

Fully cognitive transceiver for High Frequency (HF) applications

Noel Teku^a, Hamed Asadi^b, Tamal Bose^a, and Michael Marefat^a

^aUniversity of Arizona Dept. of Electrical and Computer Engineering, 1230 E Speedway Blvd,
Tucson, AZ 85721, USA

^bFord Motor Company, Dearborn, MI 48124, USA

ABSTRACT

Ionospheric conditions are variable in nature and can cause destructive interference to transmissions made in the High Frequency (HF) band, which ranges from 3-30 MHz. This poses a problem as the HF band is a critical frequency range for various applications (i.e. emergency, military). To manage these dynamic conditions, intelligent techniques should be implemented at the transmitter and receiver to properly maintain reliable communications. In this paper, we present work deriving components of a cognitive HF transceiver with agents called cognitive engines (CEs) operating at the transmitter and receiver. At the transmitter, cognition is employed to determine the combination of modulation and coding techniques that maximize throughput. At the receiver, cognition is implemented to derive the best parameters for equalization (i.e. tap length, step size, filter type, etc.) Results are presented showing that the individual components are able to satisfy their objectives. A discussion is also provided surveying recent research efforts pertaining to the development of cognitive methods for the Automatic Link Establishment (ALE) protocol, a common networking methodology for HF stations.

Keywords: HF, Ionosphere, Reinforcement Learning, Cognitive Engine, Equalization

1. INTRODUCTION

The High Frequency (HF) band, the portion of spectrum from 3-30 MHz, has provided benefits to different emergency and military applications. The main appeal of the band has been the use of the ionosphere as the medium for signal transmission, reducing the need for additional equipment while still maintaining long-range communications. However, due to the frequent dynamic nature of the ionosphere, transmissions can often be distorted by factors outside of the user's control. As a means of compensating for these effects, different forms of equalization have been studied to be utilized in recovering signals transmitted in the HF band. However, due to frequent channel fluctuations, the optimal equalizer configuration is usually unknown. Likewise, on the transmitter side, it's possible that a specific configuration of transmission parameters (i.e. modulation, coding) may be suboptimal compared to other configurations based on the current channel conditions.

In this paper, we present a tutorial on the development of components for a cognitive transceiver to be used in the HF band. The cognition aspect is implemented at the transmitter and receiver by using a cognitive engine (CE), which is "an intelligent agent which observes the radio environment and chooses the best communication settings that meet the application's goal."¹ The cognitive engine on each side has a specific objective that it must achieve. In this work, we explain how these components can function individually to provide a means of achieving more intelligent communications in the HF band. In Section 2, we provide background on the cognitive engines used in this work. In Section 3, the experimental setup and results for verifying the transmitter design is detailed. In Section 4, we discuss the concept of cognitive equalization that's employed at the receiver and present results. In Section 5, we provide a brief overview of the Automatic Link Establishment (ALE) protocol used to establish

Further author information:

Noel Teku: E-mail: ntekul@email.arizona.edu, Telephone: 1 520 243 9176

Hamed Asadi: E-mail: hamed.asadi@ieee.org, Telephone: 1 313 845 5961

Tamal Bose: E-mail: tbose@email.arizona.edu, Telephone: 1 520 621 6193

Michael Marefat: E-mail: marefat@email.arizona.edu, Telephone: 1 520 621 4852

networks in the HF band, as well as efforts in the literature used to incorporate cognitive/spectrum sensing techniques to increase the effectiveness of the protocol. In Section 6, the work in this effort is summarized and future directions are presented.

2. COGNITIVE ENGINES

A cognitive engine (CE) is an intelligent agent that enables a radio to make decisions about communication configurations based on its experiences with an environment. In Refs. 1–5, CEs have been used to strategically search through the large search space of parameters to identify and use the most effective communication settings and parameters. To enable this functionality, one avenue through which CEs have been implemented is reinforcement learning, an approach that is based on a balance of exploration and exploitation. Exploration refers to selecting a new option from the search space, while exploitation refers to choosing the option that has shown the best performance out of the options already explored.² Rewards are “a measure of success”,⁶ used to characterize the performance of each of the configurations in the search space. The reward is usually based on the objective of the CE. In previous works (see Refs. 1–5), the objective of the CE has been to maximize throughput/goodput. It’s assumed that the transmitter can obtain instantaneous feedback from the receiver detailing the goodput measurement based on the specific configuration used. The CE at the transmitter can then update the rewards associated with an option.

