Closing the Loop on Cognitive Radar for Spectrum Sharing

Anthony F. Martone[1], Jacob A. Kovarskiy[2], Charles E. Thornton[3], Benjamin H. Kirk[2], Jonathan W. Owen[4], Brandon Ravenscroft[4], Austin Egbert[5], Adam Goad[5], R. Michael Buehrer[3], Ram M. Narayanan[2], Shannon Blunt[4], Charles Baylis[5] and Kelly D. Sherbondy[1]

[1]U.S. Army Research Laboratory, Adelphi, MD

[2]Electrical Engineering Department, The Pennsylvania State University, University Park, PA

[3]Bradley Department of Electrical and Computer Engineering, The Virginia Polytechnic Institute and State University, Blacksburg, VA

[4]Radar Systems Lab, University of Kansas, Lawrence, KS

[5]Wireless and Microwave Circuits and Systems Program, Baylor University, Waco, TX 76798-7356

Abstract— **We investigate three emerging topics essential for the development of cognitive radar (CR) for spectrum sharing: the response time (RT) of the CR, the autonomous regulation of the perception-action cycle (PAC), and the regulation of cognition. The RT measures the latency of all algorithms / hardware and is examined with respect to enabling capabilities of software-defined systems for rapid flexibility and responsiveness. The autonomous regulation of the PAC determines "how fast the CR can interact with the environment" as well as "how fast the CR should interact with the environment." The regulation of the PAC is explored with respect to pulse-to-pulse waveform agility to coexist successfully with dynamic radio frequency (RF) emitters in the ambient electromagnetic environment (EME) and the consequence of modifying the waveform within the coherent processing interval (CPI). Finally, the regulation of cognition determines how to select a particular CR technique appropriately for a given dynamic environment. This selection requires a high-level, or meta-decision process to identify the appropriate cognitive radar technique as the EME changes over time and we therefore concentrate discussion on the newly emerging topic for radar called metacognitive radar. The exploration of these three topics include a review of past and current research with discussion of possible future research.**

I. Introduction

A current major challenge in the radar community is operationalizing cognitive radar (CR). Although multiple definitions and taxonomies of CR exist [1-5], the Institute of Electrical and Electronic Engineers (IEEE) has defined it as a "system that in some sense displays intelligence, adapting its operation and its processing in response to a changing environment and target scene. In comparison to adaptive radar, cognitive radar learns to adapt operating parameters as well as processing parameters and may do so over extended time periods [6]." The model commonly explored for adjusting CR operation in a changing environment and target scene is the perception-action cycle (PAC) [1,3], which is a closed-loop model that actualizes the behavior above. The PAC has served as a foundation to categorize CR techniques and to assess their level of cognition. The implementation of the PAC for CR has been explored for several applications [4, 7-18]. The definition of CR for future systems will continue to evolve in order to incorporate the unrecognized practical aspects of operationalizing models into emerging software-defined and intelligent systems for rapid flexibility and responsiveness [19]. The goal of this paper is to establish a framework of research topics that the radar community can explore to address the practical aspects of CR development, which motivates the integration of multiple, disparate CR techniques onto a flexible software-defined radar architecture. These CR techniques represent a toolbox of solutions that must be chosen autonomously and under uncertainty by a high-level decision process within dynamic environmental scenarios. One can make the case that operationalizing CR is subject to three factors:

- *System Response Time (RT):* A measure of latency (for some degree of accuracy) of all standard algorithms, CR techniques, hardware, and adaptive hardware components within the CR system. The RT is measured with respect to the rate the CR interacts with the environment via the PAC.
- *Autonomous PAC Regulation*: Determines the rate (i.e. how fast) the CR interacts with the environment thereby setting the appropriate PAC speed. PAC regulation also depends on how quickly the environment and target scene evolve over time.
- *Autonomous Regulation of Cognition*: Because a single CR technique or solution has advantages and disadvantages for different environmental conditions, multiple different CR techniques are needed for dynamic scenarios. Hence, there is no single optimal CR solution. The regulation of cognition requires a "meta" process to select the appropriate CR technique as needed. We therefore concentrate here on the newly emerging topic of metacognitive radar.

These three factors imply that a CR is not composed of a single PAC operating on a single timeline, but instead consists of a hierarchy of PACs with different response times and rates. For example, PAC models operating on shorter timescales (i.e. pulse-to-pulse adaptations) have the capability to respond quickly to events in the environment, while PAC models operating on longer timescales (scan-to-scan) have the capability to determine model failure in dynamic environments. These three factors are necessary elements for future CR models that require a hierarchy of PAC loops. This paper extends the discussion in [20], where the factors above were introduced, by providing further analysis, outlining potential hardware requirements and limitations, discussing solutions to challenges that occur when the PAC changes too fast or too slow relative to the environment, presenting a scalable model for the practical implementation of CR, and highlighting how metrics are impacted by this model. This model is intentionally generic to address several radar applications; however, our focus here is the spectrum sharing application in order to demonstrate implementation for autonomous, real-time radar (described in Section II). Section III then provides a discussion of the system RT for spectrum sharing, while Section IV discusses factors contributing to PAC regulation. Finally, Section V expands upon the metacognitive radar (MCR) model presented in [21] that integrates a wide variety of CR techniques into a single software-defined radar platform for multi-function capabilities, illustrating the implementation of the MCR model for a real-time radar tracking scenario.

II. Cognitive Radar Techniques for Spectrum Sharing

Spectrum is presently one of the most precious commodities. Government regulatory agencies throughout the world have taken initiatives for more effective spectrum sharing practices between commercial and government radio frequency (RF) systems [22-25]. The need for spectrum sharing is exacerbated by the emergence of the fifth generation (5G) of wireless communications technology, with 5G a paradigm shift from its fourth generation (4G) predecessor, offering the potential to revolutionize the way people and machines communicate and exchange information [26]. Regulators are proposing new policy practices that allocate additional frequency bands [27,28]. Similar practices originated with older communication technology [29] via the spectrum auction of the advanced wireless service 3 (AWS-3) bands and spectrum sharing between radar and unlicensed wireless devices. These changes will adversely affect government radar systems that have legacy rights to the spectrum and were not engineered to share the spectrum, thus necessitating CR solutions for efficient and effective spectrum access and utilization.

To address these spectrum sharing challenges for radar, we consider a coexistence strategy for dynamic spectrum access (DSA). In this context, coexistence implies that the radar avoids interfering with other wireless devices in the spectrum without use of a common signaling protocol [30]. This arrangement requires the radar to employ spectrum sensing to determine interference levels induced by other wireless devices and subsequently reconfigure the radar waveform to avoid mutual interference. This coexistence approach involves passive monitoring of the radar's operating band, with bandwidth B, over time to determine an appropriate radar frequency allocation that simultaneously optimizes signal to interference plus noise ratio (SINR) and bandwidth. The full bandwidth solution (i.e. radar operates over all of B) clearly provides the finest range resolution; however, selection of the full bandwidth incurs co-channel interference with other wireless devices, thus decreasing the radar's SINR while generating mutual interference. The more spectrally efficient alternative is to select a frequency sub-band, with bandwidth $\beta \leq B$ and center frequency f_{SB}, to minimize the expected interference [30]. In this development the passive monitoring procedure utilizes the front-end of the radar; hence, radar operations cease during monitoring. Specialized timing is required to interleave radar and sensing operations for efficient DSA [31].

