ARL-SR-0429 • MAY 2020

Artificial Intelligence in Synthetic Biology, Cyber Defense, and Aeromechanical Design

by Bryan Glaz, Cliff Wang, Margaret Hurley, and Alexander Kott

NOTICES

Disclaimers

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

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

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

ARL-SR-0429 ● MAY 2020

Artificial Intelligence in Synthetic Biology, Cyber Defense, and Aeromechanical Design

Bryan Glaz
Vehicle Technology Directorate, CCDC Army Research Laboratory

Cliff Wang
Army Research Office, CCDC Army Research Laboratory

Margaret Hurley
Sensors and Electron Devices Directorate, CCDC Army Research Laboratory

Alexander Kott
Office of the Director, CCDC Army Research Laboratory

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

1. REPORT DATE *(DD-MM-YYYY)*	2. REPORT TYPE	3. DATES COVERED (From - To)
May 2020	Special Report	September 2019–May 2020

4. TITLE AND SUBTITLE
Artificial Intelligence in Synthetic Biology, Cyber Defense, and Aeromechanical Design

5a. CONTRACT NUMBER

5b. GRANT NUMBER

5c. PROGRAM ELEMENT NUMBER

6. AUTHOR(S)
Bryan Glaz, Cliff Wang, Margaret Hurley, and Alexander Kott

5d. PROJECT NUMBER

5e. TASK NUMBER

5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
CCDC Army Research Laboratory
ATTN: FCDD-RLD
2800 Powder Mill Rd
Adelphi, MD 20783

8. PERFORMING ORGANIZATION REPORT NUMBER
ARL-SR-0429

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

10. SPONSOR/MONITOR'S ACRONYM(S)

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

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

13. SUPPLEMENTARY NOTES
ORCID ID(s): Kott, 0000-0003-1147-9726

14. ABSTRACT
This report summarizes the outcomes of the three FY20 Army Science Planning and Strategy Meetings led by the US Army Combat Capabilities Development Command Army Research Laboratory. The intent of these annual meetings, conducted since 2013, is to explore novel scientific opportunities that may lead to providing the Army with an advantage in future conflicts. The meetings focus on identifying research gaps and barriers that may hinder the achievement of potential novel capabilities and exploring possible approaches to overcoming these gaps and barriers. This report covers the findings and recommendations developed during meetings held in fall 2019 and early 2020. These meetings concentrated on the use of artificial intelligence in such areas as synthetic biology, cyber defense, and the design of aeromechanical flying systems.

15. SUBJECT TERMS
artificial intelligence, machine learning, autonomy, cyber defense, synthetic biology, aeromechanical design

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			Alexander Kott
Unclassified	Unclassified	Unclassified	UU	36	19b. TELEPHONE NUMBER (Include area code) (301) 394-1507

Contents

Acknowledgments

The authors gratefully acknowledge the support of the US Army Combat Capabilities Development Command Army Research Laboratory Director's Office and the many participants, all of whom contributed to a thought-provoking and highly productive series of meetings.

The authors thank the editor Carol Johnson for organizing and editing the manuscript of this report.

Additional acknowledgments are included in individual sections of the report.

1. Introduction

Author: Alexander Kott

During the first half of fiscal year 2020 (FY20), the US Army Combat Capabilities Development Command (CCDC) Army Research Laboratory (ARL) organized and conducted four meetings aimed at exploring novel scientific opportunities that may lead to providing the US Army with an advantage in future conflicts. Such meetings have been conducted regularly since 2013 and are called Army Science Planning and Strategy Meetings (ASPSMs).

The temporal scope of these explorations is strategic in nature, with the time horizon being about 20 to 30 years. The meetings focus on identifying research gaps and barriers that may hinder the achievement of potential novel capabilities and exploring possible approaches to overcoming these gaps and barriers. Numerous research efforts—in-house, collaborative, and extramural—have been initiated or revectored based on insights developed during the ASPSMs.

This report covers the findings and recommendations developed during three meetings held in the first half of FY20. One other meeting is documented in a separate classified publication.

1.1 Motivation

For the last 10 years, artificial intelligence (AI) and especially neural-network-based machine learning (ML) have experienced a wave of dramatic growth in their popularity, funding, applications, and research and development (R&D) efforts. As a result, every field of science and technology (S&T) faces the question, Is there an opportunity for remarkable advances at the intersection of that field and AI?

In particular, synthetic biology (SB) has also demonstrated extraordinary growth in the last decade. Potential applications of AI to SB range from the immediate and obvious (detection and mitigation of SB-enabled weapons), to those with more far-reaching and less-understood implications such as modification of the human genome, a topic that might be upon us sooner than anticipated. It is important for the Army S&T community to explore use of ML and related techniques to accelerate SB design in environmental remediation, robotic hybrids, and energy generation for longer-term applications such as the design of artificial brains.

Current air vehicle design paradigms are based on combinations of human intuition and empirical iteration coupled with existing designs. Although this approach has led to important successes, small unmanned aerial systems (UASs) severely lag the performance of biological flyers. Realizing the former could provide the Army with

a disruptive capability. Research offers growing evidence that in the future we will be able to synthetically fabricate actuators and materials rivaling biological functionality, energy efficiency, and information processing, and with low signatures. This leads to the question, Do the AI-based design paradigms exist to put such components together into a system that ultimately rivals a biological flyer?

In the field of cyber warfare, both defensive and offensive operations continue to rely largely on human involvement. Future cyber systems will have to rely more on varying levels of automation for both decision making and response, especially in Army systems, since autonomous systems will fight cooperatively and collaboratively with Soldiers. Such systems will need to withstand attacks in both the electronic warfare and cyber domains without demanding much attention from human Warfighters. Furthermore, future Army autonomous systems will need to learn, adapt, and maintain awareness specific to Army missions.