In this work, three reinforcement learning algorithms are utilized based on three different strategies: ϵ -greedy, Softmax Strategy, and Gittins strategy. The ϵ -greedy algorithm stipulates that an option is explored with probability ϵ and exploited with probability $1 - \epsilon$.⁶ A downside of this algorithm is that based on the choice of ϵ , the algorithm may not always find the optimal configuration. To provide a better balance between exploration and exploitation, an annealing version of the ϵ -greedy algorithm based on a schedule is developed, where the value of ϵ is initially set higher (to encourage more exploration), and decremented as the algorithm progresses (to encourage more exploitation).¹ In addition, we also utilize a version of the ϵ -greedy algorithm with a-prior information, as will be explained in Section 3. Softmax Strategy handles the issue of exploration/exploitation differently. It uses a probability value for selecting an option, and this probability is based on the past performances of the option (transceiver settings and parameters), i.e. the more rewards an option obtains the more likely it will be selected for usage.⁶ The probabilities are calculated using a softmax function as shown below:

$$P_k(x) = \frac{e^{\overline{\mu}_k/T}}{\sum_i e^{\overline{\mu}_i/T}} \quad (1)$$

where k represents an option, $\overline{\mu}_k$ represents the average reward obtained by an option, and T represents the temperature parameter. T is responsible for determining how much exploration/exploitation is conducted; a large choice of T causes more exploration, but a smaller choice of T causes more exploitation.⁶ The Gittins strategy calculates a Gittins index for each option as shown below,⁷ with the objective being to select the option with the highest Gittins index:

$$v_k(\pi_0) = \sup_{N>0} \frac{E\{\sum_{n=0}^{N-1} \gamma^n R_k(n) | \pi_k(1) = \pi_0\}}{\sum_{n=0}^{N-1} \gamma^n | \pi_k(1) = \pi_0} \quad (2)$$

where π_k represents the belief state of an option k at a specific time step, γ represents the discount factor, $R_k(n)$, denotes the reward and N represents the stopping time. The belief state “represents our knowledge about the underlying reward distribution at a time step n ... [and consists of] estimates of the mean μ_k and the standard deviation σ_k , using n' samples, of the underlying reward process.”¹ It was assumed in this implementation of the Gittins strategy, that the rewards were governed by a normal distribution. More information regarding the Gittins strategy can be found in Refs. 7,8.

Additionally, at the transmitter side, the CE is also able to differentiate between eligible and ineligible methods. The eligibility of a method is determined by first calculating the lower and upper bound estimates of the average reward distribution using the following equations with the t distribution:⁹

$$R_{lk}/R_{uk}(n) = \overline{\mu_k}(n) \mp (t(\frac{1-C}{2}, n' - 1) * \frac{\overline{\sigma}}{\sqrt{n'}}) \quad (3)$$

In this work, the t distribution is assumed if the option has been selected for no longer than 30 iterations. However, after 30 iterations, we then assume a Gaussian distribution. The following equation is used to check if an option is now ineligible:¹

$$R_{uk}(n) < \operatorname{argmax}_{j[1,K]} \overline{\mu_j}(n) \quad (4)$$

The approach essentially classifies all eligible methods as those with an upper bound estimate greater than the highest average reward assuming K total options. Once a method is classified as ineligible it is no longer considered by the CE in the decision process. It's important to note that in this work, the process of verifying eligibility have been implemented for the ϵ -greedy and Gittins Strategy CEs at the transmitter (in section 3); this has not yet been implemented for the receiver, as discussed in Section 4.

3. TRANSMITTER

In this section, we describe the methodology for which a CE can be employed at the transmitter to select the optimal combination of modulation, inner, and outer codes. The objective of the CE at the transmitter was to maximize throughput. The CEs implemented at the transmitter were ϵ -greedy (the standard, annealing, and a-priori versions) and Gittins strategy. Results pertaining to the statistical performance of each of the CEs are presented, indicating that over time the CEs are able to learn the optimal transmitter configurations. The simulations for the transmitter design were implemented in Matlab.