A generalized block diagram of the DSA approach is illustrated in Fig. 1. The authors have explored this model for five different CR techniques, where each implements the PAC with different learning models and waveform types to determine the best radar performance for coexistence, resulting in different advantages and disadvantages based on the type of RF interference (RFI). Each technique forms a time sequence of power spectra based on passive measurements, with particular focus on the percentage of occupied spectrum in B and the randomness of this occupied spectrum (i.e. how often it changes). This spectrum monitoring approach is based on energy detection and the channel state is determined via threshold techniques. The authors chose to explore energy detection for its 1) low latency, a necessary requirement for real-time DSA, 2) scalability to multiple radar applications that occur in different frequency bands, and 3) capability to detect a wide variety of wireless signal types with minimal *a priori* information.

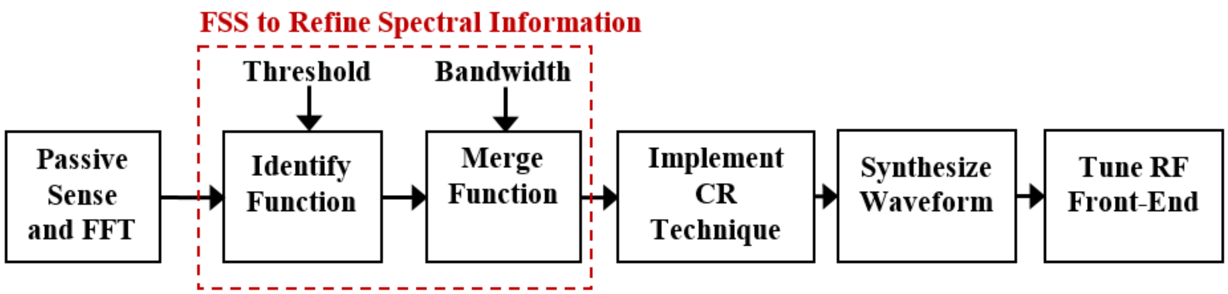

Fig. 1: Block diagram of the DSA approach that illustrates the general processing used for five different CR techniques.

The power spectra are next processed by the fast spectrum sensing (FSS) algorithm [30], which distills information about the power spectrum into representative metadata to reduce the input to a given CR technique. Prior work has shown that the latency of a multi-objective optimization decision process is significantly reduced (by orders of magnitude) using FSS. The power spectrum of Fig. 2 is used to illustrate the concept of FSS and was estimated by the fast Fourier transform (FFT) using the data in [30] with B = 100 MHz, frequency resolution Δ_B = 100 kHz, and N = 1000 samples. The first FSS function, "Identify," combines interference within the power spectrum into alternating low (green) and high (red) powered groups using an energy detection threshold that is estimated by training data and set to be 10 dB above the noise floor (-110 dBm in this case). Fig. 2 depicts 11 total groups, of which six are low-power and five are high-power. The second function, "Merge," then merges high-power groups into spectral clusters to be avoided depending on the waveform structure to be employed (e.g. determining the single largest contiguous bandwidth or via spectral notching).

The output of FSS is the metadata comprised of center frequencies, bandwidths, and power levels of the merged clusters. This metadata then informs the operation of the particular CR technique to be deployed. The techniques that have been explored thus far include Sense-React-Avoid (SRA), Sense-Predict-Avoid (SPA), Sense-Learn-Avoid (SLA), Sense-React-Notch (SRN), and Sense-Predict-Notch (SPN), with the spectral parameters determined by each used to synthesize a waveform for DSA.

The SRA technique uses the FSS metadata to quickly react to RFI by selecting the center frequency and bandwidth of the particular low-power group having the largest contiguous bandwidth [30,31]. The

SPA approach likewise seeks the largest bandwidth, but does so in an anticipatory manner by employing a stochastic model that represents the RFI as an alternating renewal process [32], with the RFI statistics learned online in real-time. Similarly, SLA is a reinforcement learning approach that treats the CR's waveform selection as a Markov decision process (MDP) for decision making under uncertainty [33], realizing a decision rule (or policy) that maximizes the expected value of a user-defined reward based on an application-specific combination of SINR and bandwidth. Recent research examines Deep Q-Learning, which can handle a larger problem space than the MDP algorithm, including additional memory states [34]. An adversarial bandit model has also been examined which maintains favorable performance even in the presence of intelligent interference sources (i.e. intelligent jammers) [35]. In [36] the SPA and SLA CR techniques were compared and found to yield similar performance in many stochastic scenarios.

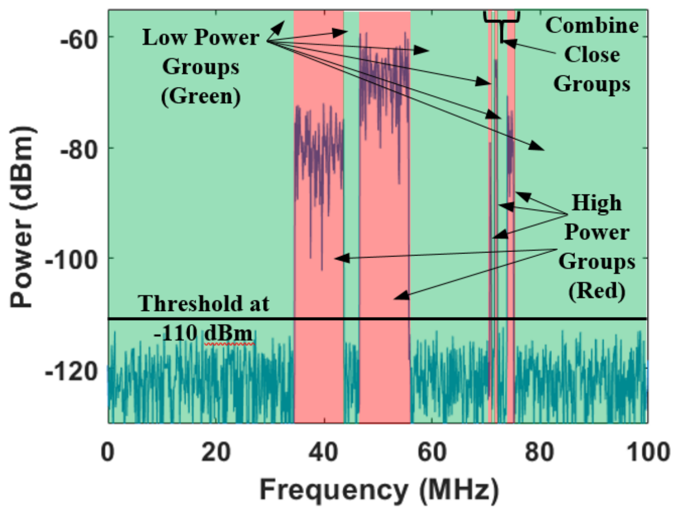

Fig. 2: Example spectrum with low and high power RFI. FSS refines the number of frequency bins handled by a decision process.

In contrast to the avoidance approaches that seek contiguous bandwidth, the SRN approach uses the FSS metadata to inform subsequent spectral shaping optimization that realizes instantiations of random FM (RFM) waveforms that are continuous, have a constant time-domain envelope, and possess an enforced spectral mask containing notches [37,38]. The main advantage of this approach is that the full bandwidth can be utilized, sans notched regions, though a higher computational cost is incurred for waveform generation. It has also recently been shown that SRN can likewise be combined with prediction [39] to facilitate the CR technique of SPN, which addresses real-time computational demand as well reducing the achievable latency in responding to RFI changes.

A summary of these CR strategies, all of which can adapt on a pulse-to-pulse basis, is provided in Table I. The RT of each method is denoted using T_{CR}, which measures the processing time (or latency) that is quantified according to the implementation strategy and hardware discussed in Section III. The Performance Metric denotes the means used to evaluate each CR technique, where the prime metrics of SINR and bandwidth are self-explanatory, while collisions (CO) and missed opportunities (MO) respectively indicate unintentional mutual interference and available spectrum not utilized by the radar. The Clutter Modulation performance metric evaluates the phenomenology (Section IV) that arises when sub-coherent processing interval (CPI) waveform adaptations occur, which causes pulse compression sidelobe and mainlobe modulation across slow-time that can mask true targets locations and increase false alarms. The Response Strategies are classified into reactive or predictive approaches, where the former signifies waveform modification based solely on the most recent passive measurements and the latter relies on a learning or policy-based model to anticipate RFI behavior.

Table I: Summary of CR techniques for DSA. All methods can adapt on a pulse-to-pulse basis.