1.2 Key Findings

The meeting titled "Artificial Intelligence and Synthetic Biology" recommended, inter alia, the following:

- Increase automation in the laboratory to improve throughput. This will include implementation of a laboratory information management system (LIMS) to improve data/metadata capture.
- Improve database capabilities. This will require an increase in available archival storage size and speed, and increased maintenance. This will enable multiple uses of data, rather than our traditional single-use experimentation protocol. This also involves a culture change to open sharing of "unsuccessful" results. Knowing what did not work can be just as useful as knowing what did.
- Improve networks for improved data sharing between remote sites. Multi-site collaborative efforts are the norm and data flow must be unrestricted.

The meeting titled "Artificial Intelligence in Design of Aeromechanical Systems" noted and recommended the following:

- The application of big data and ML has been effective in understanding and modeling some complicated mechanical systems; however, the ability to generate mechanical systems from the ground up, rather than learn from data produced by a built-up system, has not yet been demonstrated.
- In order to reach the long-term goal of AI/ML design of novel platforms, the scientific research community must first develop the approaches for the

AI/ML design of subsystems underpinned by novel theoretical mechanics formalisms.

- The CCDC Army Research Laboratory should focus on formulation of example problems that will illustrate how going beyond classical mechanics constructs can enable novel capabilities. These mechanics and theoretical foundations will serve as benchmarks that can then be used in an AI-inspired mechanical actuator/platform design challenge problem.

The meeting titled "AI in Cyber Autonomy" yielded a number of findings and recommendations, including the following:

- Even highly autonomous cyber-defense systems cannot exist in an environment where they operate independently of human defenders. For the actions of an AI system to be trusted by humans, these actions have to be adequately explained in terms that humans can understand without rigorous technical training. More importantly, these autonomous systems must be designed and maintained with the operational constraint of seamless teaming with human analysis. This quite often requires the creation of a joint mental model shared among autonomous systems and human operators, and a supervised decision process with the goal of enhancing cyber-defense effectiveness.
- Automated, coordinated detection and response offer opportunities to leverage the abilities of better-resourced or better-placed (in the network or in the software ecosystem) organizations to gain an advantage in fighting adversaries and help their coalition partners.
- Autonomous cyber-defense systems need a high level of cyber situational awareness. The goal should be a self-aware, autonomous cyber system that is capable of introspection, is state aware, and can take corrective actions to mitigate or recover from compromises. The ability to identify past failures and recover from them will bring a new level resiliency to our cyber systems.
- Data-driven learning through AI, combined with human expertise, offers an opportunity to reverse the asymmetric advantages that our adversaries have enjoyed for a long time so that future cyber defense will be transformed from being more reactive to being more proactive in nature.

2. Artificial Intelligence and Synthetic Biology

Author: Margaret Hurley

2.1 Introduction

SB and AI have been intertwined in the human imagination long before the recognition of either as a distinct scientific field. This entanglement is to some extent unavoidable as "intelligence" itself is recognized as a biological construct with considerable overlap in language and logic. It is also an acknowledgment that higher "intelligence" and extreme measures are required to harness the seeming chaos of living systems. How much chaos is by natural design and how much is a matter of perception remains a persistent question. It is hoped that the knowledge accumulated from well-planned, well-executed studies in AI/SB will help answer this question once and for all.

Both AI and SB have undergone a very long evolutionary period, culminating in simultaneous explosive growth over the last decade. The global SB market was estimated at approximately USD $1.1 billion in 2010. Current estimates put the 2020 market value at USD $6.8 billion and five-year growth to 2025 at USD $19.8 billion. The latter is a low-end estimate that has been projected to be much higher by other analysts. However, among these rosy projections, there is a call to arms. Multiple reports, including the most recent National Academies of Sciences report on "Safeguarding the Bioeconomy" (2020), conclude that the United States, long a leader in biological sciences, is set to lose primacy as a result of strategic Chinese investment and the sheer size of their trained workforce.

Current applications of SB have been successful and transformative but have largely been confined to areas immediately suited to the technology, namely, commodity chemical production, bio-based materials (including biofuels), and medicine and health. These applications are self-evident "low hanging fruit" and required no great predictive powers to forecast. However, interviews with 25 SB experts at the Synbiobeta 2018 conference (an annual meeting dedicated to celebrating progress, establishing networks, and providing support for SB practitioners in all areas of industry, academia, and government) show a much broader vision for potential use even in the near- to mid-term. Twenty-five SB subject-matter experts (SMEs) were asked to comment on potential areas of impact, current challenges, and the most outstanding application of the technology to date. The SME response on potential areas of greatest impact included further advancement in health and medicine (including neuroscience, regenerative medicine, tailored medicine, and tissue engineering), agriculture (including pest control, genetically modified organisms [GMOs], and improvements in natural

fertilizers), biomaterials (including production of biocomposites and biofuels, as well as biomining for remediation and repair and design of electronic devices), and microbial community engineering. The assembled list of greatest challenges included finding sufficient funding for development of innovative technologies, ethical use of a powerful tool, the innate difficulties in control alluded to previously, the multidisciplinary nature of the field, and limitations in vision for a burgeoning field whose potential is only starting to be tapped. The list of most outstanding current applications truly showed the breadth of interests in the field, as answers ranged from current production of biomaterials (the conference stage was literally made from mushroom-produced materials by Ecovative Design), to universal influenza vaccine design (a highly relevant topic during the coronavirus disease 2019 [COVID-19] pandemic of spring 2020, where the SB community data sharing and support has been outstanding), to attempts to clone the wooly mammoth based in the Church laboratory at Harvard University, to development of the Impossible Burger. An additional potentially powerful use was demonstrated in 2017 when the Church laboratory published results encoding an animated GIF into and retrieving it from bacterial DNA, designated in the popular media as a "living hard drive".