3.1 Experimental Setup

Table 1 summarizes the modulations, inner and outer codes that were accessible to the CE. The spectral efficiency of these options were calculated using: $SE = k * IR * CR$, where k represents the modulation order of the chosen option, IR represents the inner code rate, and CR represents the outer code rate. The theoretical throughputs are used as a-priori information for the Gittins Strategy and the version of the ϵ -greedy with a-priori information. They were calculated using: $THR = SE * 200kHz$, where a bandwidth of 200 kHz is assumed in this work. In order to incorporate real-world behavior, an experiment utilizing a Universal Software Radio Peripheral (USRP) was conducted to determine the effectiveness of each of the configurations listed in Table 1. For each configuration, 100 packets were sent over-the-air using a USRP in the HF band, and the packet success rate (PSR) was stored. While PSRs were measured at different SNRs for each configuration, for the results presented in this work, the measurements used were conducted at an SNR of 45 dB. Subsequently, it was assumed in this work that the CEs only had access to one channel. At each time step, based on the configuration chosen, this PSR was used to simulate the channel conditions on a random vector of 30 elements, to determine the fraction of packets that were successfully transmitted. The *goodput* was then determined by applying this fraction to the theoretical throughput.

3.2 Results

We characterize the CEs using learning curves based on the statistical metrics of their performance in Figures 1 and 2, which were obtained over 1000 time steps per experiment and averaged over 300 experiments. Figure 1 shows the average goodput obtained by the Gittins strategy and the aforementioned versions of the ϵ -greedy algorithm. In these simulations, ϵ for the standard ϵ -greedy CE was set to 0.1, for the annealing version ϵ was initially set to 0.5 and decremented by 0.0005 at each time step until it was less than or equal to 0.001, and for the a-priori version ϵ was set to 0.0001. The discount factor of the Gittins strategy was set 0.9. Figure 1 shows that prior to the 34th time step, the standard and annealing versions of the ϵ -greedy algorithm outperform the ϵ -greedy with prior information and Gittins CEs. After the 34th time step, the ϵ -greedy with prior information and Gittins CEs begin to experience a spike in performance. Additionally, for this specific channel scenario (i.e. one channel with an SNR of 45 dB and bandwidth of 200 kHz), it appears that the spikes are as aggressive as each other. The ϵ -greedy with prior information CE, due to ϵ being set to 0.0001, has a low rate of exploration but a high rate of exploitation. This means that while it does have access to the theoretical rewards of each option, a number of iterations may be required before it selects the option with the highest theoretical reward.

Table 1. Modulations, Inner/Outer Codes available at CE from Ref. 10

Modulation	Inner Code	Outer Code
BPSK	None	None
QPSK	Conv-V27 (1/2 Rate)	Golay (24,12)
8-PSK	Conv-V27 (2/3 Rate)	ReedSolomon-M8
16-PSK	Conv-V27 (4/5 Rate)	Hamming (7,4)
DBPSK	Conv-V27 (5/6 Rate)	Hamming (12,8)
DQPSK		SECDED (22,16)
8-DPSK		SECDED (39,32)
4-ASK		SECDED (72,64)
16-QAM		
32-QAM		
64-QAM		

However, once it does, the CE will likely stick with this reward for the majority of the remaining time steps as evidenced through the spike in Figure 1. Figure 1 also indicates that the annealing ϵ -greedy is able to catch up to the two CEs near the 200th time step as it transitions to utilizing more exploitation than exploration as time progresses.




Figure 1. Average Goodput of ϵ -greedy Variations and Gittins Strategy; x-axis shows time steps and y-axis shows throughput (kbps). Indicates that ϵ -greedy with prior information and Gittins Strategy CEs experience a spike in performance after underperforming initially.

