Towards Human-Machine Integration for Signals Intelligence Applications

Jonas D. Rockbach,
Luka-Franziska Bluhm
Human-Machine Systems
Fraunhofer FKIE
Wachtberg, Germany

Isabel Schlangen, Laura Over,
Sabine Apfeld
Sensor Data and Information Fusion
Fraunhofer FKIE
Wachtberg, Germany

Lukas Henneke,
Kevin Wilkinghoff
Communication Systems
Fraunhofer FKIE
Wachtberg, Germany

***Abstract*—Recent advancements in radar technology and telecommunications have led to a high population of increasingly complex emitters in the electromagnetic spectrum. From the perspective of a signals intelligence (SIGINT) operator, assessing the general situation comes with several major challenges, e.g. the detection and extraction of signals of interest, the elimination of unwanted interferences, or the meaningful representation of increasingly complex modulation patterns. This paper provides a general overview of the challenges of SIGINT and suggests possible ways to support operators with automated signal processing.**

***Index Terms*—Human-Machine Integration, Signals Intelligence, Electronic Intelligence, Communications Intelligence**

I. Introduction

Gathering intelligence information is a central skill of modern societies. In particular, signals intelligence (SIGINT) describes the extraction of information based on intercepted electromagnetic signals. Signal collection may be implemented at fixed locations or on moving platforms, such as on an high altitude long endurance unmanned aerial vehicle (HALE UAV). The main goal is to collect information about the capabilities of adversaries and assess the current situation using passive sensing techniques, hence avoiding deployment of units in access-restricted or hostile areas. SIGINT therefore plays a vital role in the collection of intelligence information. However, the electromagnetic spectrum (EMS) is increasingly populated with both civil and military users, such that relevant signals become more difficult to find and are often superimposed with other, less relevant transmissions. Therefore, efficient and reliable methods need to be found in order to cope with the increasing complexity of the EMS.

In general, SIGINT is divided into electronic intelligence (ELINT) and communications intelligence (COMINT). The goal of ELINT is to gather information about radar emitters to, e.g., infer their function or recognise previously seen radars by identifying characteristic features. Traditional ELINT analysis heavily relies on emitter databases which contain information about the operational modes of emitters that usually serve specific functions. However, agile multifunction radars do not make use of fixed modes anymore but choose their waveform parameters tailored to the encountered situation. Hence, static databases are not able to efficiently represent agile signals and new methods for the representation and analysis of modern multifunction radars are needed. Moreover, due the increasing occupancy of the EMS, human ELINT operators might also be overstrained with the amount of received signals and could therefore miss important activities. Here, machine learning provides solutions for representing and analysing highly complex information, as well as reducing repetitive manual processing tasks.

Corresponding author's contact email: jonas.rockbach@fkie.fraunhofer.de.

COMINT, on the other hand, is concerned with gathering intelligence from foreign communications signals transmitted over cable, radio or satellites [1], [2]. Besides analysing the pure content of a transmission which might be encrypted and therefore difficult or impossible to access, the transmission's metadata such as sender/receiver identification or their location is also of great interest. Similar to ELINT, COMINT operators need to extract relevant information from a vast amount of transmissions which might be corrupted by different kinds of noise or superimposed with other signals due to high occupancy of frequency bands. Moreover, the analysis of speech requires specialised personnel who are highly proficient in the target language and specifically trained to convert speech to writing and/or extract important information from the communication signals. The modular automatisation of speech analysis via machine learning can enable operators to process communications signals faster, possibly without the need of profound linguistic knowledge.

Apart from the benefits of automation, however, the application of machine learning to SIGINT also comes with challenges. The introduction of artificial intelligence (AI) can cause a fundamental change of the work dynamics, that may, in the worst case, even lead to a deterioration of performance [3], [4]. It is therefore important to selectively support human cognition with AI, rather than blindly automating tasks based on technical feasibility alone. If human operators are supported by AI in selected tasks, their workload is generally expected to decrease, nonetheless the supervision of the introduced machine agent also creates additional effort. Supervising an artificial agent requires the ability to interpret an algorithm to assess its performance and adjust system parameters accordingly. Therefore, meaningful integration of the *human in the loop* and explainable AI are complex but important design issues.