The applications listed represent a snapshot of the short-term beneficial impact of SB. However, while the benefits of the technology have been long been recognized, the dual-use nature of the field and the potential for darker purposes have also been widely recognized. Until recently, the concerns surrounding SB have been similar to concerns for traditional biology, heightened by the power of SB and its popularity as methods became more affordable and more accessible. These concerns have largely centered on biosafety, biosecurity, and ethical concerns. A spectacular example of the need for development and enforcement of regulation of the technology occurred in 2018 with the announcement of the birth of the first genome-edited human babies. While the 3-year prison sentence given to biophysicist He Jiankui sends a strong message against misuse, the CRISP-Cas9 (clustered regularly interspersed short palindromic repeats associated protein 9) gene-editing technology remains widely available and easy to use, and community lab spaces and grassroots initiatives for biohackers (such as biocurious.org and counterculturelabs.org) abound.

This has also resulted in an expansion of the concerns for safety and control of SB into the cyber realm, and the advent of the field of cyberbiosecurity. To underscore this, we note that in 2017, the same year that SB was used to develop a "living hard drive", biohackers demonstrated the ability to use SB to encode malware into DNA and control the computer attached to the sequencer. While the success of this hack required modifications to the system, it still underscores the need for vigilance in understanding and controlling of all aspects of the technology. It also underscores

the importance of teams (or individuals) with expertise in both the biological *and* cyber realms.

The discussion so far has focused largely on accomplishments in SB without the catalyst of AI. The need to incorporate AI into the SB "Design–Build–Test–Learn" cycle is widely recognized as a necessary step to streamline the entire process. However, the combination of AI and SB, the empowerment of already powerful technologies, heightens all of the promise and all of the perils of each. Furthermore, the line between the fields is blurring, as demonstrated by the development of synthetic biochemical circuits capable of information processing.

2.2 Objective and Scope

The potential impact of AI/SB touches a wide range of the Army's interests both immediate and long term. From sensing and mitigation of SB-enabled weapons to in situ production of bio-based materiel to implanted microchips to enhance Warfighter physical and cognitive capabilities, the Army has a use for and a need to monitor the technology. With this in mind, the ASPSM on Artificial Intelligence in Synthetic Biology was held at the ARL NE offices in Burlington, Massachusetts, on November 13, 2019. This full-day event brought together researchers from industry, academia, and other government agencies (OGAs), as well as diverse Army laboratories to discuss the state of the practice. Speakers included the following:

- Representatives of Argonne National Laboratory and the US Department of Energy Agile BioFoundry, a national-lab-led consortium to improve scalability, predictability, and cost-effectiveness of bio-based production of commodity chemicals and biofuels. This effort specifically incorporates improvements in database applications and ML.
- Representatives of the US Department of Energy Joint BioEnergy Institute (JBEI), collaborators in the Agile BioFoundry, whose mission includes development of improved engineered bioenergy crops and development of methodologies to streamline and optimize biofuel production and performance. As part of this effort, the JBEI and Agile BioFoundry introduced the Experimental Data Depot (EDD), an online repository of experimental data and metadata.
- Representatives from Carnegie Mellon University's (CMU's) Computational Biology Department, now offering a master's of science degree in automated science and proponents of model-driven protocols to

cover the entire system parameter space while reducing necessary experimentation.

- Representatives from Raytheon BBN technologies, working on standardization of parts and procedures for automated experimental design and knowledge sharing.
- Representatives from Colorado State University and GenoFAB Inc., who are working to promote usage of AI in SB with their LIMS and Electronic Notebook.
- Representatives from the California Institute of Technology and developer of the first DNA-based artificial neural network (ANN).

The meeting concluded with a round-table discussion among all attendees of long-term goals, hazards, and a "wish list" from each of the speakers.

2.3 Gaps, Requirements, and the Path Forward

Proper implementation of AI in SB requires achieving a critical mass of data. This is a central feature of *any* implementation of AI, but the extreme variability in system response inherent in biological systems exacerbates the requirement. Accordingly, a mandatory feature for laboratories striving to increase throughput is to use automation to streamline all facets of the process from experimentation to data collection to data analysis. Several practitioners made the analogy to self-driving cars, as self-driving laboratories and learn-guided experiments will become prevalent. There are multiple benefits to automating the process in addition to increasing the overall amount of data. Although, as noted, biological response is an inherently variable entity, multiple practitioners observed that variations between sites and operators account for a major percentage of data variation. Automation will minimize this variability and may allow researchers to home in on the roots of natural variation. As one attendee noted, "The future is in control theory".

Concomitant with a need for improved data quality (numerically speaking) is a growing need for improved metadata curation to make optimal use of the result. Automation combined with improvements in data quality and quantity are expected to facilitate the move toward ML-guided experimentation and reduce the overall numbers of runs required for a given project, as demonstrated in the recent literature where researchers were able to develop a predictive model of the effect of a set of small molecule drug candidates on protein subcellular localization with only 29% of all possible experimentation performed. Given sufficient efficiency in automation, attendees envisioned "real-time bioengineering" and a future where

researchers are freed from the laboratory and allowed to focus on experimental design and analysis.

In addition to improvements in data quantity and quality, implementation of AI/SB requires improvements in the type of available data. Improvements in instrumentation continually add to the list of available experimental measurables, including recent noteworthy advances in single-cell omics technology, which allow researchers to study natural cellular heterogeneity, thus reinstating one of biology's central features (diversity) as a strength rather than a weakness. In addition to technological requirements in data typing, there is a requirement for cultural change. Overwhelming amounts of viable data remain unreported as scientists are taught to focus on and publish only positive results. Researchers tend to internalize the process of learning from Thomas Edison's "10,000 ways that won't work". Publication of these negative results will open that parameter space to ML and other algorithms, and allow future work to draw from that data. It is also acknowledged that the current model for scientific success revolves around high-impact results published in high-impact journals. Success of the community as a whole also depends on producing and publishing more pragmatic results in standardization, scaling up production, and the scanning protocol variable space.