In Figure 2, the standard deviations obtained by the four CEs are shown. These values are calculated in reference to the means obtained at each time step as shown in Figure 1. It's also important to note that the annealing version of the ϵ -greedy CE initially has a larger standard deviation as it focuses more on exploring options, causing drastic variations in the mean calculation. However, as time progresses, the standard deviation is reduced because the CE focuses on exploiting the options, resulting in a smaller deviation with respect to its average goodput. The Gittins strategy and ϵ -greedy (with prior information) CEs have a very low standard deviation for a majority of the time steps, indicating that their decisions stay close with their average goodputs. This makes sense because due to having access to the prior information, they are exploiting the options with the highest rewards. It's worth noting that around the 34th time step, a spike occurs for both CEs. This is reflective of the spike shown similarly in Figure 1, where both of the CEs had a significant increase in throughput as a result of their decisions.




Figure 2. Standard Deviation of Goodputs obtained by ϵ -greedy Variations and Gittins Strategy; x-axis shows time steps and y-axis shows throughput (kbps). Indicates that ϵ -greedy with prior information and Gittins strategy CEs have a low standard deviation at all time steps except when they experience a performance spike.

4. RECEIVER

As stated in Section 1, a CE can also be used at the receiver to select the optimal parameters for equalization. In Ref. 11, we presented results of using an ϵ -greedy CE to vary the tap lengths and step size of a Least Means Square Decision Feedback Equalizer (LMS-DFE). In this work, we expand on this principle by enabling the CE to also have a choice of equalizer type and adaptive algorithm. The CEs implemented at the receiver were ϵ -greedy and Softmax Strategy with the objective of selecting the equalizer configurations that were effective in minimizing errors. The simulations for the equalizers and CEs were implemented in Matlab (See Refs. 6, 12–14) with the built-in International Telecommunication Union (ITU) HF channel models used for Mid-Latitude, Quiet and Disturbed Conditions.¹⁵ The delay and frequency spreads of the two channel models are shown in Table 2:

Table 2. ITU Channel Model Specifications from Ref. 15

Channel Condition	Delay Spread	Frequency Spread
Mid Latitude, Quiet (MQ)	0.5 ms	0.1 Hz
Mid Latitude, Disturbed (MD)	2 ms	1 Hz

4.1 Experimental Setup

The CE had the choice of Decision-Feedback Equalizer (DFE) and Decision-Directed Equalizer (DD) for filter types as well as the LMS and Recursive Least Squares (RLS) algorithms for adapting the weights of these filters. Additionally, the CEs had the choice of 10-60 feedforward taps (for both DFE/DD options) and 10-60 feedback taps (when an option with the DFE was chosen), with the choice of taps being restricted to increments of five (i.e. 10, 15, 20, ...). Lastly, the CEs had a choice of 0.01, 0.001, and 0.0001 for step sizes/regularization parameters based on usage of either the LMS/RLS algorithm respectively. For simplicity, it was assumed in the case of the DFE, that the feedforward and feedback filters were implemented with the same step size/regularization parameter for each configuration selected. It was also assumed that, when the RLS algorithm was chosen, a forgetting factor of 0.99 was used in all simulations. However, in future works, we plan to investigate if performance improvements can be obtained by using cognitive equalization to vary this as well. In addition, to simplify the generation of the sample space for the CE to choose from, we assume that each choice of DFE has the same number of feedforward and feedback taps (i.e. Configuration 1: 10 feedforward taps, 10 feedback taps, regularization parameter of 0.01; Configuration 2: 15 feedforward taps, 15 feedback taps, regularization parameter of 0.01; etc.).

For verifying the receiver design, it was assumed that only BPSK modulation was used at the transmitter. The CEs have no a-priori knowledge as to which equalizer settings are more effective. The objective of the cognitive engines is to minimize BER. Because BER is not feasible to calculate, the cognitive engines utilize

a metric called the average minimum distance. The minimum distances used to assign each incoming symbol is averaged over the length of the received sequence. This metric is then used as the reward the CE uses to characterize each equalizer configuration; subsequently, the CE's objective will be to select the configurations that minimize the average minimum distance. For each SNR of AWGN, 1000 iterations were completed, where an iteration refers to sending a training and testing sequence. The CEs were able to store the average minimum distance in each of these iterations, but to give it a chance to learn prior to measuring its performance, the actual probability of error was averaged over the last 500 iterations. The error of the fixed equalizer was also measured over the last 500 iterations to give it a fair comparison. The training and testing sequences each consisted of 500 symbols. For the results shown in this section, ϵ was set to 0.1 and T was set to 0.01.