Given the potential benefits, this work considers the balanced integration of human and artificial agents for AI-supported SIGINT. Sec. II first describes the abstract computational nature of SIGINT that is true for both human and artificial agents. In Sec. III, the potential of human-machine resource allocations and issues of automation supervision are discussed for SIGINT human-machine integration. A more detailed view on different automation techniques in the ELINT and COMINT task models is given in Sec. IV and V, respectively. Sec. VI provides a summary.

II. Computational Nature of SIGINT

A. Formalisation of SIGINT

SIGINT can be viewed from the perspective of hybrid intelligence for human-machine integration, which designs an integrated human-machine agent rather than isolated humans and machines [3], [5]. The computational objective of this SIGINT hybrid is the extraction and report of tactical information (e.g. identity, location, affiliation) based on passive measurements of the EMS (Fig. 1). The OODA loop (observe, orient, decide, act, Fig. 1) is an intelligent agent model that has been applied to hybrid agents [6]. Thus, the SIGINT agent needs to *observe* the EMS, update its model of the world (*orient*), and decide where to search in the EMS and what to report (*decide* and *act*). Note that the active part of the OODA loop in a SIGINT system refers solely to the adaptation of its sensors to the situation to improve the surveillance performance. However, the SIGINT agent generates reports that may be used by other systems to trigger effector employment, which requires specific methodological and ethical considerations [7] that fall outside the scope of this paper.

More formally, a SIGINT agent should strive to minimise the gap between the real world and its current model of it based on its mission objectives (effectiveness) while minimising the time and resources until relevant observations are reported (efficiency). This process can be formulated in terms of a so-called Quality of Service resource allocation problem with $K \in \mathbb{N}$ tasks or objectives [8], [9]: Each mission objective can be described as a utility function $u_k(q_k)$, where $q_k(r_k, e_k)$ denotes the quality of task performance given the assigned resources r_k and possible environmental parameters e_k. In case of the SIGINT agent, q_k is the similarity between the real world and the extracted model of the situation. Assuming that the tasks are independent, the SIGINT agent needs to solve the following constrained optimisation problem:

$$\text{maximise total utility } u(r) = \sum_{k=1}^{K} w_k u_k(q_k(r_k, e_k)) \qquad (1)$$

$$\text{subject to limited resources } \left(\sum_{k=1}^{K} r_k\right) - r_{\max} \leq 0, \qquad (2)$$

where $r = \{r_1, \ldots, r_K\}$ is the collection of assigned human and artificial resources, $r_{\max}$ is the total of available resources, and the weights $w_k \in [0, 1]$ satisfy $\sum_{k=1}^{K} w_k = 1$.

The mission objectives (or utility functions u_k) act as an attention filter that states which observations are of interest to the SIGINT agent. The derived SIGINT model of the world can be defined as a set of platforms (either stationary stations or moving land, air, space, or sea vehicles) that must carry at least one active emitter on them in order to be detected by the SIGINT agent (Fig. 1). Such sets of interacting platforms form specific scenarios in an area and timeframe of interest, e.g., an air defence training.

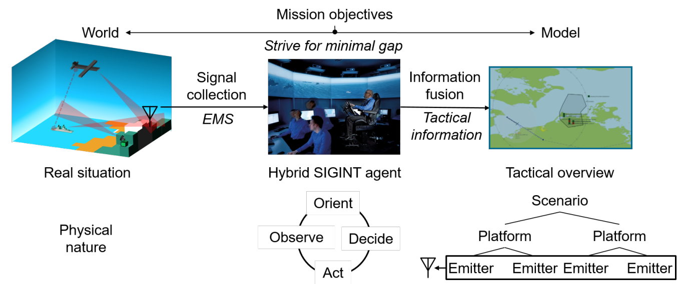

Fig. 1. Illustration of SIGINT.

B. Fusion and Control Architecture of SIGINT

Fig. 2 shows the abstract information fusion and control architecture of the SIGINT agent which is based on the Joint Directors of Laboratories (JDL) data fusion model [10]. At the interface to the world, participating agents of the SIGINT hybrid search for signals, measure these, perform direction findings of emitters bearings, and analyse the measured signals (signal acquisition, Sec. II-C). This level can be seen as the sensor organs of the hybrid agent [5]. On the next level, the individual results are fused in either ELINT or COMINT resulting in a specialised tactical overview. The third level, that represents the cognitive level, correlates ELINT and COMINT outputs to a fused SIGINT overview of the situation which is the basis for mission assessment and management.