Technique	T_{CR}	**Performance Metric**	**Response Strategy**	**Ref**
SRA	164 µs	SINR / Bandwidth / Clutter Modulation	Reactive	[30,31,40,41]
SPA	410 µs	CO / MO	Predictive	[32,36]
SLA	410 µs	SINR / Bandwidth	Predictive	[33-35]
SRN	451 µs	CO / SINR / Bandwidth / Clutter Modulation	Reactive	[37,38,42]
SPN	451 µs	SINR / Sidelobes	Predictive	[39]

The CR framework in Fig. 1 also encompasses the use of a high-power tunable matching network [43] to optimize the transmitter front-end power amplifier in terms of power added efficiency (PAE), spectral containment, or output power, the latter of which has been shown to increase the radar's maximum detection range. The goal here is to continually retune this matching network to maintain the optimal RF front-end configuration for the MCR system while the particular DSA parameters are varied. Because the current instantiation is mechanical and the adjustment time is measured in the tens of milliseconds, this technique aims to improve the average performance at the moment, though future versions (e.g. plasma based) are expected to be far faster.

III. Response Time (RT)

A particularly important topic for practical implementation of CR is the collective RT of the algorithms and hardware within the architecture. Despite being a key measure for human cognition [44,45], RT is often over-looked in the context of cognitive RF systems. For example, if two radars have equal levels of cognition (based on some ontology) and arrive at the same result, would the radar that converges to a solution faster have greater cognitive ability? The faster solution could lead to additional interactions with the environment, thereby offering the opportunity to learn at a quicker rate. Measuring the RT for human cognition involves both the time to complete a task and the accuracy of the result [45]. This joint time/accuracy consideration for radar poses a trade-off between a highly accurate (and possibly complex) CR technique and its corresponding latency, which is further exacerbated when multiple CR techniques, multiple objectives, and/or other radar signal processing algorithms can be employed. Finally, as adaptive hardware components are used in future CR designs to expand flexibility for transmitted waveforms, the trade-off between hardware tuning speed and fidelity must also be examined [43].

Our definition of RT reflects the earliest research involving the implementation of artificial intelligence (AI) in radar by addressing the limitations of digital hardware and the algorithm complexity for real-time CR systems [11]. One such example dates back to 1967 for adaptive antenna systems to reject unwanted RFI [46], the solution of which offered a computationally efficient algorithm for real-time implementation. Other examples of RT for CR include knowledge-based systems for radar surveillance where the trade-off is examined between optimal system design features (performance accuracy) and the system capability / latency based on available computational resources (time) [4,13,47]. Other CR research between 2006 and 2014 investigated the PAC based on Fuster's paradigm [1,3] that was later combined with components of knowledge-based systems to form the Fully Adaptive Radar for Track Update-Interval Control approach [48], which itself served as the foundation for the cognitive radar engineering workspace

(CREW) programmable radar platform [49]. CREW demonstrated improved radar performance in tracking a moving target by adapting the pulse repetition frequency and number of pulses per CPI in real-time.

As discussed in Section I, a CR system should be capable of housing a confluence of algorithms and technologies to address a wide variety of radar scenarios. One hardware platform that can support this capability is the universal software radio peripheral (USRP). Although intended for radio communications, the USRP can be programmed to realize radar functionality and operation as well. In [10] the capability of CREW was extended to the USRP, while in [9] the USRP x310 platform was considered for the application of real-time DSA [9], later referred to as a software-defined radar (SDRadar).

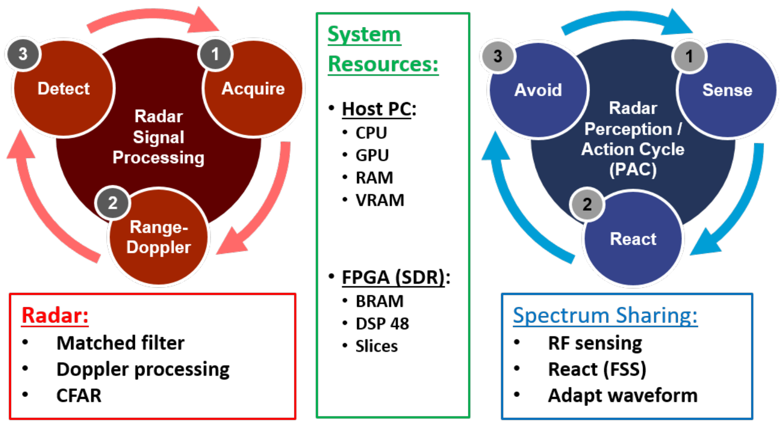

Fig. 3: Resource sharing components of radar signal processing and spectrum sharing.

The original SDRadar architecture of [9] implemented SRA and regulated the timelines of multiple spectrum sensing and radar signal processing algorithms between the field programmable gate array (FPGA) and the graphic processing units (GPUs), while having the appropriate data rate connections to deliver information at the proper time. As summarized in Fig. 3, this basic architecture shares resources between the radar processing and spectrum sensing components, and is foundational for the efficient implementation of multiple CR techniques discussed in Section II, where the goal is to maximize performance in real-time. The SDRadar shares the available resources on the USRP motherboard (i.e. FPGA) and host-PC as illustrated in Fig. 4. This firmware implementation aims to efficiently perform the most fundamental radar functions (i.e. matched filtering, Doppler processing, and constant false alarm rate detection) along with spectrum sensing to achieve real-time operation. As shown in Fig. 4, this functionality mostly uses hardware-accelerated resources like the GPU and FPGA, leaving the central processing unit (CPU) to perform more complex operations that are difficult to implement in hardware (e.g. machine learning techniques). Not only does this hardware-accelerated processing allow room for more advanced algorithm development, it also significantly improves the SDRadar's performance. For example, the original SDRadar architecture described in [9] used the CPU for range-Doppler processing and target detection, resulting in added latency (data processed every other CPI cycle). With the GPU implementation, every CPI is processed in real-time. The original architecture in [9] likewise implemented FSS [31] on the CPU, resulting in a 3.4 ms reaction time to changing RFI. Implementation of FSS on the FPGA has reduced that reaction time to 164 µs [40], a more than 20-fold reduction.

To ascertain the RT for each CR approach it is important to identify the particular timing bottleneck(s). For the SRA approach, the RT-limiting operation is FSS, which involves FFTs (~160 µs) plus a few multiplies and additions. By comparison, the SRA generation of an appropriate chirp waveform via direct digital synthesis requires only a few clock cycles (on the order of nanoseconds). Operations such as range-Doppler processing and CFAR detection are asynchronous from the PAC loop on the FPGA, and thus can be performed on the host-PC without limiting the RT.

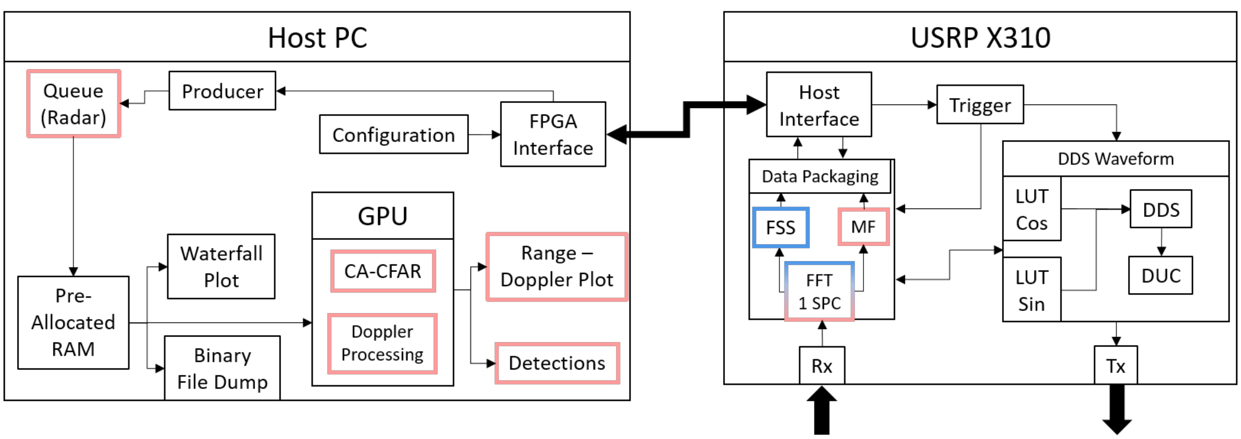

Fig. 4: Block diagram of the basic SDRadar for SRA. Red boxes denote radar functions while blue denotes spectrum functions.