4.2 Results

Figure 3 shows the BERs obtained with fixed (i.e. able to adapt weights but unable to change tap-length, step size, filter type, and adaptation algorithm) equalizer configurations as well as configurations chosen by the standard ϵ -greedy and Softmax Strategy CEs for the Mid-Latitude Quiet (MQ) channel model. The fixed configurations chosen were 10 taps (for both feedforward and feedback filters) and a step size/regularization parameter of 0.01. Figure 3 shows that the RLS-DFE is able to outperform the LMS-DFE at higher SNR, while the RLS-DD performs comparably to the LMS-DD. Additionally, Figure 3 shows that while the DFEs are able to outperform the CEs at low SNR, the Softmax Strategy CE is able to attain a similar performance to the RLS-DFE (which has the lowest error of the fixed configurations in Figure 3) at higher SNR. Additionally, the ϵ -greedy CE is able to perform comparably to the LMS-DFE at higher SNR, while also performing comparably/better than the RLS/LMS-DDs. This indicates that due to the small delay and doppler spreads of the MQ channel, 10 taps may be sufficient to recover the signal. Because the CEs have no prior information regarding which configurations are more effective, they may spend time exploring options with larger tap sizes, which may cause the equalizers to require more time to converge to the optimal weights. Figure 4 shows the performance of the CEs and fixed configurations when the Mid-Latitude Disturbed (MD) channel is used, where the CEs are shown to have outperformed the fixed LMS and RLS equalizers. This indicates that despite 10 taps being sufficient to handle the MQ channels small delay and doppler spread, because the MD channels conditions are relatively higher, more taps are required to overcome these impairments. Thus, while there are instances where fixed configurations can outperform the CEs, it's important to note that the CEs are starting with no prior information as to which configurations are optimal, and therefore considering this, the CEs seem to have a decent performance compared to fixed configurations.




Figure 3. Probability of Error of CEs and Fixed Equalizer Configurations in Mid-Latitude Quiet Conditions; x-axis shows SNR of AWGN (dB) and y-axis shows probability of error. Softmax-Strategy CE outperforms fixed LMS-DFE and attains similar performance to RLS-DFE (best-performing fixed configuration) at higher SNR.

5. IMPROVEMENTS TO ALE

The Automatic Link Establishment (ALE) protocol is used heavily in creating links between HF stations. A detailed explanation of 2G and 3G (i.e. second and third generation) ALE, can be found in Refs. 16,17. In




Figure 4. Probability of Error of CEs and Fixed Equalizer Configurations in Mid-Latitude Disturbed Conditions; x-axis shows SNR of AWGN (dB) and y-axis shows probability of error. ϵ -greedy and Softmax Strategy CEs outperform fixed equalizer configurations.

this work we summarize a few main distinctions between the two revisions for brevity. In the ALE protocol, HF stations have a finite number of frequencies on which a link can be established. The objective of the protocol is to determine which of these frequencies is most suitable for creating a link. A method common to both 2G and 3G ALE for this is Link Quality Analysis (LQA).¹⁷ This process involves transmitting sounding signals on each channel to acquire a score based on a specific metric (i.e. SNR) that characterizes its quality.¹⁶ The calling station then attempts to establish a link in order of the highest to the lowest LQA scores as needed. Another similarity between the protocols is the usage of “a form of Listen Before Transmit [LBT] in an effort to avoid interfering with on-going communications.”¹⁷ Once a frequency is chosen, 2G ALE utilizes a three-way handshake prior to transmitting data, which consists of the following steps:¹⁷

1. Transmitter initiates call to desired receiver.
2. Receiver sends acknowledgement confirming receipt of transmitter’s initial request.
3. Transmitter sends confirmation of receiver’s acknowledgement

In 3G ALE, however, the handshake process is reduced to the first two steps.¹⁷ In addition, 2G ALE systems are asynchronous, meaning that “a station attempting to communicate with a scanning receiver generally will not know when that receiver will dwell on any particular frequency”.¹⁶ The receivers scan the available channels, remaining on each for a specific amount of time, until it detects the transmitter’s linking call. Contrarily, 3G ALE systems can function either synchronously or asynchronously.^{16,17} The synchronous mode enables receivers to “change frequency at the same time, to within a relatively small timing uncertainty... The current dwell channel of every station in a network can always be computed from the time of day and the address of the desired station.”¹⁶ Subsequently, 3G ALE systems are able to attain smaller average linking times.¹⁷ Additionally, a 4G ALE protocol has been standardized fairly recently. It supports 2G and 3G ALE functionality, asynchronous and synchronous operation, and incorporates spectrum sensing capabilities (more information regarding 4G ALE can be found in Ref. 18). For the remainder of this section, we present a few fairly recent works in the literature using machine learning and spectrum sensing/prediction algorithms for frequency selection in HF.