In the OODA design, feedback loops modulate operations of previous tasks, e.g. based on the optimisation (1)–(2). This holds for all hierarchical control levels: The SIGINT level adjusts the behaviour of the ELINT and COMINT modules that in turn modulate individual signal acquisition activities, always striving to maximise the utility u_k. In general, the OODA feedback loop may be performed at each level of the data fusion model, i.e. control actions do not necessarily need to come from the highest level of SIGINT control but also from subordinate ELINT and COMINT knowledge levels. Moreover, utility can be provided in different ways, e.g. by mathematically defined functions for artificial agents (hard data) or via performance feedback by the operator (soft data).

C. Task Structure of Signal Acquisition

Signal acquisition is the basic skill of the SIGINT hybrid agent at the low level of the layered JDL information fusion (Fig. 2). The task structure of signal acquisition can be divided into four steps, as shown in Fig. 3. First, the EMS is searched for signals of interest that are defined by the mission objectives (e.g., confirm the location of a certain platform). Using the direction of arrival of the collected signals, it is possible to

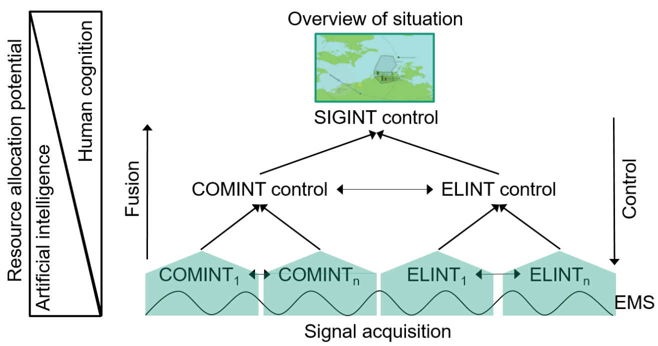

Fig. 2. Illustration of the abstract SIGINT fusion and control architecture.

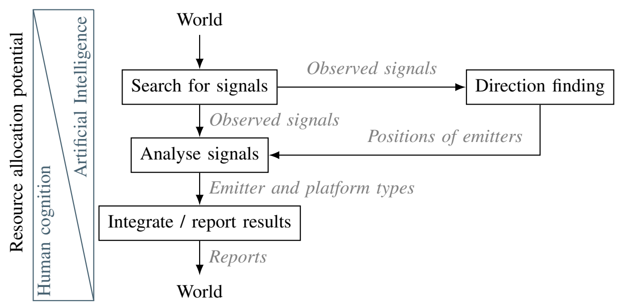

Fig. 3. High-level task model of signal acquisition.

determine the position of the emitter over time, given that the motion of the observing platform is of higher order than the emitter's [11]. The signal itself is analysed for its technical attributes and tactical information where identified emitter locations might be used as contextual knowledge. Finally, the extracted information is integrated with other results and reported (e.g., the platform is at location x in a defence mode).

III. Human-machine Integration

A. Resource Allocation Potential

The integration of human and machine resources to a hybrid agent requires the selective allocation of tasks to the involved parties [5], [12] to optimise (1)–(2). The level of automation describes the degree to which a task is executed by a machine agent. Three basic classes of automation can be identified [12]:

- In *manual mode*, the human operator performs the task without automated assistance.
- In *supervised mode*, the machine agent performs the task under supervision by the human operator who can influence the behaviour of the machine.
- In *autonomous mode*, the machine agent performs the task on its own while the operator has no possibility to influence the operation.

In comparison, modern (weak) AI methods often outperform humans in specialised tasks close to the sensor such as waveform analysis, while humans provide a broad spectrum of competencies that are most efficient for high-level tasks such as strategic reasoning [5]. Human cognition is especially tuned to fuse heterogeneous information to a joint assessment of the world, integrating experiences and handling expectations [13] that are not available to machines. These natural cognitive capabilities can be supported by human-in-the-loop automation (supervised mode) for observation and action [3] because human cognition builds on the organism's observation-action capabilities [14]. Orientation and decision may as well be supported by AI in the form of decision support, however these tasks benefit largely from natural human cognition that is integrative and flexible (manual and supervised mode). A purely autonomous task allocation is almost never a preferred choice in complex scenarios since unpredictable situations may decrease the performance of the artificial agent over time [4]. In addition, in-the-loop control is also central for understanding machine dynamics [3], [15] (Sec. III-B).

