Practical Aspects of Cognitive Radar

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

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

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

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

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

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

***Abstract*—In this paper we examine some of the practical aspects of implementing cognitive radar (CR) techniques onto software defined radar (SDRadar) platforms. These aspects include: 1) the response time (RT) of algorithms and components to determine latency bottlenecks, 2) autonomous regulation of the perception-action cycle (PAC) to determine "how fast the CR *can* interact with the environment" as well as "how fast the CR *should* interact with the environment," and 3) regulation of the cognition level to understand how to select a particular CR technique appropriately for a given dynamically-changing environment. To provide concrete examples of these three implementation aspects for CR, we will focus on the specific application of target tracking in a congested spectral environment.**

Keywords—cognitive radar, spectrum sharing, metacognition, spectrum sensing, waveform diversity, machine learning

I. Introduction

The CR research area has produced a wide variety of techniques that have the potential to improve radar performance [1]. Applying artificial intelligence to radar has predominantly focused on target tracking, though more recent applications have emerged in the context of spectrum sharing and passive radar [2-5]. Each CR research area employs an ontology for modeling and algorithm development. Several ontologies have been proposed to map the fundamental building blocks of cognition to realizable radar solutions [1, 6, 7], though a standard has not yet been established within the community.

Despite the insightful ontologies and the wide variety of available CR research, the real-time implementation of CR techniques onto a radar sensor is still in its infancy. The big-picture view of how a *confluence of algorithms and technologies* are used to address real-life radar scenarios is often missed. Practical development of CR technology requires the integration of multiple, disparate CR techniques onto a flexible software-defined radar architecture. These CR techniques as a whole represent a tool-box of solutions that must be chosen autonomously, and in real-time, by a high-level decision process to address multiple fast-changing target and environmental scenarios. As researchers have evolved their models into this larger perspective, a new set of essential research topics has begun to emerge, which include: 1) the *response time (RT) of algorithms and components* to determine latency bottlenecks within the CR architecture, 2) *autonomous regulation of the perception-action cycle (PAC)* to determine "how fast the CR *can* interact with the environment" as well as "how fast the CR *should* interact with the environment," and 3) *regulation of the cognition level* to understand how to select the appropriate CR technique for dynamically-changing scenarios.

In this paper we explicitly describe these three essential research topics for practical implementation of cognitive radar. To provide concrete examples we will focus on the application of target tracking in congested spectral environments. Section II describes the application of cognitive radar to spectrum sharing. Section III then discusses the RT for CR, which is often overlooked as a key component in cognitive RF systems even though it is a crucial measure of human cognition [8]. Based on definitions used for human cognition, the RT for CR is assessed with respect to the *time* a technique converges to a solution and the subsequent *accuracy* of the technique's result in improving radar performance, thereby posing a trade-off between a highly accurate (and possibly complex) CR technique and its associated latency.

Section IV discusses the *autonomous regulation of the PAC* and the importance of "how fast" the CR should interact with the environment. Slow adaptation could result in inaccurate target information while fast adaptation could result in a mismanagement of radar resources. Adapting quickly, i.e. within the coherent processing interval (CPI) of the radar, can also have consequences for clutter cancellation [9, 10]. Section IV likewise discusses factors contributing to the *regulation of cognition*. Much of the published work in the literature investigates CR techniques to increase the level of cognition; however, why apply a cognitive solution when an adaptive approach, or deterministic approach, will suffice? The key consideration for CR is knowing when to apply the appropriate algorithm for a given environmental and target scenario.

II. Cognitive Radar for Spectrum Sharing

CR for spectrum sharing has received increased attention due to the reallocation of spectrum from government RF systems to the commercial communications industry. This challenge is of particular concern for radar since most legacy systems do not have the frequency agility to share the spectrum [5, 11]. In this paper we examine the application of non-cooperative radar coexistence for real-time dynamic spectrum access (DSA). The DSA approach requires that the electromagnetic environment (EME) be passively measured over time to determine an appropriate center frequency and bandwidth for the radar to occupy. The measured EME power spectra are used to detect the presence of other RF users within the overall operating band B, thus enabling the radar to decide how best to access the band. This frequency allocation choice should mitigate mutual interference by maximizing the radar's signal to interference plus noise ratio (SINR). Ideally the radar would occupy the entire bandwidth B to enhance range resolution; however, this arrangement incurs mutual interference with other RF users, and thus degrades SINR. A compromise for coexistence is to have the radar reduce its bandwidth and identify a sub-band (within the overall operating band) that incurs minimal interference [4]. A frequency sub-band (SB), with bandwidth $\beta \leq B$ and center frequency f_{SB}, is defined as a contiguous set of frequency bins within a measured power spectrum representing a feasible frequency allocation for the radar. Multiple CR techniques have been explored by the authors to promote effective DSA for radar. These techniques implement the PAC with different learning models and waveform types to determine the best performance for radar under this coexistence scenario. Each technique also has advantages and disadvantages based on the nature of the RF interference (RFI).