The SPA implementation uses the host-PC and memory to quantify the statistics of RF activity and subsequently predict the likelihood of future RFI. Similarly, SLA uses the host-PC to learn RF activity via policy iteration and select the best future option to avoid RFI. To facilitate construction of each of these models, the host collects and stores a history of RFI metadata states over a given evaluation interval that serves as training data. The SPA training process runs on the timescale of 10s of milliseconds, thus satisfying real-time requirements. Since policy iteration is more computationally complex, the SLA model is necessarily trained on a smaller subset of the stored data. In both of these cases, the predicted/learned RFI state determined on the host then informs chirp waveform generation performed on the FPGA. The history of RFI metadata also continues to be updated in case retraining is needed.

To maximize the utilization of bandwidth the SRN and SPN approaches involve the generation of spectrally-notched waveforms that are forms of random frequency modulation (RFM) [50], which facilitates incoherent sidelobe averaging during slow-time processing to offset the increase in sidelobes incurred by notching, all while preserving the RFM structure to ensure compatibility with high-power transmitters. With regard to RFI determination for use in waveform generation, that component of the RT for these approaches is analogous to SRA and SPA, respectively. The common attribute of notched, random RFM waveform generation within the FPGA is achieved by using the zero-order reconstruction of waveforms (ZOROW) method [51], which involves an analytical spectrum representation of the physical waveform that accounts for the modest rate of the digital-to-analog converter (DAC) and overall lower fidelity available in USRPs. Because ZOROW is an iterative process, there is a relationship with spectral notch depth and RT (e.g. ~25 dB relative to the spectral peak was achieved with an overall RT of 451 μs).

Finally, tuning of the RF front-end to maximize output power is not currently fast enough to adapt dynamically on CR time-scales. Consequently, the algorithm that controls transmit circuit tuning must adhere to a set of known configurations based on expected CR conditions (see [43]) and therefore presently has no impact on the RT. That said, the prospect of increased tuning speeds may change this arrangement in the future.

IV. PAC Regulation

The autonomous regulation of the PAC changes how the CR interacts with the environment. To illustrate this point, consider the echo-location capability of bats [1]. As a bat approaches a target (e.g. an insect), it changes its transmitted sound frequency and duration. This manner of range-dependent adaptation involves an increase in the rate of transmitted waveforms (and received echoes) to pin-point the insect's location during the final moments of engagement. In a similar way, the CR interacts with the environment to determine the rate at which the PAC should change, which could be implemented at various timescales such as the pulse repetition rate, the dwell rate, the scanning rate, or the convergence rate of a high-level decision process. If adaptation is too slow, target information could be lost. If adaptation is too fast, radar resources could be mismanaged.

Fast adaptation can also introduce nonstationary effects. As discussed in [37,40-42], intra-CPI waveform adaptation can translate into a modulation of the clutter, making standard cancellation less effective and ultimately necessitating proper compensation within the receive processing. Likewise, real-time transmit circuit optimization can introduce amplifier impedance modulation if adjustments coincide with transmitted pulses [52]. Similar to the RT discussion, a trade-off between slow and fast PAC adaptation rates exists and should be balanced to best fit evolving target and environmental conditions.

Consider a scenario where a radar and at least one other system are concurrently sharing a common frequency band in a rapidly varying manner. If these systems are to operate without mutual interference, accurate sensing and a short RT are needed. Pulse-agile techniques for spectrum sharing have been shown to perform well at minimizing the mutual interference between non-cooperative, RF systems [37,40-42]. Of course, the efficacy of these techniques is dictated by the speed with which the PAC is completed. For example, Fig. 5 illustrates two different PAC rates for the SDRadar sharing the spectrum with RFI that randomly hops frequency every 2 ms. Specifically, Figs. 5a and 5b show spectrograms in which the SDRadar adaptation rate is 312.5 Hz (or 3.2 ms) and 6.25 kHz (or 164 μs), respectively. When the SDRadar adapts too slowly relative to the RFI (Fig. 5a), the collision rate can get as high as 100%. In contrast, Fig. 5b shows that the higher adaptation rate allows the SDRadar to respond over 10 times faster than the RFI is hopping, realizing a collision rate below 10%. This lower collision rate not only reduces the radar's impact on other signals sharing the band, it also improves the radar's performance.

Fig. 6 shows the resulting receiver operating characteristic (ROC) curves when performing range-Doppler processing based on the spectrum sharing scenarios in Fig. 5. Notice that both slow and fast

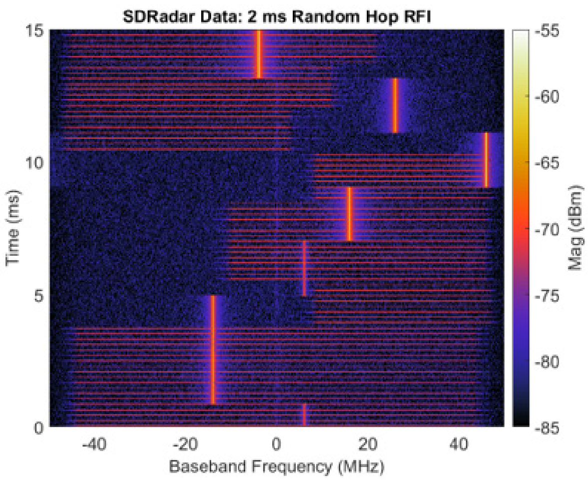

Fig. 5a: 3.2 ms adapt rate

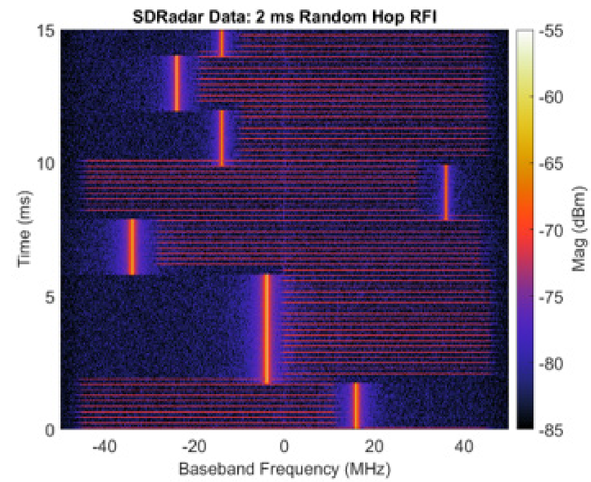

Fig. 5b: 164 µs adapt rate

Fig. 5: Spectrogram of SDRadar sharing a 100 MHz band with a random frequency hopping signal. Collisions are reduced with a faster PAC adaptation rate.