Ref. 19 compares the usage of three different techniques to observe if improvements can be attained in ALE users’ ability to sense primary users: energy detection, matched filter detection, and cyclostationary feature detection. Results indicate that cyclostationary feature detection outperforms matched filter and energy detection; however, the authors note that all three methods enhance the effectiveness of sensing primary users in ALE.¹⁹ Ref. 20 develops a new method of obtaining LQA scores by incorporating statistical parameters about the channel and spectrum prediction. Two parameters called the average idle duration and probability are introduced and, along with SNR, are used to calculate LQA scores for each channel. Spectrum prediction is also used to give higher priority to idle than occupied channels. Results show that the technique is able to provide

smaller linking times.²⁰ Ref. 21 utilizes the Upper Confidence Bound (UCB) algorithm, a reinforcement learning technique, as a means of adaptively detecting/utilizing idle channels for transmission and is shown to be more effective than randomly selecting channels. The authors note that “[ALE Stations] only select the transmission channels that are expected to be available according to their internal ranking based on propagation characteristics, and they do not have any adaptive mechanism to move to another channel.”²¹ Hidden Markov Models (HMMs) have also been an algorithm investigated to predict channel occupancy for HF in Refs. 22, 23. The authors of Refs. 21, 23 combined their reinforcement learning and HMM approaches to form a hybrid method for spectrum prediction/frequency selection in Ref. 24. The authors implement the reinforcement learning aspect of the hybrid system using a modified version of the UCB algorithm UCB₁-M to choose multiple channels at a time instead of one. The authors note that the availability of the channels affects which of the two algorithms in the hybrid technique is used more, stating “that if there are plenty of available channels for transmission during a minute, the hybrid UCB-HMM system will switch to long-term transmission slots following the predictions of M HMMs working in parallel (M-HMMs)... Nevertheless, if the environment changes and most of the channels are unavailable or partially-available, it will transmit in a short-term slot following the decisions of the UCB₁-M algorithm.”²⁴

6. CONCLUSION/FUTURE DIRECTION

In this effort, we have shown initial designs for implementing a functioning cognitive transmitter and receiver for the HF band. We have shown that individually, the CEs at the transmitter and receiver are capable of meeting their objectives of maximizing throughput/minimizing BER respectively. We have also presented a short summary of the ALE protocol and some efforts to implement cognitive/spectrum sensing techniques for HF networks. One future direction we plan to take is to connect the transmitter and receiver designs together to construct a functioning HF transceiver. We plan to develop a prototype on software-defined radios. One of our future works involves creating a methodology based on which information about current ionospheric conditions can be used to eliminate configurations from the search space to reduce the time of the exploration/exploitation process. We also plan to investigate deep learning techniques to function as cognitive engines. In addition, we plan to implement more robust equalizer algorithms to provide smaller BERs.

ACKNOWLEDGMENTS

This project was partially supported by the Broadband Wireless Access and Applications Center (BWAC); NSF Award No. 1265960.