The first technique is Sense-React-Avoid, which quickly responds to RFI in a computationally efficient manner. As shown in Fig. 1 Sense-React-Avoid uses the power spectrum to estimate a snapshot of the EME that is then processed by the Fast Spectrum Sensing (FSS) algorithm [12]. FSS refines this spectral information into representative metadata by merging "closely-spaced" frequency bins into clusters of low and high power RFI, with the low-power clusters representing available sub-bands for radar operation. Sense-React-Avoid then selects the center frequency and bandwidth of the low-power cluster possessing the largest bandwidth for radar operation. This spectral allocation is then used by the radar for waveform synthesis.

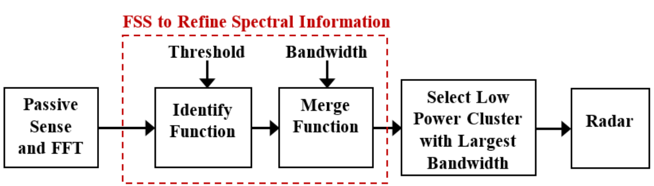

Fig. 1: Block diagram of Sense-React-Avoid approach.

The Sense-Predict-Avoid approach in Fig. 2 seeks to adapt radar waveforms in anticipation of changes in RFI rather than reacting to changes after they occur [13]. Consequently, while the Sense-React-Avoid strategy performs sub-optimally during RFI state transitions, the Sense-Predict-Avoid strategy can mitigates error during these transitions. Sense-Predict-Avoid employs a stochastic model that quantifies RFI activity over time as an alternating renewal process. In this approach, the RF spectrum is channelized into N equally spaced sub-bands and a history of RFI *ON* and *OFF* states are observed in each sub-band. The time spent by each sub-band in each state is used to estimate distribution functions for each respective sub-band and state. Spectrum sensing paired with cumulative distribution functions therefore determines the likelihood of future spectrum availability. Low computational resources for carrying out these tasks allows the system to learn RFI statistics online in real-time [13]. After the model is trained online, Sense-Predict-Avoid observes the most recent spectral activity and selects a contiguous set of sub-bands with a low likelihood of RFI.

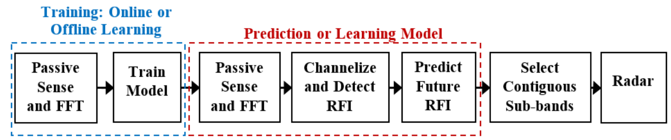

Fig. 2: General block diagram of the Sense-Predict-Avoid and Sense-Learn-Avoid approaches.

The general block diagram in Fig. 2 is also used for Sense-Learn-Avoid, a Reinforcement Learning approach that treats the CR's environment as a Markov Decision Process (MDP) [14]. An MDP is a mathematical model for Reinforcement Learning characterized by the set of states, actions, state transitions, and reward functions in which pairs of states and actions are mapped to numeric rewards. The goal of Reinforcement Learning is to find a policy that the radar can use to select appropriate actions given observed states, where policies can be indirectly determined using function approximations via deep neural networks [15]. The potentially high dimensionality of the policy iteration problem was later reduced using a Deep Neural Network within a Q-learning framework (termed a Deep Q-Network). The Deep Q-Network allows for larger dimensionality of the state space during an offline training period. After the training period has ended, learned behavior then guides the CR waveform selection with minimal complexity. Similar to Sense-Predict-Avoid, the Sense-Learn-Avoid approach observes the most recent spectral activity and selects a contiguous set of sub-bands with a low likelihood of RFI. Unlike Sense-Predict-Avoid, however, Sense-Learn-Avoid can also learn other behaviors, e.g. when it is actually acceptable to allow RFI in the radar band. Additionally, when applying a Deep Q-Network, Sense-Learn-Avoid can also apply online learning and potentially adapt to a changing environment.