AI protocols are continuing to improve with their ever-wider application in a wider variety of fields. However, it is advisable for practitioners to more openly discuss their choice of algorithm in publications, and we particularly wish to draw attention to results presented in this workshop leveraging the strengths of eight different AI techniques. Work continues to improve algorithms to disentangle cause from correlation and practitioners are reminded that in the AI world one size does not fit all. Along with this, improved treatment of uncertainty is necessary, as has been noted in other ASPSM reports.

Collaboration (both experimental and analytical) across multiple physical sites will become more common in the future. This will lead to requirements for improved networking quality and capability, data storage and management, and visualization capabilities, as well as greater requirements in cybersecurity and cyberbiosecurity.

One point of contention that surfaced during discussion was training of scientists in the future. While broader training in AI and data analytics among the scientific community as a whole was generally agreed upon as beneficial, no consensus was reached on whether it is better to field a team of specialists or individuals with a broader background, including general computational biologists. A corollary to that is the discussion of where to draw the line on automation of analysis. Proponents contend that a well-mannered AI and well-structured database relieve researchers of the requirement to be programmers as well. Others see that as a goal for the

future, as trained experts in computational biology and data analysis are still a much needed component of any team.

The need for responsible research and ethics training was not explicitly called out, although it was inherent in the afternoon discussion. It is recognized in the community that in addition to natural intelligence and AI, we are on the verge of something beyond either, a potential superintelligence. There is an obvious reluctance among researchers to develop anything beyond human control, with an intelligence surpassing that of the designer. If Allen Turing can be taken by surprise with great frequency by machines of his own devising, then what hope have we?

The final wish-list discussion was actually centered on near- and mid-term capabilities rather than long-term vision. Attendees wished to develop designer microbiomes, solutions to metabolomics, and artificial cells, and discover cancer biomarkers. These capabilities require attaining the levels of system control that we currently strive for and are necessary to move into the rich application fields (bioproduction, biomaterials, medicine, and the rest) discussed in the introduction. With these tools, the path to tailored medicine, artificial organs, living prosthetics, artificial skin, in situ biomateriel production and remediation, and enhanced human capabilities (strength, vision, intelligence) becomes open to us. Devices (including sensors) may be expanded from wearables to ingested or implanted. Biological computing may be the next wave past quantum computing, and biological circuits may be the ultimate in ecofriendly hardware. Longer-term applications for AI/SB are difficult to forecast as only those comfortable with a technology wield it with any proficiency. We focus instead on the immediate goal, to develop the technology.

2.4 Recommendations

Recommendations for ARL and the Army at large are the following:

- Increase automation in the laboratory to improve throughput. This is actively being done under the ARL's TRANSFORME Essential Research Program (ERP), but the requirement will be ongoing as technology improves. This will include implementation of a LIMS to improve data/metadata capture.
- Improve database capabilities. This will require an increase in available archival storage size and speed, and increased maintenance. This will allow multiple uses of data rather than our traditional single-use experimentation protocol. This is also currently being done under the TRANSFORME ERP.

This also involves a culture change to open sharing of "unsuccessful" results. Knowing what did not work can be just as useful as knowing what did.

- Improve networks for improved data sharing between remote sites. Multi-site collaborative efforts are the norm and data flow must be unrestricted.
- Work with the AI *and* SB communities to standardize protocols and improve data quality.
- Continue working with AI practitioners in other technical areas to follow improvements in the field and find areas of overlap. Foster discussions of algorithm choice, limitations, and needs.

2.5 Acknowledgments

The organizers gratefully acknowledge the spirited participation and substantial contributions from all attendees. The organizers particularly wish to thank the invited speakers for their efforts:

- Jacob Beal: Raytheon BBN Technologies
- Fusun Yaman: Raytheon BBN Technologies
- Robert Murphy: CMU
- Hector Garcia Martin: Lawrence Berkeley Laboratory
- Philip Laible: Argonne National Laboratory
- Jean Peccoud: Colorado State University
- Lulu Qian: California Institute of Technology

We also wish to thank attendees from ARL, Meagan Small, Deborah Sarkes, Jessica Terrell, Asha Hall, Yelena Sliozberg, and Dean Culver; Matthew Lux from CCDC Chemical Biological Center (CBC); Kathleen Swana of CCDC Soldier Center (SC); and Zaira Martin Moldes and Rachel Parker of Tufts University. We wish to thank Robert Jensen and the staff of ARL NE for their assistance. Drs Sarkes and Small are gratefully acknowledged for organizational help, documentation assistance, and service as an invaluable sounding board.

2.6 Bibliography

Bartley BA, Beal J, Karr JR, Strychalski EA. Organizing genome engineering for the gigabase scale. Nature Communications. 2020;11:689. doi: 10.1038/s41467-020-14314-z.

Beal J, Adler A, Yaman F. Managing bioengineering complexity with AI techniques. Biosystems. 2016;148:40–46. doi: 10.1016/j.biosystems.2015.08.006.

Bianchini F. The problem of prediction in artificial intelligence and synthetic biology. Complex Systems. 2018;27(3):249–265. doi: 10.25088/ComplexSystems.27.3.249.

Bostrom N. Superintelligence: paths, dangers, strategies. Oxford (UK): Oxford University Press; 2014.

Carbonell P, Radivojevic T, Garcia Martin H. Opportunities at the intersection of synthetic biology, machine learning, and automation. ACS Syn Bio. 2019;8:1474–1477. doi: 10.1021/acssynbio.8b00540.