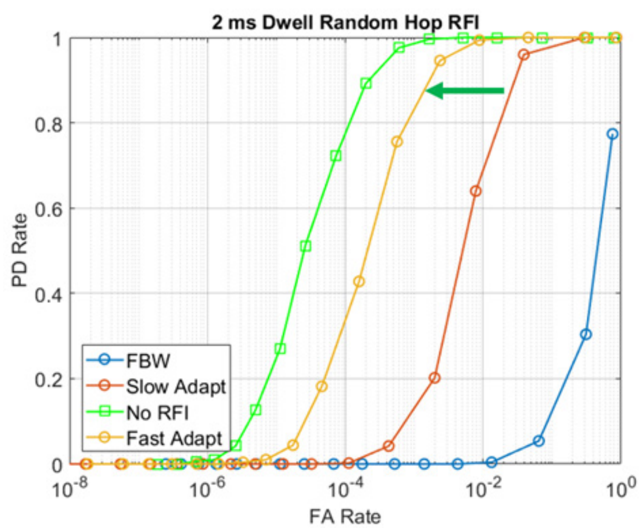

Fig. 6: ROC curves for SDRadar adapting at 3.2 ms ("Slow Adapt") and 160 µs ("Fast Adapt"), while coexisting with randomly hopping RFI.

adaptation show improvement relative to full bandwidth (FBW) operation that involves no adaptation (blue trace), though fast adaptation is clearly much closer to the ideal case of no RFI (green trace). While increasing adaptation speed certainly improves the ability to avoid mutual interference, intra-CPI adaptation can also introduce the modulation effects noted above [37,40-42]. For instance, the delay-Doppler smearing of clutter induced by waveform changes to accommodate dynamic RFI can translate into an increase in the rate of false alarms by a factor of 100 or more [40] if not appropriately compensated.

One such compensation approach is deconvolution using the delay-Doppler point-spread function [41], which borrows from distortion removal in image processing. By treating the clutter-modulated range-Doppler response as the corrupted image and the full bandwidth response as that provided by the ideal ‘lens’, then deconvolution using the point-spread function is effectively an additional filtering ‘lens’ that removes unwanted artifacts from the scene.

A separate approach performs fast-time (range) and slow-time (Doppler) processing jointly since the modulation effect is a result of fast/slow-time coupling (conceptually similar to angle/Doppler coupling for airborne clutter cancellation). The non-identical multiple pulse compression (NIMPC) formulation [42] performs this manner of joint processing, in so doing realizing a multiplicative (instead of additive) increase in the available range and Doppler degrees-of-freedom. The limitation to NIMPC has been high computational cost due to joint operation, though efficient versions are currently have recently been demonstrated [53].

V. Metacognitive Radar (MCR)

The final topic of consideration for practical implementation of CR is the regulation of cognition. The key is knowing *when* to apply the *appropriate* algorithm for a given environmental and target scenario. For example, why apply a cognitive solution when an adaptive, or deterministic, approach is adequate? This question is particularly pertinent for the increasing congested and constrained electromagnetic environment where new cognitive capabilities are so important to address growing spectral/operational complexity [54]. This section therefore discusses the merits and challenges of implementing MCR for algorithm selection.

The MCR framework is a comprehensive feedback and control process to regulate the appropriate CR strategy for changing scenarios, i.e. a best-of-breed approach that recognizes when to use certain CR strategies and when to change between them. The techniques discussed in Section II represent a dichotomy of cognitive solutions ranging from purely adaptive with simple waveform synthesis (e.g. SRA) to training-based models with sophisticated waveform designs (e.g. SPN). Each technique has advantages for enhancing radar performance in specific classes of environmental conditions. Consequently, a comprehensive DSA approach is to select the appropriate algorithm based on real-time observations of RFI and radar performance. This approach treats the CR techniques in Section II as a database, or toolbox, of solutions that can be selected for implementation based on the conditions at hand. That said, this strategy also requires a higher-level decision process to regulate use of the individual CR approaches.

The higher-level cognition process considered here is the bio-inspired metacognitive radar model [21]. Metacognition is a widely studied topic in the fields of human learning, neurological impairments, computation, automation, and control [55]. Metacognition for radar was explored briefly for waveform diversity [56], but more recently for CR applications [21,57]. A block diagram of the MCR model is illustrated in Fig. 7, which involves selecting the appropriate CR technique based on long-term observations of radar performance and the electromagnetic environment (EME). The main components of this model are MCR Knowledge, MCR Monitoring, MCR Control, and the CR techniques [58]. In the context of spectrum sharing, the MCR Knowledge defines the learning rate and capabilities for each CR technique. MCR Monitoring classifies the spectrum to identify a subset of appropriate CR strategies matched to the given EME. MCR Control then defines the regulation of the learning process for implementing CR techniques.

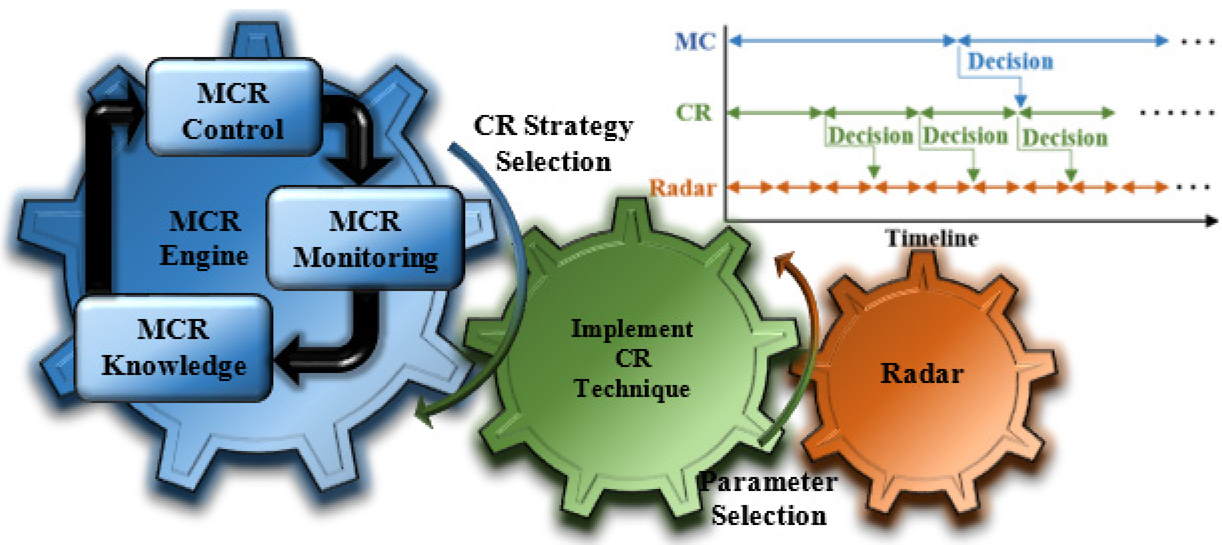

Fig. 7: Metacognitive radar model to select the appropriate CR technique based on real-time observations of radar performance and the EME.