Recent work has compared Sense-Predict-Avoid and Sense-Learn-Avoid [16]. While Sense-Learn-Avoid performs worse on random, intermittent RFI scenarios, this mode shows improvement in real measured RFI environments. Sense-Predict-Avoid demonstrates good performance in the intermittent scenarios, with similar performance to Sense-Learn-Avoid for highly variable, real RFI. The MDP-based strategy of Sense-Learn-Avoid tends toward higher bandwidth utilization and fewer waveform adaptations compared to prediction via Sense-Predict-Avoid, which elicits fewer collisions and more frequent waveform adaptation. At the cost of higher computational resources and longer training time, Sense-Learn-Avoid could potentially be able to recognize patterns that are not captured by the statistics in the Sense-Predict-Avoid approach.

The Sense-React-Notch approach in Fig. 3 generates waveforms possessing spectral notches rather than adapting the center frequency and bandwidth of a single contiguous band. This technique relies on spectral shaping optimization to realize instantiations of random FM waveforms [17] that are continuous, have a constant time-domain envelope, and possess an enforced spectral mask containing notches based on the FSS response [9], where notch depth depends on the particular shaping optimization and corresponding number of iterations.

The main advantage of this approach is that the RF spectrum can be utilized at the full bandwidth, sans notched regions.

Fig. 3: Block diagram of the Sense-React-Notch approach.

It has also very recently been shown that Sense-React-Notch can likewise be combined with prediction [18] to facilitate the CR technique of Sense-Predict-Notch in Fig. 4. Employing notching with prediction can improve bandwidth utilization compared to avoidance while maintaining prediction accuracy. Generating notched waveforms requires additional computation time that may limit the adaptation rate. Further, highly random unpredictable RFI may result in the erroneous placement of notches and reduce bandwidth utilization. The ability to use the Sense-Predict-Notch approach when appropriate becomes essential for rapidly changing RFI environments.

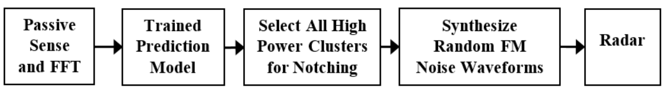

Fig. 4: Block diagram of the Sense-Predict-Notch approach.

The above CR strategies modify the radar waveform from CPI to CPI and possibly from pulse to pulse. The proposed approach of Fig. 5 considers a high-power tunable matching network [19, 20] to maximize the Power Added Efficiency, or output power, of the transmitter front-end power amplifier. The goal of this technique is to retune a matching network to maintain the best amplifier efficiency and spectral containment of the transmitted waveform as transmit parameters are varied by other DSA implementations. It has been shown that the output power of the radar transmission is significantly improved with this procedure [21], leading to increased maximum detection range by the radar. A current disadvantage is that the digital DSA approaches described above can adapt much faster than the physical high-power tuner, which could cause sub-optimal matching. The present high-power tuner technology relies on mechanical tuning techniques, so it requires at least tens of milliseconds to perform a tuning operation.

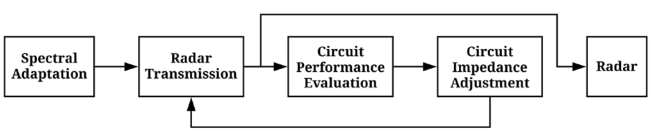

Fig. 5: Block diagram of the impedance tuning process.

III. The Response Time of Cognitive Radar

Some of the earliest research on implementing artificial intelligence (AI) for radar required an examination of the RT, where it was shown that the limiting factors of integrating AI with radar were based on algorithm complexity and the available digital hardware [2, 22-24]. These implementation methods considered a measured trade-off between optimal system design to maximize performance (*accuracy*) and system capability / latency based on available computational resources to obtain a solution (*time*). More recent models have examined implementation of CR techniques on software-defined platforms [25-28]. A challenge with several of these designs is the size, weight, power, and cost (SWaP-C) of the radar platform. A possible hardware platform that can support this low SWaP-C capability is the universal software radio peripheral (USRP). Although intended for radio communications, the USRP can be programmed to fully support radar functionality and operation [29]. Two research groups recently applied cognition to the USRP: one extending the capability of the cognitive radar engineering workspace [30], and the other using the USRP x310 platform for the application of real-time DSA [25]. The latter platform has subsequently been referred to as software-defined radar (SDRadar).