Fig. 2 includes a summary of the expected benefit of human and machine resource allocations in the fusion and control architecture. It shows that AI can provide substantial support on tasks that are closer to the sensor, while high-level tasks in the JDL model and the related decision making should remain with human operators. Likewise, the task allocation potential is highlighted for the signal acquisition model in Fig. 3.

B. Factors for Automation Supervision

The introduction of machine intelligence introduces both a reduction of operator workload in the form of support and an overhead in the form of machine supervision [3]. It is therefore important to understand the contributing factors of machine supervision [12] in order to judge whether an introduced automation is actually beneficial to the hybrid agent.

Fig. 4 shows a simplified model of automation supervision that also follows the OODA-loop. First, the operator actively interacts with the machine to gain an understanding of its state (situation awareness) [3], [16]. This requires sufficient information and possibilities to manipulate the machine agent, e.g. through dynamic reparametrisation or explainable AI [5], [17], [18]. Second, based on the observed machine state, operators calibrate their trust towards the provided results in accordance with the general automation capabilities. Here, both under-trust and over-trust may lead to non-optimal hybrid performance [15]. With over-trust, false machine results might be used, while under-trust can lead to the dismissal of the machine agent although its output would be beneficial. In trust calibration, sufficient human-in-the-loop experience with the system is essential. Finally, the operator exerts control over the automation based on the selected trust [3], either by overwriting automation results or changing machine responsibilities which should both be integrated in the system design [5].

To summarise, the introduced automation should not only have a high accuracy, but should also:

- provide information about the state of the automation and its capabilities (explainability),
- facilitate the calibration of operator trust towards the automation (human in the loop),
- provide the possibility to manipulate automation operations (adjustability and control).

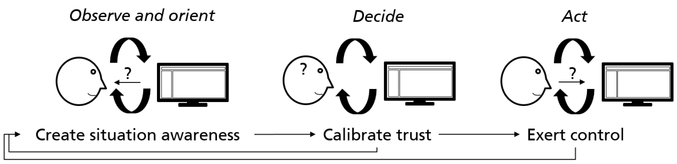

Fig. 4. Simplified model of automation supervision

IV. Electronic Intelligence

An ELINT system consists of one or more passive receivers intercepting radar signals for further analysis. Usually, such a receiver is able to cover at least the frequencies from 2 GHz to 18 GHz which are divided into bands of the receiver's instantaneous bandwidth. This bandwidth is chosen as a compromise between probability of intercept, sensitivity, size, weight, and cost, and ranges from 0.5 GHz to 2 GHz. When a signal is received, the measured electromagnetic waves are transcribed into so-called pulse descriptor words (PDWs) which encode at least the radio frequency (RF), pulse width, and time of arrival leading to the respective pulse repetition interval (PRI).

Reconnoitring the full frequency range manually is very demanding. On the one hand, the spectrum is occupied with a variety of sources, many of which are irrelevant for the operation. On the other hand, sources might be relevant but analysed previously, hence claiming the operator's attention regularly without yielding new information. As a result, many frequencies are monitored redundantly and at once, which increases the risk of missing relevant signals. In line with the discussed human-machine resource allocation potential (Sec. III-A), these shortcomings can be addressed via automation techniques that adopt repetitive low-level tasks. With this, the operator potentially gains more time to monitor unusual activities and perform situation or mission assessment tasks. Following Fig. 3, this section gives an overview of ELINT-specific automation methods for signal search and analysis.

A. Radar Signal Search

In order to intercept a radar signal, the ELINT receiver must be tuned to the correct frequency band at the time it is illuminated by the radar. If the emitter and ELINT receiver both use periodic scanning patterns, synchronisation might occur, i.e. the receiver might always be tuned to a specific frequency band while the radar is transmitting away from it, hence impeding an intercept. To maximise the probability of intercept and avoid synchronisation, approaches for designing search strategies have been proposed, e.g. [8], [19]–[22]. However, these methods are only useful if the radars to be intercepted perform a periodic scan of the environment with a constant RF. If the radar is beam-agile, synchronisation is very unlikely and hence the probability of intercept does not depend on the search strategy [22]. An electronically steered array antenna might even be easier to intercept than a mechanically rotating radar. Therefore, an automatic signal search should use a strategy that is optimised to avoid synchronisation with periodically scanning emitters and beam-agile radars will be intercepted as by-catch. If a frequency-agile radar with an agility band larger than the ELINT receiver's instantaneous bandwidth should be observed over a longer period of time, information about the expected next emission is useful to adapt the receiver's search dwells. The works [23]–[26] suggest neural-network-based approaches for predicting upcoming radar signals that could be used for optimisation of the search strategy and hence aid the operator in gaining a better overview of the signal spectrum. However, ensuring the transparency of neural networks is challenging from a human-machine integration perspective (Sec. III-B). Possible solutions lie in both global and local approaches to interpretability [17].