Fig. 7 also depicts the timeline of radar operation, which is much shorter than that of a particular CR technique. In the classical definition of CR, waveforms and other parameters are modified to control the behavior of the radar and improve performance. For the tracking application, the radar would operate over several CPI cycles before the CR can estimate target position and velocity information, and then modify the appropriate parameters. In a similar way, the MCR engine observes CR operations and monitors radar performance over an extended time period while simultaneously monitoring the EME. Based on the radar performance and changing conditions of the EME, the MCR engine explores different CR techniques over time. The CR technique that is expected to produce the "optimal" performance is then selected (or "exploited") for use until either the radar performance degrades or the EME conditions change. To achieve convergence to the optimal CR strategy, i.e. one that yields the best radar performance over time, we develop a sequential decision process analogous to the multi-armed bandit problem. By receiving feedback about each algorithm, the goal is to balance the exploration and exploitation trade-off. During an exploration phase, each technique is selected and radar performance metrics are observed. The exploitation phase then selects the best technique based on the expected performances. For instance, in [59] the "explore-first" approach was used initially, whereby the MCR engine first sequentially samples CR strategies to evaluate radar performance before exploiting the option with the most favorable expected performance for the remaining decision rounds.

Radar performance evaluation during the explore-first phase considers a normalized linear combination of SINR and bandwidth. These metrics allow for simple estimates that correspond to radar target detectability and range resolution, respectively. Past work heavily weighted SINR [21], though the relative weighting of SINR and bandwidth highly depends on the particular radar application. For example, those that require tracking of several closely-spaced targets need higher bandwidth to achieve sufficient range resolution. In contrast, the tracking of distributed moving targets favors SINR, instead of bandwidth, to mitigate range-bin migration during a CPI. Prior work in [60] has demonstrated that short-range moving targets may trade-off SINR for bandwidth while still preserving target detection performance.

The MCR proposed in [21] also explores how to rearrange the CR techniques into simpler categories for implementation. Fig. 8 illustrates the toolbox of solutions that are used for SDRadar. This architecture allows for the combining of different cognitive methods, waveform types, and radar signal processing with the tunable matching network. The "Cognitive Strategy" component uses the fundamental DSA information from the CR techniques discussed in Section II, which are here rearranged into the generic categories of Fixed, React (SRA), Learn (SLA), and Stochastic (SPA). The "Waveform Type" selection determines whether an LFM or RFM waveform is transmitted, and whether the radar will attempt to avoid sub-bands with RFI, implement narrowband notches to mitigate the RFI, or select the entire bandwidth. The maneuverability to avoid and notch is a form of frequency hopping that is dependent on the RFI. Finally, the signal processing is tailored for a generic tracking radar application by using range-Doppler processing. A decision is made in each of these four categories to provide a comprehensive strategy for DSA.

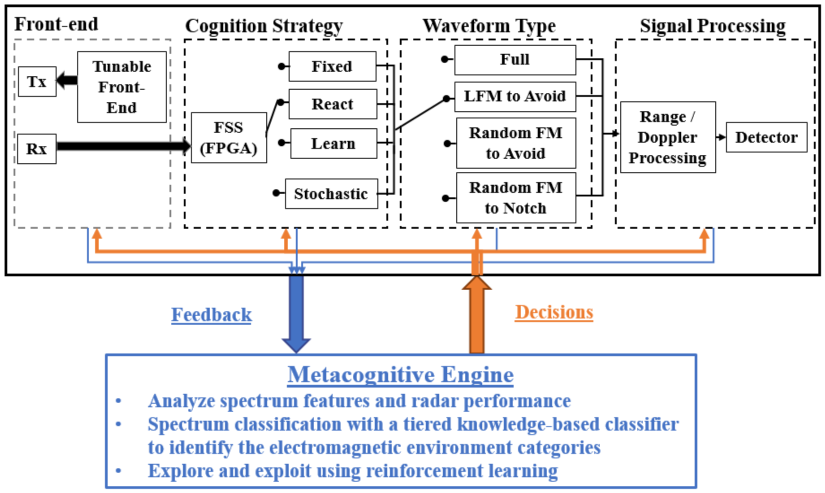

Fig. 8: Implementation architecture of the SDRadar.

This MCR process has been demonstrated to optimize radar performance in changing RFI scenarios via simulation in [21], and has been implemented on the SDRadar hardware for real-time DSA and evaluated on a sequence of changing RFI scenarios. The RFI scenarios include: I) a swept tone interferer that sweeps over a 100 MHz bandwidth, II) a random frequency hopper that moves around the same 100 MHz bandwidth, and III) an intermittent 20 MHz 4G Long Term Evolution (LTE) uplink signal. A spectrum evaluation interval (SEI) is defined to reevaluate the environment for changes in the EME. Each SEI reinitiates the explore-first process every 20 CPIs. Fig. 9a shows the mode of operation driven by the MCR engine over the course of these RFI scenarios while Fig. 9b shows the respective SINR performance. Each RFI scenario lasts for 80 CPIs.

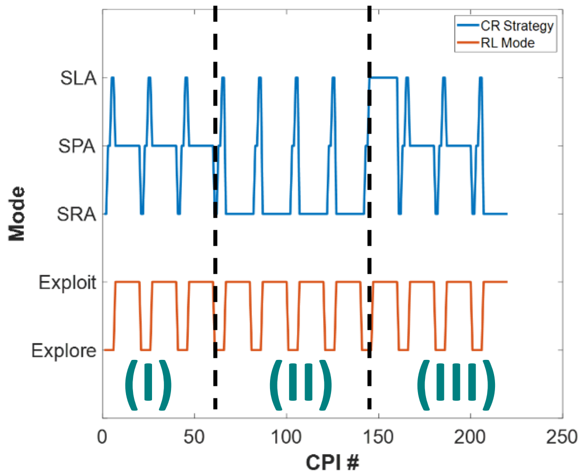

Fig. 9a: MCR mode of operation.

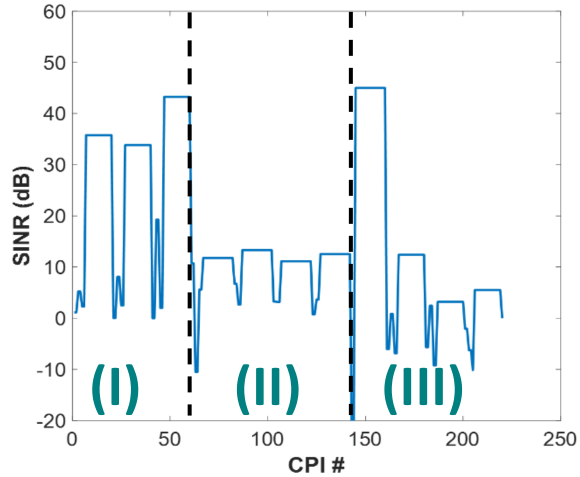

Fig. 9b: Average SINR result.

Fig. 9: MCR example using real-time implantation on the SDRadar. Each figure denotes three RFI scenarios: I) a tone that sweeps over a 100 MHz bandwidth, II) a random frequency hopper that moves around the same 100 MHz bandwidth, and III) an intermittent 20 MHz LTE uplink signal.

RFI Scenario I (swept tone) considers a deterministic interferer that is well suited for SPA. The explore-first process consistently converges to the SPA solution, with reevaluation occurring every 20 CPIs. Fig. 9b shows the exploit phase consistently achieving higher performance during the MCR process. The SINR variation during the exploit cycles are the result of occasional collisions with the high-power swept RFI. While a few collisions occur in the first two exploit cycles, SPA perfectly avoids the RFI in the third cycle, which results in a SINR increase. RFI Scenario II (random hopper) shows that the MCR converges to SRA each SEI. The predictive SPA and SLA CR techniques incur a degradation in performance for random RFI scenarios.