The basic SDRadar architecture [31] (see Fig. 6) implements Sense-React-Avoid and regulates the timelines of multiple spectrum sensing and radar signal processing algorithms between the Field Programmable Gate Array (FPGA) and Graphic Processing Units (GPUs), while having the appropriate data rate connections to deliver information at the proper time. 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 V, where the goal is to maximize performance (*accuracy*) in real-time.

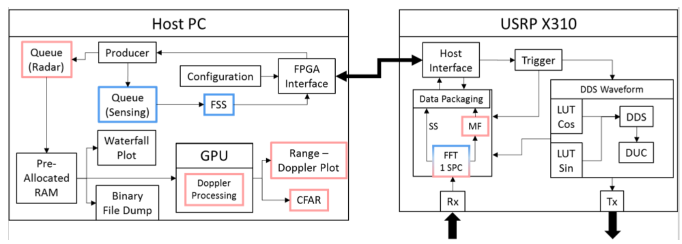

Fig. 6: Block diagram of the basic SDRadar for Sense-React-Avoid.

The SDRadar sensor shares the available resources on the USRP motherboard (i.e. FPGA) and host-PC as described in Fig. 6. The essential radar functions are highlighted in red and the spectrum sensing functions highlighted in blue. This implementation aims to efficiently implement the most fundamental functionality of radar processing (i.e. matched filtering, Doppler processing, and constant false alarm rate detection) along with spectrum sensing (i.e. fast Fourier transform (FFT), signal detection, and clustering) for real-time operation. As shown in Fig. 6, this functionality mostly utilizes hardware-accelerated resources like the GPU and FPGA, leaving the CPU to perform more complex operations that are difficult to implement in hardware (i.e. 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 [25] 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 [25] likewise implemented FSS [12] on the CPU, resulting in a reaction time of 3.4 ms to changing RFI. Implementation of FSS on the FPGA has reduced that reaction time to 160 μs, which includes the time needed to sense the spectrum, process the FFT, decide on the center frequency and bandwidth for DSA, and then transmit and receive 1 radar pulse.

IV. Regulation of the PAC

The *autonomous regulation of the PAC* changes the rate that the CR interacts with the environment. To illustrate this point, consider the echo-location capability of bats [23]. As the bat approaches a target (e.g. an insect), it will change its transmitted sound frequency and duration, which is referred to as range-dependent adaptation. Reducing pulse duration as range decreases subsequently increases the rate of the transmitted (and received) waveforms in order to pin-point the insect location during the final moments of engagement.

In a similar way, the CR must interact with the environment to determine the rate at which the PAC should change. If adaptation is too slow, target information could be lost; if adaptation is too fast, radar resources could be mismanaged. Adapting quickly, i.e. within the CPI of the radar, can also have adverse consequences for range-Doppler processing. As illustrated experimentally in [9, 10, 32], intra-CPI adaptation of waveforms introduces both sidelobe and mainlobe variation to individual pulse compression responses, which collectively translate into nonstationary modulation of the clutter. Consequently, standard clutter cancellation is much less effective and requires some means of compensation. Similar to the RT discussion, a trade-off between slow and fast PAC adaptation rates therefore exists and should be balanced to best fit evolving target and environmental conditions.

To effectively share the spectrum for DSA, accurate spectrum sensing and fast adaptation are needed, which does require pulse-agile techniques that have been shown to perform quite well at minimizing the mutual interference between non-cooperative RF systems [9, 10]. Of course, the efficacy of these techniques is limited by the speed with which the PAC is completed. For example, Fig. 7 illustrates two different PAC speeds for the SDRadar sharing the spectrum with RFI that randomly hops frequency every 2 ms. Specifically, Figs. 7a and 7b show spectrograms in which the SDRadar adaptation rate is 3.2 ms and 160 µs, respectively.