B. Radar Signal Analysis

Since more than one emitter might be active simultaneously, the received PDW stream needs to be deinterleaved first, i.e. the PDWs need to be sorted by emitter. If available, bearing measurements facilitate deinterleaving in a first step. After that, the PDWs are usually sorted by common statistical properties, see e.g. [27] for an introduction. Deinterleaving is a necessary step that needs to be completed before the emitters can be further analysed and classified. Traditional methods mostly rely on histogram techniques, e.g. [28], [29], which do not provide satisfactory results if the waveform parameters are not constant. More recently, methods that use neural networks for deinterleaving have been proposed, e.g. [30], [31].

A common goal of ELINT analysis is to identify the emitter type. Traditionally, this is achieved by matching the parameters of the measured signal against a database. For agile emitters, several methods based on machine learning have been proposed, e.g. [23], [26], [32]–[34]. Automatic human-in-the-loop emitter type identification can decrease the workload of the operator, who can focus on the signals from those emitters that are relevant to the mission. Hence, the risk of missing important signals is reduced. However, automatic emitter type identification might not always be sufficiently accurate, e.g. due to a low signal-to-noise ratio (SNR) or the presence of unknown emitters or signals. Instead, received signals might be classified according to relevance, e.g. based on decision trees, without necessarily assigning an ID. Although this requires manual identification, the workload can be reduced as only the relevant signals need to be processed.

Usually, classification techniques using neural networks are trained to recognise a set of known classes. Detecting that an intercepted signal belongs to an unknown emitter requires modifications to the usual approach. A possible solution using open set recognition is described in [26], [35]. If a signal unknown to the automatic identification system is detected, manual analysis can be supported by machine learning approaches. One example is the automatic recognition of the PRI modulation which provides useful information about the possible function and capabilities of the radar. The publications [36]–[38] present such approaches using neural networks. As discussed above, it is important to consider the transparency of the neural network approaches as well [17].

V. COMMUNICATIONS INTELLIGENCE

To extract intelligence information from communication signals, COMINT operators follow the high-level task model of SIGINT shown in Fig. 3. The use of machine learning promises to support the operator. In this section, we discuss the steps *signal search* and *signal analysis* of RF signals.

A. Signal Search

The interception of RF signals poses several challenges. First, signal impairments like fading due to the radio propagation environment and interfering signals in the same or adjacent frequency bands have to be considered. Furthermore, in non-cooperative scenarios, substantial path losses can occur, resulting in low SNRs and therefore making signals of interest difficult to detect. Also, low probability of intercept approaches like spread spectrum techniques may be used by communication partners for obfuscation purposes.

Signal search can be divided into two subtasks: monitoring a single RF and searching in a wide frequency band. For both tasks, there are several automated detection approaches, which can be extended by AI methods to improve performance [39]. For frequency monitoring, threshold-based energy detectors are commonly used but often falsely detect interfering signals. Using a deep neural network to identify the typical waveform of signals of interest can increase accuracy [39]. Especially for signal detection in wide frequency bands, signal feature detectors are used in addition to energy detectors [40]. Instead of being dependent on expert knowledge for the time-consuming derivation of suitable signal features, neural networks can learn these characteristics from signal data directly [41]. For example, AI-based object detection in images can be adapted to detect and classify signals in broadband spectrograms [42] which can also provide explainability [17] (Sec. III-B).

B. Signal Analysis: Digital Transmissions

In non-cooperative communications, in-depth signal analysis is necessary to gain signal knowledge. The analysis starts with recognizing the modulation technique. Then, symbol rate estimation and synchronisation are necessary to extract modulation symbols (demodulation) and decode symbols to bit sequences. When error correction codes are used, error correction has to be performed as well before interpreting the bit stream according to the transmission protocol. However, encryption can prevent extraction of information.