Finally, RFI Scenario III poses a challenging LTE RFI scenario. Due to the amount of variation present in these LTE patterns, the optimal solution appears to vary over time. From CPI 140 to 160, the MCR is trained on the spectral data from the prior SEI (CPI 120 to 140), which causes SLA to converge to a solution that avoids transmitting in bands with RFI altogether. By avoiding transmitting in active LTE bands, SLA results in a high SINR from CPI 140 to 160. Once the MCR has learned the new spectral patterns, the process favors higher bandwidth in exchange for lower SINR and begins to alternate between SPA and SRA. During the two next SEIs, the MCR converges to SPA. The last SPA cycle shows an SINR degradation compared to the previous one, which results in the MCR selecting SRA to slightly recover the losses.

VI. Conclusions

There are numerous practical aspects involved with implementing CR techniques onto an SDRadar platform, with the various techniques possessing different advantages and disadvantages according to the statistical EME conditions. Response time of a CR capability requires careful consideration, as does regulation of the perception-action cycle relative to the dynamic nature of the EME. Intra-CPI waveform adaptation introduces another performance trade-space between adequately avoiding/notching in-band RFI and the modulation of clutter that can hinder effective cancellation, which subsequently requires compensation. At a higher level, management of the overall CR process necessitates the automation of CR technique selection in a metacognitive manner that balances all the various trade-offs to select a comprehensive approach for a given situation.

References

[1] Haykin, "Cognitive radar: a way of the future," *IEEE Signal Processing Mag.*, vol. 23, no. 1, pp. 30-40, Jan. 2006.

[2] A. F. Martone, "Cognitive radar demystified," in *URSI Radio Science Bulletin*, vol. 2014, no. 350, pp. 10-22, Sept. 2014.

[3] S.Z. Gurbuz, *et al.*, "An overview of cognitive radar: past, present, and future," *IEEE Aerospace & Electronic Systems Mag.*, vol. 34, no. 12, pp. 6-18, Dec. 2019.

[4] J.R. Guerci, E.J. Baranoski, "Knowledge-aided adaptive radar at DARPA: an overview," *IEEE Signal Processing Mag.*, vol. 23, no. 1, pp. 41-50, Jan. 2006.

[5] M.S. Greco, F. Gini, P. Stinco K. Bell, "Cognitive radars: on the road to reality: progress thus far and possibilities for the future," *IEEE Signal Processing Mag.*, vol. 35, no. 4, pp. 112-125, July 2018.

[6] *IEEE Standard for Radar Definitions*, IEEE Standard 686, 2017.

[7] A. Huizing, A. Charlish, S. Brüggenwirth, "Response to the letter by Fred Daum," *IEEE Aerospace and Electronic Systems Mag.*, vol. 35, no. 6, pp. 79-79, 1 June 2020.

[8] J.P. Stralka, G.G.L. Meyer, "OFDM-based wideband phased array radar architecture," *IEEE Radar Conf.*, Rome, May 2008, pp. 1–6.

[9] B.H. Kirk, et. al., "Cognitive software defined radar: waveform design for clutter and interference suppression," *SPIE Radar Sensor Technology*, Anaheim, CA, May 2017.

[10] J.M. Christiansen, G.E. Smith, K.E. Olsen, "USRP based cognitive radar testbed," *IEEE Radar Conf.*, Seattle, WA, May 2017, pp. 1115-1118.

[11] R.M. Narayanan, "Aspects of radar waveform design," *Proc. Intl. Radar Symp. India*, Bangalore, India pp. 8–14, Dec. 2009.

[12] W.D. Blair, G.A. Watson, S.A. Hoffman, G.L. Gentry, "Benchmark problem for beam pointing control of phased array radar against maneuvering targets," *American Control Conf.*, vol. 4, June 1994, pp. 2071–2075.

[13] V.C. Vannicola, J.A. Mineo, "Applications of knowledge based systems to surveillance," *IEEE National Radar Conf.*, Ann Arbor, MI, April 1988, pp. 157–164.

[14] A. Charlish, K.L. Bell, C. Kreucher, “Implementing perception-action cycles using stochastic optimization,” *IEEE Radar Conf.*, Florence, Italy, Sept. 2020.

[15] F. Barbaresco, J.C. Deltour, G. Desodt, B. Durand, T. Guenais, C. Labreuche, “Intelligent m3r radar time resources management: Advanced cognition, agility autonomy capabilities,” *Intl. Radar Conf.*, Oct. 2009, pp. 1–6.

[16] G. E. Smith, Z. Cammenga, A. Mitchell, K. L. Bell, J. Johnson,M. Rangaswamy, and C. Baker, “Experiments with cognitive radar,” *IEEE Aerospace and Electronic Systems Mag.*, vol. 31, no. 12, pp. 34–46, Jan. 2016.

[17] R. Oechslin, S. Wieland, S. Hinrichsen, U. Aulenbacher, P. Wellig, “A cognitive radar testbed for outdoor experiments,” *IEEE Aerospace and Electronic Systems Mag.*, vol. 34, no. 12, pp. 40–48, Dec. 2019.

[18] G. Montavon, W. Samek, K.R. Muller, “Methods for interpreting and understanding deep neural networks,” *Digital Signal Processing*, vol. 73, pp. 1–15, 2018

[19] M. Greenspan, “Potential pitfalls of cognitive radars,” *IEEE Radar Conf.*. Cincinnati, OH, 2014, pp. 1288–1290.

[20] A.F. Martone, *et al.*, “Practical aspects of cognitive radar,” *IEEE Radar Conference*, Florence, Italy, Sept. 2020.

[21] A.F. Martone, *et al.*, “Metacognition for Radar Coexistence,” *IEEE Intl. Radar Conf.*, Wash., D.C., Apr. 2020.

[22] S. Bhattarai, *et. al.*, “An overview of dynamic spectrum sharing: ongoing initiatives, challenges, and a roadmap for future research,” *IEEE Trans. on Cognitive Communications and Networking*, 2(2), pp.110-128, June 2016.

[23] M. Palola et al., “Field trial of the 3.5 GHz citizens broadband radio service governed by a spectrum access system (SAS),” *IEEE Intl. Symp. on Dynamic Spectrum Access Networks*, Piscataway, NJ, 2017, pp. 1-9.

[24] J. H. Reed, et. al., “On the co-existence of TD-LTE and radar over 3.5 GHz band: an experimental study,” *IEEE Wireless Communications Letters*, vol. 5, no. 4, pp. 368-371, Aug. 2016.

[25] F. Liu, C. Masouros, A.P. Petropulu, H. Griffiths, L. Hanzo, “Joint radar and communication design: applications, state-of-the-art, and the road ahead,” *IEEE Transactions on Communications*, vol. 68, no. 6, pp. 3834–3862, Feb. 2020.

[26] M. Shafi, *et al*., “5G: a tutorial overview of standards, trials, challenges, deployment, and practice,” *IEEE Journal on Selected Areas in Communications*, vol. 35, no. 6, pp. 1201-1221, June 2017.

[27] Federal Communications Commission, (Aug. 2015). *The FCC’s 5G FAST Plan* [Online]. Available: https://www.fcc.gov/5G.

[28] Federal Communications Commission, “FCC starts first 5G mid-band spectrum auction today,” Wash. D.C., Commission Documents, July, 2020.

[29] Federal Communication Commission, “Auction of Advanced Wireless Services (AWS-3) Licenses Closes,” Wash. D.C., DA 15-131, 2015.

[30] A.F. Martone, K.I. Ranney, K. Sherbondy, K.A. Gallagher, S.D. Blunt, “Spectrum allocation for noncooperative radar coexistence,” *IEEE Trans. Aerospace & Electronic Systems*, vol. 54, no. 1, pp. 90-105, Feb. 2018.