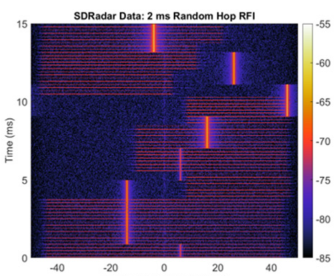

Fig. 7a: 3.2 ms adapt rate

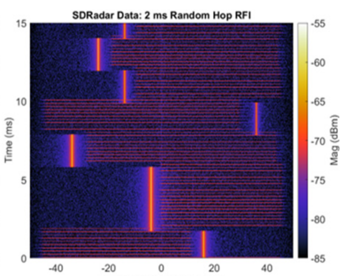

Fig. 7b: 160 µs adapt rate

Fig. 7: Spectrogram of SDRadar sharing a 100 MHz band with a random frequency hopping signal. Collisions are reduced with a fast adaptation rate of the PAC.

When the SDRadar adapts too slowly relative to the RFI (Fig. 7a) the collision rate can get as high as 100%. In the scenario shown in Fig. 7b, however, because the accelerated adaptation rate allows the SDRadar to respond over 10 times faster than the RFI is hopping, the collision rate is now < 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. 8 shows the resulting receiver operating characteristic (ROC) curves when performing range-Doppler processing based on the SDRadar sharing the spectrum according to the adaptation rates from Fig. 7. Notice that both slow and fast adaptation show improvement over just transmitting a full-bandwidth waveform (blue trace), but the fast adaptation clearly lies much closer to the ideal case of no RFI (green trace).

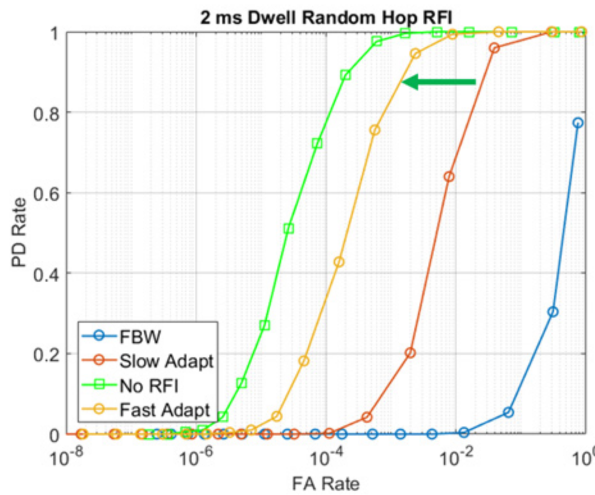

Fig. 8: ROC curves where the SDRadar is adapting at 3.2 ms ("Slow Adapt") and 160 µs ("Fast Adapt"), while coexisting with randomly hopping RFI.

While increasing the adaptation speed of the system improves the ability to avoid mutual interference, intra-CPI waveform adaptation does introduce the above-noted modulation effect that degrades clutter cancellation [9, 10, 32]. Specifically, a smearing of the delay-Doppler point spread function is induced that results in an increase in the residual clutter after cancellation, which subsequently translates into an increase in the rate of false alarms by a factor of 100 or more [31] if not appropriately compensated. Various approaches have thus far been developed and demonstrated to address aspects of this effect, including deconvolution [10], combined mismatched filtering and clutter borrowing/filling [32], and joint range-Doppler processing [33], with varying efficacy and computational cost.

V. Regulation of Radar Cognition

The techniques discussed in Section II represent a dichotomy of cognitive solutions ranging from purely adaptive with simple waveform synthesis (Sense-React-Avoid) to learning models with sophisticated waveform designs (Sense-Predict-Notch). It has been shown that each technique has advantages with regard to enhancing radar performance in particular EME conditions. A comprehensive approach for DSA therefore involves the selection of an appropriate algorithm based on the real-time observations of radar performance and the RFI. This perspective treats the individual CR techniques in Section II as components in a database, or tool box, of solutions that are selected for implementation based on the conditions at hand. Implementation of this strategy 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 [34]. Metacognition is a widely studied topic in the fields of human learning, neurological impairments, computation, automation, and control [35]. Interestingly enough, metacognition for radio has only been explored in a limited capacity and is a relatively new topic for radar [36, 37]. An illustration of the metacognitive radar (MCR) model is depicted in Fig. 9. This model selects the appropriate CR technique based on real-time observations of radar performance and the EME. The main components of this model are MCR Knowledge, MCR Monitoring, MCR Control, and the CR techniques [36]. In the context of spectrum sharing [34], 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 defines the regulation of the learning process for implementing CR techniques.

