the left. The formation of Eq. 2 requires evaluating a contiguous set of frequencies. The top row in Fig. 2 contains the power estimates with $\beta_1 = \Delta_r$ (i.e., the power spectrum resolution). The next row is formed by summing two sequential power estimates from the top row and has $\beta_2 = 2\Delta_r$ (two power estimates are used). Without loss of generality, the bottom row is formed by summing all power estimates so that $\beta_5 = 5\Delta_r = B$.

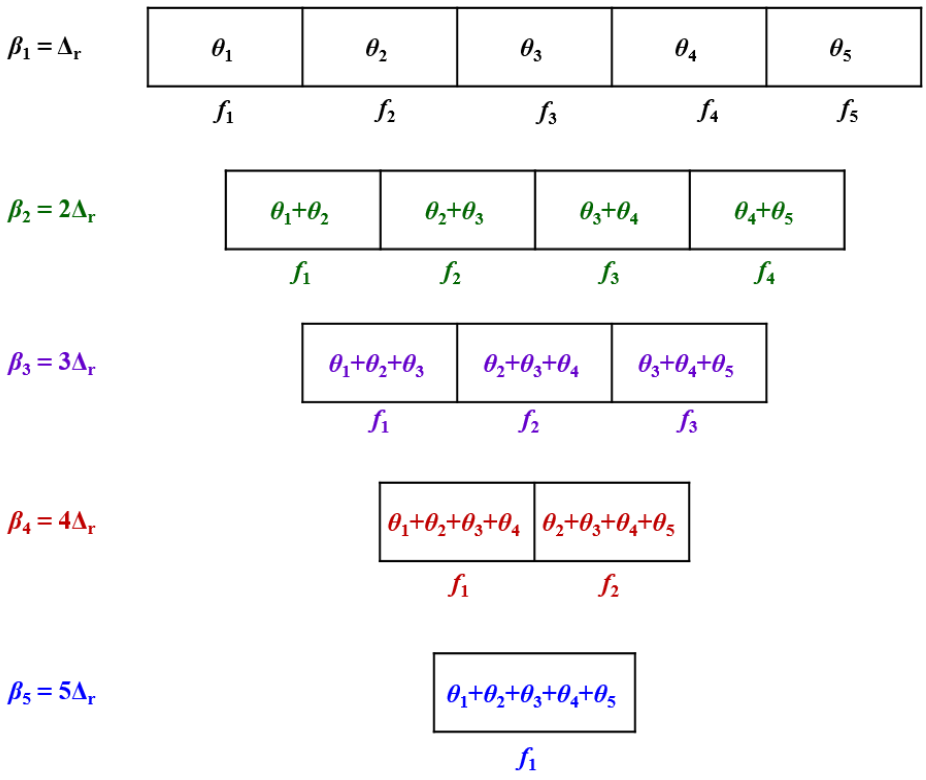

Fig. 2 Formation of Eq. 2 requires combining contiguous power spectral estimates

The receive power is next estimated using the return pulses (the echos) of the CPI of the radar. The radar provides feedback (Fig. 1) to SS-MO with updated estimates of the target range, R, and the target radar cross section, σ. The receive power is modeled using the radar range equation and is defined as

$$P_r = P_t G^2 \lambda^2 \sigma N_P / [(4\pi)^3 R^4], \tag{3}$$

where G is the antenna gain (for both transmit and receive antennas), P_t is the peak transmit power, λ is the wavelength of the waveform, and N_P is the number of pulses within a CPI. The SINR objective function is then defined as

$$Z_1(\beta_i, f_j) = P_r \tau \beta_i / \Gamma(\beta_i, f_j), \tag{4}$$

where τ is the pulse width and $\tau\beta_i$ as the time-bandwidth product of the pulse compression waveform. The second objective function is defined as

$$Z_2(\beta_i, f_j) = \beta_i, \tag{5}$$

and contains $\bar{N}$ elements representing the bandwidth of each sub-band. The goal of the SS-MO approach is to simultaneously maximize both Eqs. 4 and 5, but the presence of β_i in both equations sets up a fundamental conflict. A maximum bandwidth requires that $\beta_i = B$, but this results in a higher likelihood of encountering interference and thereby reduces the SINR.

Multiobjective optimization is used to resolve the conflicting objectives in Eqs. 4 and 5. The optimization is defined using the following linear weighting function:[14]

$$Z(\beta_i, f_j) = \alpha \acute{Z}_1(\beta_i, f_j) + (1 - \alpha)\acute{Z}_2(\beta_i, f_j), \tag{6}$$

where $0 \leq \alpha \leq 1$ is a user-defined weighting parameter. The function $\acute{Z}_1(\beta_i, f_j) = Z_1(\beta_i, f_j)/\max[Z_1(\beta_i, f_j)]$ is the normalized objective function of $Z_1(\beta_i, f_j)$, and the function $\acute{Z}_2(\beta_i) = Z_2(\beta_i, f_j)/\max[Z_2(\beta_i, f_j)]$ is the normalized objective function of $Z_2(\beta_i)$.