Cherry KM. Qian L. Scaling up molecular recognition patterns with DNA-based winner-take-all neural networks. Nature. 2018;559:370–376. doi: 10.1038/s41586-018-0289-6.

Cole SD, Beabout KK, Turner KB, et al. Quantification of interlaboratory cell-free protein synthesis variability. ACS Syn Bio. 2019;8(9):2080–2091. doi: 10.1021/acssynbio.9b00178.

Costello Z, Garcia Martin H. How to hallucinate functional proteins. Q Bio QM. 2019 Mar. Preprint, arXiv:1903.00458.

Cyranoski D. What CRISPR-baby prison sentences mean for research. Nature. 2020;577:154–155. doi: 10.1038/d41586-020-00001-y.

Deans T. Parallel networks: synthetic biology and artificial intelligence. ACM Journal on Emerging Technologies in Computing Systems. 2014;11(21). doi: 10.1145/2667229.

El Karoui M, Houyos-Flight M, Fletcher L. Future trends in synthetic biology—a report. Front Bioeng Biotechnol. 2019. doi: 10.3389/fbioe.2019.00175.

Larsen PE, Zerbs S, Laible PD, Collart FR, Korajczyk P, Dai Y, Noitor P. Modeling the pseudomonas sulfur regulome by quantifying the storage and communication of information. mSystems. 2018;3:e00189—17. doi: 10.1128/mSystems.00189-17.

Mesko B. The future of extremism: artificial intelligence and synthetic biology will transform terrorism. Santa Clara (CA): Futurism; 2016 Nov 14 [accessed 2020 May]. https://futurism.com/the-future-of-extremism-artificial-intelligence-and-synthetic-biology-will-transform-terrorism.

Morrell WC, Birkel GW, Forrer M, et al. The Experimental Data Depot: a web-based software tool for biological experimental data storage, sharing, and visualization. ACS Syn Bio. 2017;6:2248–2259. doi: 10.1021/acssynbio.7b00204.

Murch RS, So WK, Buchholz WG, Raman S, Peccoud J. Cyberbiosecurity: an emerging new discipline to help safeguard the bioeconomy. Front Bioengin Biotechnol. 2018;6:39. doi:10.3389/fbioe.2018.00039.

Naik AW, Kangas JD, Sullivan DP, Murphy RF. Active machine-learning driven experimentation to determine compound effects on protein patterns. eLife. 2016. doi: 10.7554/eLife.10047.

National Academies of Sciences, Engineering, and Medicine. Safeguarding the bioeconomy. Washington (DC): The National Academies Press; 2020. https://doi.org/10.17226/25525.

Opgenorth P, Costello Z, Okada T, et al. Lessons from two design-build-test-learn cycles of Dodecanol production in *Escherichia coli* aided by machine learning. ACS Syn Bio. 2019;8:1337–1351. doi: 10.1021/acssynbio.9b00020.

Petersen P, Tikhomirov G, Qian L. Information-based autonomous reconfiguration in systems of interacting DNA nanostructures. Nat Comm. 2018;9:5362. doi: 10.1038/s41467-018-07805-7.

Prabhune M. The power of synthetic biology: 25 thought leaders opine. The Bench. Redwood City (CA): Synthego; 2019 Jan 17 [accessed 2020 May]. https://www.synthego.com/blog/synthetic-biology-applications.

Shipman SL, Nivala J, Macklis JD, Church GM. CRISPR-Cas encoding of a digital movie into the genomes of a population of living bacteria. Nature. 2017;547:345–349. doi: 10.1038/nature23017.

Stuart T, Satija R. Integrative single-cell analysis. Nat Rev Gen. 2019;20:257–272. doi:10.1038/s41576-019-0093-7.

Wang F, Zhang W. Synthetic biology: recent progress, biosafety and biosecurity concerns, and possible solutions. J Biosafety Biosecurity. 2019;22–30. doi: 10.1016/j.jobb.2018.12.003.

3. AI in Design of Aeromechanical Systems

Author: Bryan Glaz

3.1 Introduction

Current air vehicle design paradigms are based on combinations of human intuition and empirical iteration upon existing designs. While this approach has led to large manned platforms that have no biological peers, it has led to small UASs (smaller than class 2) that severely lag the performance of biological flyers. However, if platforms rivaling biological adaptation and performance could be designed, they would provide the Army with a disruptive capability in Multi-Domain Operations (MDO).

The move toward bio-inspired flight for UASs has been active for approximately 30 years. Much has been learned over this time and much data has been collected (ranging from biologists to engineers). However, there are recent basic science movements ongoing (e.g., biohybrid robotics, complex systems[1]) that indicate that one day we will be able to synthetically fabricate actuators and materials rivaling biological functionality, energy efficiency, and information processing, with low signatures. This leads to the question—even if one could synthetically fabricate such materials and embedded controllers, do the AI-based design paradigms exist to put such components together into a system that ultimately rivals a biological flyer? This ASPSM focused on this question.

3.2 Objectives and Scope

This ASPSM brought together scientists in aeromechanics, robotics, AI and ML, and evolutionary design to explore the scientific opportunity for developments in AI (e.g., Silver et al.[2]) to lead to new paradigms in mechanics, with the ultimate goal of novel platform designs rivaling biological capabilities. The focus of the meeting was on aeromechanical systems, as there is a large capability gap between fielded and currently envisioned state-of-the-art platforms and biological flyers. The objectives of the discussion were to formulate new basic research directions that would ultimately lead to design synthesis tools that are orthogonal to current design approaches and would offer the possibility for outside-the-box platform designs underpinned by novel mechanics. Furthermore, exploring design studies will help to inform where advancements in fields like materials, biochemistry, controls, actuators, and so on are critical to overall design.