Machine learning can also enhance signal analysis capabilities, e.g. in modulation classification [41], [43]. Moreover, transformer architectures are able to learn the synchronisation of RF signals [44]. Many digitally modulated signals are highly structured due to repeating bit sequences used for synchronisation or resulting from packet headers. Such structures allow to identify transmission protocols using only raw signal data without demodulation. Neural networks can learn defining features of the signal structure which help linking detected transmissions to known emitters. In addition, approaches for protocol-agnostic RF device fingerprinting are emerging [45].

C. Signal Analysis: Speech

When analysing communication signals containing speech, several challenges need to be overcome. Noisy radio transmissions result in the loss of information. Moreover, speakers often have varying dialects and are not cooperative [46]. This is more challenging for machines in comparison to human experts who are specifically trained to understand spoken content in difficult acoustic conditions. However, training personnel is costly and human resources are usually limited, making it difficult to analyse all signals in full detail. Here, automatic systems can help to harness human competences more efficiently by filtering for relevant information.

Communication signals containing speech can be automatically analysed for specific content with varying degrees of abstraction. For this purpose, multiple algorithms are combined in a flexible processing chain [47]. Usually, the first step is to detect the presence of speech, called voice activity detection [48]. In a second step, voiced segments can be assigned to languages (language recognition) [49]. Specific speakers can be searched for (speaker recognition) [50] or changes of speakers can be detected (speaker diarisation) [51]. Furthermore, one can search for specific keywords (keyword spotting) [52] or recognise all words being spoken (automatic speech recognition) [53], [54]. Finally, the recognised text can be translated from a source to a target language such that operators can easily understand all spoken content (machine translation) [55]. Note that these steps get increasingly more complex and thus more difficult, each requiring more labeled training data to be able to yield useful results. Especially in the military domain, acquiring sufficient amounts of training data in the right acoustic conditions can be difficult to achieve while ensuring an acceptable level of performance [46]. Therefore, human cognition should be utilised at the complex steps as proposed by the resource allocation potential (Sec. III-A).

VI. CONCLUSION AND OUTLOOK

This work discusses possible approaches for AI-supported signals intelligence. In general, AI support at tasks close to the sensor is assumed to improve signals intelligence performance, while higher-level tasks benefit from natural human cognition. Potential AI support is discussed for both electronic and communications intelligence. Importantly, the application of AI should be evaluated not only in terms of machine performance, but also in its potential for human-machine integration such as explainability. A central challenge in SIGINT hybrid agent design is the adaptive allocation of human and artificial resources.

REFERENCES

[1] H. Martinek, *Communications Intelligence: Transcriber Output Enhancement*. US Army Research Institute for the Behavioral and Social Sciences, 1980, vol. 1268.

[2] U. Uebler and H.-J. Kolb, "Signals intelligence-processing-analysis-classification," MEDAV GmbH Uttenreuth (Germany), Tech. Rep., 2009.

[3] E. Hollnagel and D. D. Woods, *Joint Cognitive Systems. Foundations of Cognitive Systems Engineering*. CRC Press, 2005.

[4] D. D. Woods and E. Hollnagel, *Joint cognitive systems: Patterns in cognitive systems engineering*. CRC Press, 2006.

[5] J. D. Hasbach and T. E. Witte, "Human-Machine Intelligence: Frigates are Intelligent Organisms," in *2021 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*, 2021.

[6] N. Stanton, C. Baber, and D. Harris, *Modelling command and control: Event analysis of systemic teamwork*. Ashgate Publishing, Ltd., 2008.

[7] F. Santoni de Sio and J. Van den Hoven, "Meaningful human control over autonomous systems: A philosophical account," *Frontiers in Robotics and AI*, p. 15, 2018.

[8] S. Apfeld, A. Charlish, and W. Koch, "Quality of Service Resource Management for Search Strategy Design in Electronic Support," *Sensor Signal Processing for Defence (SSPD)*, 2017.

[9] A. Charlish, F. Hoffmann, C. Degen, and I. Schlangen, "The development from adaptive to cognitive radar resource management," *IEEE Aerospace and Electronic Systems Magazine*, vol. 35, no. 6, pp. 8–19, 2020.

[10] D. Hall and J. Llinas, *Handbook of multisensor data fusion*. CRC press, 2001.

[11] S. Blackman and R. Popoli, *Design and analysis of modern tracking systems*. Norwood, MA: Artech House, 1999., 1999.

[12] T. B. Sheridan, *Telerobotics, automation, and human supervisory control*. MIT press, 1992.