REFERENCES

- [1] Asadi, H., Volos, H., Marefat, M., and Bose, T., “Experience level analysis for a cognitive radio engine,” *Analog Integrated Circuits and Signal Processing* (May 2018).
- [2] Volos, H. I. and Buehrer, R. M., “On balancing exploration vs. exploitation in a cognitive engine for multi-antenna systems,” in [*GLOBECOM 2009 - 2009 IEEE Global Telecommunications Conference*], 1–6 (Nov 2009).
- [3] Volos, H. I. and Buehrer, R. M., “Cognitive engine design for link adaptation: An application to multi-antenna systems,” *IEEE Transactions on Wireless Communications* **9**, 2902–2913 (Sep. 2010).
- [4] Asadi, H., Volos, H., Marefat, M. M., and Bose, T., “Metacognitive radio engine design and standardization,” *IEEE Journal on Selected Areas in Communications* **33**, 711–724 (April 2015).
- [5] Asadi, H., Volos, H., Marefat, M. M., and Bose, T., “Metacognition and the next generation of cognitive radio engines,” *IEEE Communications Magazine* **54**, 76–82 (January 2016).
- [6] White, J., [*Bandit Algorithms for Website Optimization: Developing, Deploying, and Debugging*], O’Reilly Media (2012).
- [7] Gittins, J. C., [*Multi-armed Bandit Allocation Indices*], Wiley, Chichester, NY (1989).
- [8] Gittins, J. and Jones, D., “A dynamic allocation index for the sequential design of experiments,” *Progress in Statistics*, 241–266 (1974).

- [9] Pishro-Nik, H., [*Introduction to Probability, Statistics, and Random Processes*], Kappa Research, LLC (2014).
- [10] Asadi, H., *Robust Intelligent Agents for Wireless Communications: Design and Development of Metacognitive Radio Engines*, PhD dissertation, University of Arizona (2018).
- [11] Teku, N. and Bose, T., “Cognitive equalization for hf channels,” *International Telemetering Conference 2018* (2018).
- [12] Proakis, J., Salehi, M., and Bauch, G., [*Contemporary Communication Systems Using MATLAB*], Cengage Learning, 3 ed. (2012).
- [13] Haykin, S., [*Adaptive Filter Theory*], Pearson, 5 ed. (2014).
- [14] Diniz, P., [*Adaptive filtering: Algorithms and practical implementation*], Kluwer international Series in Engineering and Computer Science, Springer, 3 ed. (2008).
- [15] ITU-R Rec. F.1487, “Testing of HF modems with bandwidths of up to about 12 kHz using ionospheric channel simulators,” (2000).
- [16] Johnson, E., Koski, E., Furman, W., Jorgenson, M., and Nieto, J., [*Third-generation and Wideband HF Radio Communications*], Artech House mobile communications series, Artech House (2013).
- [17] Furman, W. N. and Koski, E., “Next generation ale concepts,” in [*The Institution of Engineering and Technology 11th International Conference on Ionospheric radio Systems and Techniques (IRST 2009)*], 1–5 (April 2009).
- [18] MIL-STD-188-141D, “Interoperability and performance standards for medium and high frequency radio systems,” U.S. Dept. of Defense (2017).
- [19] Shahid, A., Ahmad, S., Akram, A., and Khan, S. A., “Cognitive ale for hf radios,” in [*2010 Second International Conference on Computer Engineering and Applications*], **2**, 28–33 (March 2010).
- [20] Yu, L., Chen, J., Ding, G., and Qin, Z., “Fast automatic link establishment: A new metric and the value of spectrum prediction,” in [*2016 8th International Conference on Wireless Communications Signal Processing (WCSP)*], 1–6 (Oct 2016).
- [21] Melin-Gutierrez, L., Modi, N., Moy, C., Prez-lvarez, I., Bader, F., and Zazo, S., “Upper confidence bound learning approach for real hf measurements,” in [*2015 IEEE International Conference on Communication Workshop (ICCW)*], 381–386 (June 2015).
- [22] Lin, G., Cheng, Y., and Jiang, H., “A hmm-based spectrum access method used for hf communication,” in [*2015 8th International Symposium on Computational Intelligence and Design (ISCID)*], **2**, 426–431 (Dec 2015).
- [23] Melian-Gutierrez, L., Zazo, S., Blanco Murillo, J. L., Perez-Alvarez, I., Garcia-Rodriguez, A., and Perez-Diaz, B., “Hf spectrum activity prediction model based on hmm for cognitive radio applications,” *Physical Communication* **9**, 199211 (12 2013).
- [24] Melin-Gutierrez, L., Modi, N., Moy, C., Bader, F., Prez-lvarez, I., and Zazo, S., “Hybrid ucb-hmm: A machine learning strategy for cognitive radio in hf band,” *IEEE Transactions on Cognitive Communications and Networking* **1**, 347–358 (Sep. 2015).