3.3 Background

A confluence of scientific developments sparked the formation of this ASPSM. Recent developments in AI/ML for general reinforcement learning algorithms have demonstrated approaches that learn without human expert data and eventually exceed human strategies.[2] However, typically, such advancements have been demonstrated in the context of board games (at least in the open literature). Clearly, it would be valuable to develop a similar capability for the design of novel platforms, where AI/ML approaches would develop novel strategies (e.g., designs) that exceed human strategies. However, as in Silver et al.,[2] a clear set of rules and objectives need to be given to the AI/ML tools to allow for development of new strategies related to platform design. Such investigations would be new within the broader aeromechanics research ecosystem, and perhaps, even within the broader mechanical sciences not limited to platform design.

In parallel to developments in AI/ML, there has been an increased understanding in morphological design and adaptation,[3] and emerging biohybrid robotics perspectives on actuation and platform mechanics.[2] These developments are beginning to show pathways toward physical platform development (e.g., actuators, morphology, design of embedded control systems), as well as design principles that may explain some of the reasons why biological organisms are far superior to human-engineered unmanned platforms in areas such as adaptability, robustness, energy efficiency, silent actuation, agility, and so on. While valuable, these perspectives only give a top-down viewpoint; for example, they quantify how already-designed platforms (biological, human-engineered, or hybrid) compare in certain aspects, even those that were previously difficult to quantify, as they do not come from the physics realm but rather the information theory realm. In order to push the S&T further and exploit these developments for advanced capabilities, the science of bottom-up mechanics and platform design needs to be developed. Here, bottom-up refers to the ability to go from basic mechanics to synthesis, from the actuator and component level up to the platform design level, in order to meet performance objectives. Thus, the difference in top-down versus bottom-up is analogous to understanding an already-existing platform versus designing and building a platform from the most basic component/subsystem up to the platform design given some performance goals.

Despite advancements from these seemingly orthogonal research ecosystems, there is an opportunity to exploit for advanced Army capabilities by some convergence between these separate fields. To initiate this long-term research strategy, the ASPSM focused on the development of a challenge problem to gauge the state of the art and inform ARL leadership on the distribution of basic and applied research

investments that are needed. The challenge problem would seek to answer questions such as whether methods such as those developed in Silver et al.[2] can be applied to the aircraft design problem directly (implying applied research to technology development is the appropriate investment path), or whether basic/applied scientific advancements in areas such as machine intelligence for inference (and so on) are still needed before AI/ML generative design tools can outperform current state-of-the-art government/industry design codes. To explicitly gauge the state of the art, the ASPSM focused on the guidelines and scope of the challenge problem, which will be released by ARL to the community afterward. However, one of the key findings from the meeting after substantial discussion was that scientific research into AI/ML for synthesizing novel theoretical mechanics-based systems was needed before a platform design challenge problem could be addressed. Specifically, it was the group's opinion that an intermediate step of using AI/ML tools to synthesize novel mechanics, such as those seemingly present in nonequilibrium actuators found in biology, was needed before platform design could be addressed. Then, once the actuator/subcomponent design problem was matured, the platform design problem could then be addressed.

3.4 Gaps and Recommendations

Several observations, technical gaps, and recommendations came out of the ASPSM discussions and breakout sessions. Some general observations were the following:

- The application of big data and ML has been effective in understanding and modeling some complicated mechanical systems, particularly when augmenting empiricism. For instance, physics-informed neural networks are being used to predict crack growth and health monitoring metrics in aircraft structures, and are proving effective in connecting empirical trends with continuum mechanics models. However, the ability to generate mechanical systems from the ground up, rather than learn from data produced by a built-up system, has not yet been demonstrated.
- Morphological computation and embodied intelligence are ongoing schools of thoughts for describing adaptiveness in autonomous systems; the next step is understanding how to formulate these metrics as objective functions/constraints that AI could operate on.
- Before platforms design with capabilities that rival biological systems can be tackled with AI/ML approaches, the underlying mechanics enabling highly capable biological systems need to be understood and synthesized by AI/ML approaches. It is clear that attractive attributes such as agility, energy

efficiency, thermal management, robust actuation, among others, derive from fundamentally different theoretical mechanics foundations than classical mechanics. For example, the theoretical mechanics of nonequilibrium molecular motors give biological actuators (e.g., muscle tissue) many of their attractive mechanical performance and control system effectiveness. Thus, in order to reach the long-term goal of AI/ML design of novel platforms, the scientific research community must first develop the approaches for AI/ML design of subsystems underpinned by novel theoretical mechanics formalisms.

3.4.1 Some Key Research Questions

The following are some key research questions we determined:

- How do we design for adaptation, rather than over-designing for specific mission objectives (e.g., designing a flyer that can fly fast and straight but also can aggressively navigate through a forest like a bird)?
- How do we mathematically formalize structures that cut across domains, such as logic based to physics based? Category theory? Operator theory? Can AI synthetically form topologies in the absence of needed information? How can AI infer across "manifolds" to solve a physics problem that humans struggle to solve; for example, can AI infer that novel materials need to be combined with origami-inspired mathematics to achieve form factors and weights resembling a biological flyer?
- What are the formalisms that should be the foundation for allowing AI/ML to discover mechanics outside the bounds of classical mechanics? This mathematical formalism should be able to accommodate both physical canonical coordinates and nonphysical information theoretic coordinates.

3.5 Conclusions

There are significant challenges to achieving ARL's long-term goal of establishing the scientific foundations for future platform design tools that can go beyond human-engineered systems. However, this ASPSM has contributed to the development of the path forward. As a result of the ASPSM, ARL will focus on formulation development of example problems that will illustrate how going beyond classical mechanics constructs can enable novel capabilities. These mechanics and theoretical foundations will serve as benchmarks that can then be used in an AI-inspired mechanical actuator/platform design challenge problem. Once these example problems are formulated and published, ARL will then return to the formulation of an AI challenge problem so that the state of the art in AI/ML

approaches can be examined for their use in the novel design of platforms underpinned by nonclassical mechanics.