[31] B.H. Kirk, R.M. Narayanan, K.A. Gallagher, A.F. Martone, K.D. Sherbondy, “Avoidance of time-varying radio frequency interference with software-defined cognitive radar,” *IEEE Trans. Aerospace & Electronic Systems*, vol. 55, no. 3, pp. 1090-1107, June 2019.

[32] J.A. Kovarskiy, B.H. Kirk, A.F. Martone, R.M. Narayanan and K.D. Sherbondy, “Evaluation of real-time predictive spectrum sharing for cognitive radar,” in *IEEE Trans. Aerospace & Electronic Systems*, vol. 57, no. 1, pp. 690-705, Feb. 2021.

[33] E. Selvi, R.M. Buehrer, A.F. Martone, K.D. Sherbondy, “On the use of Markov decision processes in cognitive radar: an application to target tracking,” *IEEE Trans. Aerospace & Electronic Systems*, vol. 56, no. 5, pp. 3904-3921, Oct. 2020.

[34] C. E. Thornton, M. A. Kozy, R. M. Buehrer, A. F. Martone and K. D. Sherbondy, “Deep Reinforcement Learning Control for Radar Detection and Tracking in Congested Spectral Environments,” to appear in *IEEE Transactions on Cognitive Communications and Networking*.

[35] C. E. Thornton, R.M. Buehrer, and A.F. Martone, “Constrained online learning to mitigate distortion effects in pulse-agile cognitive radar,” *IEEE Radar Conf.*, Atlanta, GA, May 2021 (Under Review).

[36] J. A. Kovarskiy, *et al.*, “Comparing stochastic and Markov decision process approaches for predicting radio frequency interference,” *Proc. SPIE, Radar Sensor Technology XXIII*, 1100318, Baltimore, MD, May 2019.

[37] B. Ravenscroft, J.W. Owen, J. Jakabosky, S.D. Blunt, A.F. Martone, K.D. Sherbondy, “Experimental demonstration and analysis of cognitive spectrum sensing and notching for radar,” *IET Radar, Sonar & Navigation*, vol. 12, no. 12, pp. 1466-1475, Dec. 2018.

[38] J.W. Owen, C.A. Mohr, B.H. Kirk, S.D. Blunt, A.F. Martone, K.D. Sherbondy, “Demonstration of real-time cognitive radar using spectrally-notched random FM waveforms,” *IEEE Intl. Radar Conf.*, Washington, DC, Apr. 2020, pp. 123-128.

[39] J.A. Kovarskiy, et. al., “Spectral prediction and notching of RF emitters for cognitive radar coexistence,” *IEEE Intl. Radar Conf.*, Washington, DC, Apr. 2020.

[40] B.H. Kirk, A.F. Martone, K.D. Sherbondy, R.M. Narayanan, “Performance analysis of pulse-agile SDRadar with hardware accelerated processing,” *IEEE Intl. Radar Conf.*, Washington, DC, Apr. 2020.

[41] B.H. Kirk, A.F. Martone, K.D. Sherbondy, R.M. Narayanan, “Mitigation of target distortion in pulse-agile sensors via Richardson–Lucy deconvolution,” *Electronics Letters*, vol. 55, no. 23, pp. 1249-1252, Nov. 2019.

[42] B. Ravenscroft, et al., “Experimental assessment of joint range-Doppler processing to address clutter modulation from dynamic radar spectrum sharing,” *IEEE Intl. Radar Conf.*, Washington, DC, Apr. 2020.

[43] A. Dockendorf, et al., "Fast optimization algorithm for evanescent-mode cavity tuner optimization and timing reduction in software-defined radar implementation," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 56, no. 4, pp. 2762-2778, Aug. 2020.

[44] P.C. Kyllonen, J. Zu, "Use of response time for measuring cognitive ability," *Journal of Intelligence*, vol. 4, no. 14, Nov. 2016.

[45] G. Maris, H. Maas, "Speed-accuracy response models: Scoring rules based on response time and accuracy," *Psychometrika*, vol. 77, pp. 615–633, Sept. 2012.

[46] B. Widrow, P. E. Mantey, L. J. Griffiths and B. B. Goode, "Adaptive antenna systems," *Proceedings of the IEEE*, vol. 55, no. 12, pp. 2143-2159, Dec. 1967.

[47] W. Melvin, *et. al,* "Knowledge-based space-time adaptive processing for airborne early warning radar," *IEEE Aerospace and Electronic Systems Magazine*, vol. 13, no. 4, pp. 37-42, April 1998.

[48] K.L. Bell, *et. al.*, "Cognitive radar framework for target detection and tracking," *IEEE Journal of Selected Topics in Signal Processing*, vol.9, no.8, pp. 1427-1439, Dec. 2015.

[49] J.M. Christiansen, G.E. Smith, "Development and calibration of a low-cost radar testbed based on the universal software radio peripheral," *IEEE Aerospace and Electronic Systems Mag.*, vol. 34, no. 12, pp. 50-60, 1 Dec. 2019

[50] S.D. Blunt, *et al*, "Principles & applications of random FM radar waveform design," *IEEE AES Systems Mag*., vol. 35, no. 10, pp. 20-28, Oct 2020.

[51] C.A. Mohr, J.W. Owen, S.D. Blunt, C.T. Allen, "Zero-order reconstruction optimization of waveforms (ZOROW) for modest DAC rates," *IEEE Intl. Radar Conf*., Washington, DC, Apr. 2020.

[52] A. Egbert et al., "The effect of real-time radar transmitter amplifier impedance tuning on range and Doppler detection accuracy," *IEEE Texas Symp. on Wireless and Microwave Circuits and Systems*, Waco, TX, March 2019, pp. 1-4.

[53] C. Jones, B. Ravenscroft, J. Vogel, S.M. Shontz, T. Higgins, K. Wagner, S. Blunt, "Computationally efficient joint-domain clutter cancellation for waveform-agile radar," *IEEE Radar Conf*, Atlanta, GA, May 2021.

[54] Department of Defense, "Electromagnetic Spectrum Superiority Strategy," Oct. 2020.

[55] Z. Yang, K. Merrick, L. Jin, H.A. Abbass, "Hierarchical deep reinforcement learning for continuous action control," IEEE Trans. Neural Networks & Learning Systems, vol. 29, no. 11, pp. 5174-5184, Nov. 2018.

[56] G.T. Capraro, M.C. Wicks, "Metacognition for waveform diverse radar," *Intl. Waveform Diversity & Design Conf*., Kauai, HI, Jan. 2012.

[57] K.V. Mishra, *et al.*, "Toward metacognitive radars: concept and applications," *IEEE Intl. Radar Conf.*, Wash., D.C., Apr. 2020.

[58] H. Asadi, A.H. Volos, M.M. Marefat, T. Bose, "Metacognition and the next generation of cognitive radio engines," *IEEE Communications Mag.*, vol. 54, no. 1, pp. 76-82, Jan. 2016.

[59] A. Slivkins, "Introduction to multi-armed bandits," *Foundations and Trends in Machine Learning*, vol. 12 no. 1-2, Nov. 2019.

[60] A. Martone, K. Gallagher, K. Sherbondy, A. Hedden and C. Dietlein, "Adaptable waveform design for enhanced detection of moving targets," *IET Radar, Sonar & Navigation*, vol. 11, no. 10, pp. 1567-1573, 10 2017.