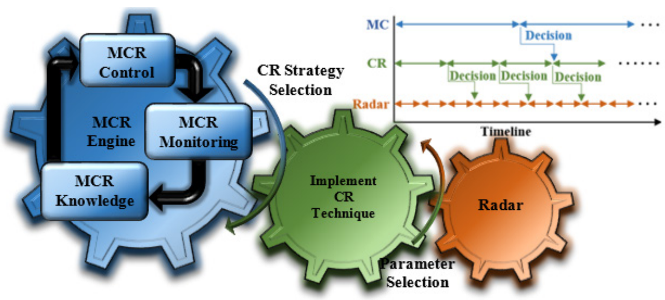

Fig. 9: Metacognitive radar model (MCR) to select the appropriate CR technique based on real-time observations of radar performance and the EME.

Fig. 9 also depicts the timeline of radar operation, CR technique implementation, and the MCR engine. The timeline of radar operation 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 a long 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 produces the "best" 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 greatest radar performance over time, consider a formulation inspired by the multi-armed bandit problem, a form of reinforcement learning that balances the exploration and exploitation trade-off. The exploration stage samples the potential options to measure the likelihood of success, while the exploitation stage uses the measured information to employ the option that optimizes performance under current conditions. For instance, [38] used the "explore-first" approach where the MCR engine first sequentially samples CR strategies to evaluate performance before exploiting the "best" option.

The method in [34] also explores how to rearrange the CR techniques into simpler categories for implementation, with Fig. 10 illustrating the tool box of solutions that are used for SDRadar. This architecture allows for the combining of different cognitive methods, waveform types, and radar signal processing, along with the tunable matching network in the Front-end. The methods in the Cognitive Strategy use the fundamental DSA approach from the CR techniques discussed in Section II that are arranged into the following generic categories: Fixed, React (Sense-React-Avoid), Learning (Sense-Learn-Avoid), and Stochastic (Sense-Predict-Avoid). The Waveform Type then determines the maneuverability of the waveform, the elements of which include full-band transmission, partial band transmission with LFM to avoid, random FM to avoid, and random FM to notch. The maneuverability for avoidance and notching is a form of frequency hopping that is dependent on the RFI. Finally, the Signal Processing is tailored for a generic tracking radar application. Each of these four categories switch between their constituent elements to provide ample combinations so that the appropriate processing path for a given RFI scenario can be chosen to maximize performance.

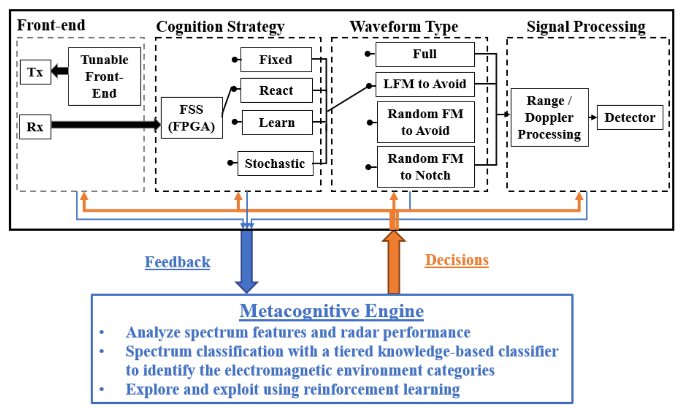

Fig. 10: Implementation architecture of CR techniques on the SDRadar.

This metacognitive engine monitors the spectral environment and radar performance. Spectral features are first extracted and classified using a tiered knowledge-based classifier to provide a "generalized classification" to down-select the CR techniques into a subset of possible candidates. The exploration/ exploitation reinforcement learning strategy is then used to implement these candidate techniques for real-time radar operation, where each candidate is evaluated and the best strategy chosen for exploitation.

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 particular EME conditions. Response time of a CR capability requires consideration, as does regulation of the perception-action cycle relative to the dynamic nature of the EME. Specifically, 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 determine the "best" approach for a given situation.

References

[1] 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.

[2] 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.

[3] K.L. Bell, C.J. Baker, G.E. Smith, J.T. Johnson, M. Rangaswamy, "Cognitive radar framework for target detection and tracking," *IEEE J. Selected Topics in Signal Proc.*, vol. 9, no. 8, pp. 1427-1439, Dec. 2015.