3.6 Acknowledgments

The ARL organizers gratefully acknowledge the substantial contributions of the many ARL, academic, and industry colleagues who participated in the ASPSM discussions.

3.7 References

1. Ricotti l, Trimmer B, Feinberg AW, Raman R, Parker K, Bashir R, Sitti M, Martel S, Dario P, Menciassi A. Biohybrid actuators for robotics: a review of devices actuated by living cells. Science Robotics. 2017;2(12). doi: 10.1126/scirobotics.aaq0495.

2. Silver D., Hubert T, Schrittwieser J, et al. A general reinforcement learning algorithm that masters chess, shogi, and go through self play. Science. 2018;362(6419). doi: 10.1126/science.aar6404.

3. Corrucci F, Cheney N, Kriegman S, Bongard J, Laschi C. Evolutionary developmental soft robotics as a framework to study intelligence and adaptive behavior in animals and plants. Front Robot AI. 2017;4(34). doi: 10.3389/frobt.2017.00034.

4. AI and Cyber Autonomy

Author: Cliff Wang

(Disclaimer: This report is based on the workshop report generated by the academic co-organizers, Profs Matt Fredrikson and Lujo Bauer, both from CMU.)

4.1 Summary

The combination of a rapidly expanding attack surface and increasingly sophisticated and determined attackers suggests a new focus in computer security on defenses that react and adapt with no human intervention and that maximize the impact of human defender effort. At the same time, AI and ML algorithms have had a transformative impact on many real-world applications. This, in turn, suggests that it could be beneficial to explore how AI and ML can be harnessed for cyber defense.

A one-day ASPSM workshop on cyber autonomy at CMU brought together leading researchers from academia and experts from government to generate ideas for what research was needed to develop a new generation of cyber-defense systems and algorithms. Some of the themes that emerged from the workshop included the following:

- Integration between AI systems and human defenders is both necessary and a key challenge; research is needed to better explain how AI systems make decisions and take advantage of human input.
- Systems might benefit from AI being an integral part of system design rather than an add-on service.
- Coordination of detection and response should take place across systems and administrative boundaries, which requires developing solutions for federated learning and collaboration.
- AI algorithms often learn from available data, but in cyber-defense settings the data may be under attacker control; approaches need to be developed that minimize the attacker's ability to influence how AI systems learn.
- Autonomous systems need to have situational awareness and introspective abilities so that they can react to unforeseen situations.
- To mitigate the asymmetric advantage adversaries typically have over defenders, autonomous cyber-defense systems should use resources efficiently and be able to efficiently revert to known secure states.

- Rigorous formal specification is needed to precisely define the goals of autonomous systems and the restrictions under which they should operate.

While some of these themes are already becoming apparent in ongoing research efforts within various academic communities, discussions at the workshop tended to acknowledge that future research will benefit from addressing these themes together rather than in isolation.

4.2 Motivation and Approach

The intersection of several trends and technical developments in computing suggests that new approaches are needed in cybersecurity. These trends and developments include the following:

- Attackers are increasingly sophisticated, motivated, and well resourced. *They also can take the time to carry out multi-step reconnaissance and planning before launching a decisive attack.*
- The attack surface that needs to be defended is expanding rapidly, with the number of computing devices connected to the Internet increasing by an order of magnitude in the space of a *decade and the vast heterogeneity of devices.*
- AI and ML have matured to the point where they present opportunities for better defenses but also bring about new risks. The risks arise because AI and ML could be used by attackers to mount more successful attacks, as well as because AI and ML algorithms could themselves be susceptible to new kinds of attacks.

This confluence of trends suggests a need to examine whether current approaches to cybersecurity will remain effective and investigate what additional approaches are needed to address new threats or take advantage of new opportunities. For example, it seems necessary to focus on developing defenses that are increasingly automated, since human (defender) involvement may often be incompatible with the scale and pace of new attacks. However, automated techniques that take advantage of opportunities enabled by AI must consider new attack vectors and classes of vulnerabilities that might follow.

Research topics relevant for this discussion and represented by workshop attendees included the following:

- autonomous or semi-automated network defenses;
- automated identification and remediation of software vulnerabilities;

- applications of game theory for cybersecurity;
- other applications of AI and ML to cybersecurity;
- adversarial ML;
- human factors in cybersecurity; and
- security analytics.

The workshop's goals were to help determine the new research needed, a roadmap and order of priorities, and barriers to and enablers of new research directions and results that will lead to advanced capabilities in cyber defense and help the Department of Defense (DOD) to defeat our adversaries.

The workshop consisted of a half day of invited short presentations and a half day of focused discussion in parallel breakout sessions on specific topics. The invited short presentations were designed to inform all workshop participants of cutting-edge research results and the perspectives of leading researchers in the fields relevant to the workshop topic; the focused discussion was designed to first brainstorm and then more closely evaluate ideas for promising research directions. The workshop culminated in a general session with presentations by breakout session leads.

4.3 Workshop Discussion and Conclusions

The working groups focused on determining what new research was needed—and what directions were most promising—to enable increased autonomy in cyber defense and what were the key challenges that needed to be overcome to accomplish this. The three parallel breakout groups largely converged on a set of topics that describe promising directions for harnessing autonomy and AI for cyber defense. A compilation of the ideas generated by the working groups follows.

4.4 Human Impact/Explainability

A pervasive theme of the discussions was that even highly autonomous cyber-defense systems cannot exist in an environment where they operate independently of human defenders. The need to include humans as part of a defense system arises in several areas:

- Humans are necessary to specify the goals that a system for defense or response should meet. Doing so implies the development of methods to provide such specifications and ensure that autonomous systems or components meet the specifications.