[13] J. Hohwy, *The predictive mind*. Oxford University Press, 2013.

[14] R. Pfeifer and J. Bongard, *How the body shapes the way we think: a new view of intelligence*. MIT press, 2006.

[15] J. D. Lee and K. A. See, "Trust in automation: Designing for appropriate reliance," *Human factors*, vol. 46, no. 1, pp. 50–80, 2004.

[16] M. R. Endsley, "Toward a theory of situation awareness in dynamic systems," *Human factors*, vol. 37, no. 1, pp. 32–64, 1995.

[17] A. Adadi and M. Berrada, "Peeking Inside the Black-Box: A Survey on Explainable Artificial Intelligence (XAI)," *IEEE Access*, vol. 6, pp. 52 138–52 160, 2018.

[18] C. M. Burns and J. Hajdukiewicz, *Ecological interface design*. CRC Press, 2004.

[19] I. V. L. Clarkson, "Optimal Periodic Sensor Scheduling in Electronic Support," *Defence Applications of Signal Processing*, 2005.

[20] ——, "Optimisation of Periodic Search Strategies for Electronic Support," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 47, no. 3, pp. 1170–1784, 2011.

[21] C. Winsor and E. J. Hughes, "Optimisation and evaluation of receiver search strategies for electronic support," *IET Radar, Sonar & Navigation*, vol. 6, no. 4, 2012.

[22] I. V. L. Clarkson, "Intercepting Beam-Agile Radar," *IEEE International Conference on Radar*, 2018.

[23] Z.-M. Liu and P. S. Yu, "Classification, Denoising, and Deinterleaving of Pulse Streams With Recurrent Neural Networks," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 55, no. 4, pp. 1624–1639, 2019.

[24] S. Apfeld, A. Charlish, and G. Ascheid, "The Value of Memory: Markov Chain versus Long Short-Term Memory for Electronic Intelligence," *IEEE Radar Conference*, 2021.

[25] S. Apfeld and A. Charlish, "Ensembles of Predictive Radar Models for Electronic Intelligence," *IEEE Radar Conference*, 2022.

[26] S. Apfeld, "Machine Learning for Electronic Intelligence," Dissertation, RWTH Aachen University, 2021, available online at https://publications.rwth-aachen.de/record/834559.

[27] R. Wiley, *ELINT: The Interception and Analysis of Radar Signals*. Artech House, 2006.

[28] H. Mardia, "New techniques for the deinterleaving of repetitive sequences," *IEE Proceedings F Radar and Signal Processing*, vol. 136, no. 4, 1989.

[29] K. Nishiguchi and M. Kobayashi, "Improved algorithm for estimating pulse repetition intervals," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 36, no. 2, pp. 407–421, 2000.

[30] S. Gasperini, M. Paschali, C. Hopke, D. Wittmann, and N. Navab, "Signal Clustering With Class-Independent Segmentation," *IEEE International Conference on Acoustics, Speech and Signal Processing*, 2020.

[31] Z.-M. Liu, "Pulse Deinterleaving for Multifunction Radars With Hierarchical Deep Neural Networks," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 57, no. 6, 2021.

[32] S. A. Shapero, A. B. Dill, and B. O. Odelowo, "Identifying Agile Waveforms with Neural Networks," in *International Conference on Information Fusion*, 2018.

[33] P. Notaro, M. Paschali, C. Hopke, D. Wittmann, and N. Navab, "Radar Emitter Classification with Attribute-specific Recurrent Neural Networks," *pre-print arXiv: 1911.07683*, 2019.

[34] X. Li, Z. Liu, Z. Huang, and W. Liu, "Radar Emitter Classification with Attention-Based Multi-RNNs," *IEEE Communications Letters*, vol. 24, no. 9, pp. 2000–2004, 2020.

[35] S. Apfeld and A. Charlish, "Recognition of Unknown Radar Emitters With Machine Learning," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 57, no. 6, 2021.

[36] J. P. Kauppi, K. Martikainen, and U. Ruotsalainen, "Hierarchical classification of dynamically varying radar pulse repetition interval modulation patterns," *Neural Networks*, vol. 23, no. 10, pp. 1226–1237, 2010.

[37] Y. Liu and Q. Zhang, "An improved algorithm for PRI Modulation Recognition," in *International Conference on Signal Processing, Communications and Computing*, 2017.