[4] A.F. Martone, K.D. Sherbondy, K.A. Gallagher, "Spectrum sensing for joint bandwidth/SINR radar waveform optimization," in *Radar and Communication Spectrum Sharing*, SciTech Publishing, Ch. 15, pp. 493 – 526, 2018.

[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] C. Horne, M. Ritchie, H. Griffiths, "Proposed ontology for cognitive radar systems," *IET Radar, Sonar & Navigation*, vol. 12, no. 12, pp. 1363-1370, Dec. 2018.

[7] S. Haykin, Y. Xue, P. Setoodeh, "Cognitive radar: step toward bridging the gap between neuroscience and engineering," *Proc. IEEE*, vol. 100, no. 11, pp. 3102-3130, Nov. 2012.

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

[9] 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.

[10] 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.

[11] M. Labib, J.H. Reed, A.F. Martone, A.I. Zaghloul, "Coexistence between radar and LTE-U systems: survey on the 5 GHz band," *USNC-URSI Radio Science Meeting*, Boulder, CO, Jan. 2016.

[12] 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.

[13] J.A. Kovarskiy, R.M. Narayanan, A.F. Martone, K.D. Sherbondy, "A stochastic model for prediction and avoidance of RF interference to cognitive radars," *IEEE Radar Conf.*, Boston, MA, Apr. 2019.

[14] 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," *to appear in IEEE Trans. Aerospace & Electronic Systems*, doi: 10.1109/TAES.2020.2987443.

[15] C.E. Thornton, R.M. Buehrer, A.F. Martone, "Experimental analysis of reinforcement learning techniques for spectrum sharing radar," *IEEE Intl. Radar Conf.*, Washington D.C., Apr. 2020.

[16] 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.

[17] S.D. Blunt, *et al.*, "Principles & applications of random FM radar waveform design," to appear in *IEEE Aerospace & Electronic Systems Mag*.

[18] J.A. Kovarskiy, J.W. Owen, R.M. Narayanan, S.D. Blunt, A.F. Martone, K.D. Sherbondy, "Spectral prediction and notching of RF emitters for cognitive radar coexistence," *IEEE Intl. Radar Conf.*, Washington, DC, Apr. 2020.

[19] A. Semnani, G.S. Shaffer, M.D. Sinanis, D. Peroulis, "High-power impedance tuner utilising substrate-integrated evanescent-mode cavity technology and external linear actuators," *IET Microwaves, Antennas & Propagation*, vol. 13, no. 12, pp. 2067-2072, Oct. 2019.

[20] A. Goad, A. Egbert, A. Dockendorf, C. Baylis, A. Martone, "Optimizing transmitter amplifier load impedance for tuning performance in a metacognition-guided, spectrum sharing radar," *IEEE Intl. Radar Conf.*, Washington, D.C., Apr. 2020.

[21] Alcala-Medel, *et al.,* "Fast frequency-agile real-time optimization of high-power tuning network for cognitive radar applications," *IEEE Intl. Microwave Symp*., Boston, MA, June 2019.

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

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

[24] V.C. Vannicola, J.A. Mineo, "Applications of knowledge based systems to surveillance," *IEEE Radar Conf.*, Ann Arbor, MI, Apr. 1988.

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

[26] 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.

[27] R. Oechslin, *et al.*, "Cognitive radar experiments with CODIR," *Intl. Conf. Radar Systems*, Belfast, UK, Oct. 2017.

[28] G. E. Smith *et al*., "Experiments with cognitive radar," *IEEE Intl. Workshop Computational Advances in Multi-Sensor Adaptive Processing*, Cancun, Mexico, Dec. 2015.

[29] A. Prabaswara, A. Munir, A.B. Suksmono, "GNU radio based software-defined FMCW radar for weather surveillance application," *Intl. Conf. Telecommunication Systems, Services, & Applications*, Bali, Indonesia, Dec. 2011.

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

[31] 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.

[32] J.W. Owen, B. Ravenscroft, S.D. Blunt, "Devoid clutter capture and filling (DeCCaF) to compensate for intra-CPI spectral notch variation," *SEE Intl. Radar Conference*, Toulon, France, Sept. 2019.

[33] 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.

[34] A.F. Martone, *et al.*, "Metacognition for Radar Coexistence," *IEEE Intl. Radar Conference*, Washington, D.C., Apr. 2020.

[35] 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.

[36] 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.

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

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