The solution to Eq. 6 is defined as

$$(\beta_i^*, f_j^*) = \underset{(\beta_i, f_j)}{\arg\max}[Z(\beta_i, f_j)], \tag{7}$$

subject to

$$Z_1(\beta_i, f_j) \geq Z_{1,\min} \tag{8}$$

and

$$Z_2(\beta_i, f_j) \geq Z_{2,\min}, \tag{9}$$

where $Z_{1,\min}$ is the minimum SINR requirement and $Z_{2,\min}$ is the minimum bandwidth requirement. The optimal solution is then defined as

$$Z^* = Z(\beta_i^*, f_j^*). \tag{10}$$

Prior investigations have examined SS-MO for small N, resulting in a small solution space that is easily solved by the brute force approach discussed previously. It is possible that wideband, high-resolution radar applications will require more advanced multiobjective optimization solutions. The need for these solutions is increased by the need to add additional objective functions to multiobjective optimization (signal-to-clutter ratio, Doppler, interference-to-noise ratio, etc.). Future research can include investigations into these optimization solutions. A key aspect

of this investigation would be to analyze the computational complexity for all proposed optimization methods, a needed analysis for real-time radar operations. For example, assuming system designs requiring fixed timelines (constant CPI), the processing time of the passive sensing and optimization results in a reduced number of pulses on target, which consequently decreases the SINR. New multiobjective optimization methods with low computational complexity would be needed to accurately estimate parameters within the sensing cycle of the radar.

5. Conclusion

This report presents a wide array of research areas that involve optimization of noise-corrupted loss functions over a potentially high-dimensional parameter space. These topics include adversarial machine learning, distributed agents, and spectrum sensing for radar. Adversarial machine learning involves deep neural networks, where we seek to solve a stochastic optimization problem subject to a minimum norm constraint. The distributed agents research area requires optimization of a global loss function, where we seek methods to minimize this loss function (i.e., minimize an aggregated set of many local loss functions). Spectrum sensing for radar requires solutions for multiobjective optimization to determine the best performance trade-off available. Our goal, in the near future, is to apply the tools and techniques of stochastic optimization to each of these research areas.

6. References

1. Bulo S, Biggio B, Pillai I, Pelillo M, Roli F. Randomized prediction games for adversarial machine learning. 2016. arXiv:1609.00804v1.

2. Papernot N, McDaniel P, Jha S, Fredrikson M, Celik Z, Swami A. The limitations of deep learning in adversarial settings. Proceedings of the IEEE European Symposium on Security & Privacy. 2016.

3. Papernot N, McDaniel PD, Goodfellow IJ, Jha S, Celik ZB, Swami A. Practical black-box attacks against deep learning systems using adversarial examples. 2016. arXiv:abs/1602.02697.

4. Papernot N. Personal communication, 2017 Dec 4 and 8.

5. Drost RJ, Sadler BM. Survey of ultraviolet non-line-of-sight communications. Semiconductor Science and Technology. 2014;29(8):084006.

6. Drost RJ, Moore TJ, Sadler BM. UV communications channel modeling incorporating multiple scattering interactions. Journal of the Optical Society of America A. 2011;28(4):686–695.

7. He Q, Sadler BM, Xu Z. Modulation and coding tradeoffs for non-line-of-sight ultraviolet communications. In: Proceedings of SPIE; Vol. 7464; p. 74640H–1–74640H–12.

8. Liao L, Drost RJ, Li Z, Lang T, Sadler BM, Chen G. Long-distance non-line-of-sight ultraviolet communication channel analysis: experimentation and modelling. IET Optoelectronics. 2015;9(5):223–231.

9. Drost RJ, Moore TJ, Sadler BM. Monte-Carlo-based multiple-scattering channel modeling for non-line-of-sight ultraviolet communications. In: Proceedings of SPIE; Vol. 8038; p. 803802–1–803802–9.

10. Drost RJ, Moore TJ, Sadler BM. Ultraviolet scattering propagation modeling: analysis of path loss versus range. Journal of the Optical Society of America A. 2013;30(11):2259–2265.

11. Martone A, Sherbondy K, Ranney K, Dogaru T. Passive sensing for adaptable radar bandwidth. 2015 IEEE Radar Conference (RadarCon). 2015. pp 0280–0285.

12. Martone A, Ranney K, Sherbondy K. Genetic algorithm for adaptable radar bandwidth. 2016 IEEE Radar Conference (RadarConf). 2016. pp 1–6.

13. Martone A, Dietlein C, Govoni M, Sherbondy K, Pulskamp J. Tuning technology for adaptable radar bandwidth. 2016 IEEE MTT-S International Microwave Symposium (IMS). 2016. pp 1–3.

14. Martone A, Ranney K, Sherbondy K, Gallagher K, Blunt S. Spectrum allocation for noncooperative radar coexistence. IEEE Transactions on Aerospace and Electronic Systems. 2018;54(1):90–105.

15. Martone A, Gallagher K, Sherbondy K, Hedden A, Dietlein C. Adaptable waveform design for enhanced detection of moving targets. IET Radar, Sonar & Navigation. 2017;11(10):1567–1573.

16. Ravenscroft B, Owen JW, Jakabosky J, Blunt SD, Martone AF, Sherbondy KD. Experimental demonstration and analysis of cognitive spectrum sensing and notching for radar. IET Radar, Sonar & Navigation. 2018;12(12):1466–1475.

List of Symbols, Abbreviations, and Acronyms

ARL	US Army Research Laboratory
CPI	coherent processing interval
EME	electromagnetic environment
JSM	Jacobian saliency map
RF	radio frequency
SINR	signal-to-interference-plus-noise ratio
SNR	signal-to-noise ratio
SS-MO	spectrum-sensing multioptimization
UV	ultraviolet
UVC	ultraviolet communications

1 (PDF)	DEFENSE TECHNICAL INFORMATION CTR DTIC OCA
2 (PDF)	DIR ARL IMAL HRA RECORDS MGMT RDRL DCL TECH LIB
4 (PDF)	ARL RDRL CIN D MJ WEISMAN RJ DROST G VERMA RDRL SER U AF MARTONE