[38] X. Li, Z. Liu, and Z. Huang, "Attention-Based Radar PRI Modulation Recognition With Recurrent Neural Networks," *IEEE Access*, vol. 8, 2020.

[39] Q.-V. Pham, N. T. Nguyen, T. Huynh-The, L. Bao Le, K. Lee, and W.-J. Hwang, "Intelligent radio signal processing: A survey," *IEEE Access*, vol. 9, pp. 83 818–83 850, 2021.

[40] K. Holladay, M. A. Koets, A. Burmeister, R. Dollarhide, and D. Varner, "A configurable signal analyzer for embedded systems," in *Military Communications Conference*, vol. 2, 2003, pp. 1173–1177.

[41] M. Kulin, T. Kazaz, I. Moerman, and E. D. Poorter, "End-to-end learning from spectrum data: A deep learning approach for wireless signal identification in spectrum monitoring applications," *IEEE Access*, vol. 6, pp. 18 484–18 501, 2018.

[42] T. O'Shea, T. Roy, and T. C. Clancy, "Learning robust general radio signal detection using computer vision methods," in *2017 51st Asilomar Conference on Signals, Systems, and Computers*, 2017, pp. 829–832.

[43] X. Li, F. Dong, S. Zhang, and W. Guo, "A survey on deep learning techniques in wireless signal recognition," *Wirel. Commun. Mob. Comput.*, 2019.

[44] T. J. O'Shea, L. Pemula, D. Batra, and T. C. Clancy, "Radio transformer networks: Attention models for learning to synchronize in wireless systems," in *50th Asilomar Conference on Signals, Systems and Computers*. IEEE, 2016, pp. 662–666.

[45] I. Agadakos, N. Agadakos, J. Polakis, and M. R. Amer, "Chameleons' oblivion: Complex-valued deep neural networks for protocol-agnostic RF device fingerprinting," in *European Symposium on Security and Privacy*. IEEE, 2020, pp. 322–338.

[46] A. Cornaggia-Urrigshardt, F. Gökgöz, F. Kurth, H.-C. Schmitz, and K. Wilkinghoff, "Speech recognition lab," in *International Conference on Military Communications and Information Systems*, ser. Procedia Computer Science, vol. 205. Elsevier, 2022, pp. 218–228.

[47] A. Cornaggia-Urrigshardt, N. Jarocky, F. Kurth, S. Urrigshardt, and K. Wilkinghoff, "SCALA-Speech: An interactive system for finding and analyzing speech content in audio data," in *52. Jahrestagung der Gesellschaft für Informatik*, ser. LNI, vol. P-326. GI, 2022, pp. 81–90.

[48] M. Sharma, S. Joshi, T. Chatterjee, and R. Hamid, "A comprehensive empirical review of modern voice activity detection approaches for movies and TV shows," *Neurocomputing*, vol. 494, pp. 116–131, 2022.

[49] M. Li, W. Cai, and D. Cai, "Survey talk: End-to-end deep neural network based speaker and language recognition," in *20th Annual Conference of the International Speech Communication Association*. ISCA, 2019.

[50] Z. Bai and X. Zhang, "Speaker recognition based on deep learning: An overview," *Neural Networks*, vol. 140, pp. 65–99, 2021.

[51] T. J. Park, N. Kanda, D. Dimitriadis, K. J. Han, S. Watanabe, and S. Narayanan, "A review of speaker diarization: Recent advances with deep learning," *Comput. Speech Lang.*, vol. 72, p. 101317, 2022.

[52] I. López-Espejo, Z. Tan, J. H. L. Hansen, and J. Jensen, "Deep spoken keyword spotting: An overview," *IEEE Access*, vol. 10, pp. 4169–4199, 2022.

[53] J. Li *et al.*, "Recent advances in end-to-end automatic speech recognition," *APSIPA Transactions on Signal and Information Processing*, vol. 11, no. 1, 2022.

[54] Z. Zhang, J. T. Geiger, J. Pohjalainen, A. E. Mousa, W. Jin, and B. W. Schuller, "Deep learning for environmentally robust speech recognition: An overview of recent developments," *ACM Trans. Intell. Syst. Technol.*, vol. 9, no. 5, pp. 49:1–49:28, 2018.

[55] S. Maruf, F. Saleh, and G. Haffari, "A survey on document-level neural machine translation: Methods and evaluation," *ACM Comput. Surv.*, vol. 54, no. 2, pp. 45:1–45:36, 2021.