- Human analysts will often have insights that automated algorithms cannot (at least until the method for deriving each "insight" is taught to the algorithm). Hence, it is necessary to develop methods to best take advantage of human experts and amplify their abilities. *It is important for autonomous systems to inherit "human expertise" as part of their foundation.*
- For the actions of an AI system to be trusted by humans, these actions have to be adequately explained in terms that humans can understand without rigorous technical training, including in cases when AI systems are fundamentally very difficult to understand. *In addition, the AI systems need to be designed or created such that they can support effective human–AI system teaming.*
- Humans are a necessary complement to AI systems (as discussed previously) but could also introduce new vulnerabilities that weaken composite human+AI systems. The ways in which AI systems interface with users may need to be accounted for (e.g., humans providing incorrect inputs to AI components; or humans interpreting or acting on the outputs of an AI system in a way that is not purely objective, for example, trusting some outputs more than others independently of their accuracy).

4.5 AI-Embedded Architecture View

The goal of autonomous cyber defense is to maintain the normal (acceptable) state of the cyber system against internal/external dynamic conditions (threats and failures) through a chain of asynchronous interactions between cyber components for distributed coordination, and without a human being in the loop. However, developing such an autonomous cyber-defense system that exhibits self-awareness and distributed coordination/collaboration will require an architectural view in which AI is integrated into the design of the system rather than as an add-on service to allow for positive and emergent behavior.

4.6 Automated, Coordinated Detection and Response

Defense and response are often implemented as separate, loosely coupled phases, both because of lack of integration between the different mechanisms that implement them and because both detection and response may involve actions that cross administrative boundaries.

The vast majority of network defenses today operate within the scope of a single administrative domain. Where cross-domain interaction occurs, it is performed among human administrators or, when automated, is limited to simple patch

distribution or information gathering by one tool (e.g., malware-encounter data by an antivirus product). Federated learning offers new opportunities for collectively training and responding to new attacks across multiple administrative domains. Going further, the prospect of allowing active, automated collaboration in defense has been rarely explored, but offers opportunities to leverage the abilities of better-resourced or better-placed (in the network or in the software ecosystem) organizations to assist others. This sharing could be part of normal operation, respecting traditional administrative boundaries and organizational privacy. Or, it might become more fluid and permissive in emergency situations, subject to audit to compensate for the extraordinary access leveraged to diagnose and remediate the threat. Disciplines core to this vision include federated (and privacy-preserving) learning, statistical data privacy, methods to support privileged access with monitoring/auditing, and robust distributed coordination protocols.

4.7 Forecasting and Learning in Adversarial Ecosystems

The cyber-ecosystem is constantly evolving—the adversaries change their strategies when they are detected, and the defense mechanisms and the network properties change over time. It is important to answer the following questions. Can we predict what attackers are likely to do? How will the attacker adapt? What kill chain is the attacker in? How do we build automated responses that countenance these predictions + environmental conditions? How do we reason how the attacker will adapt to a response? Using previously observed large-scale data of attacks and defenses gives a novel opportunity to create advanced ML models to forecast how these changes happen. This will power novel emulators that can help to model future attacks on our systems and lead to the creation of robust and defendable systems.

4.8 Introspection and "Self-Aware" Autonomous Systems

Autonomous cyber-defense systems need a high level of cyber situational awareness. To determine the best way to respond to an attack or situation, the system needs to be self-aware of its current state, then it can determine which actions to take. The goal should be a self-aware autonomous cyber system that is capable of introspection, is state aware, and can take corrective actions to mitigate or recover from compromises. The system should have a certain level of confidence to certify correctness (i.e., measurable and explainable performance) and explain the rationale of the course of action.

4.9 Mitigating the Asymmetric Advantage of the Adversary

Asymmetry is an important research area because the asymmetry of attackers versus defenders is a fundamental problem. Unless this asymmetry is changed, the attacks will continue to dominate defenses. In short, any defensive strategy must take steps to mitigate this advantage. This can be done in several ways. Autonomous cyber defense needs to be able to appropriately respond to the adversary with efficient use of resources. It should be able to return a system and/or network to a known state from an arbitrary/compromised state. At the very least, this will force the adversary to repeat the attack, perhaps at a higher cost. However, this requires the defender to establish trust in a few baseline secure states.

4.10 Specification-Driven Defense for Verification and Automation Synthesis

The use of automated and semi-automated verification methods has recently met with success in efforts to provide comprehensive security properties such as memory safety, control flow integrity, and functional correctness for critical software components such as cryptographic libraries. One can argue that the development of more rigorous cyber defenses requires the use of precise specification formalisms and techniques for verifying implementations against them. This is particularly important for autonomous defenses, where we must be able to precisely define the system's goals and restrictions to prevent the agent from performing undesired actions that are difficult to foresee at design time. Generally, formal models and specifications are needed for the autonomous agent, the system it operates in, and the adversary. In particular, these would define the objective of the agent (possibly a cost function or other quantitative objective), what actions it can take, and any constraints it must respect. Such specifications would enable not only proving the correctness and robustness of existing agents and defenses, but also synthesizing new defenses automatically. A key challenge in realizing this vision lies in identifying general classes of desirable properties for AI-based components, as their behavior often follows from trends in available data rather than human reasoning.

4.11 Acknowledgments

The organizers gratefully acknowledge the substantial contributions of ARL, academic, and other colleagues that participated in the ASPSM discussions. The organizers would particularly like to thank the breakout session leads—Engin Kirda, Srijan Kumar, Mike Reiter, and Tom Ristenpart—and Patrick McDaniel,

who summarized the work of the 2019 Networking and Information Technology Research and Development (NITRD) workshop on AI and cybersecurity.

The following individuals attended the workshop and contributed to the workshop discussions:

Academia:

- Ehab Al-Shaer, University of North Carolina (UNC) Charlotte
- Lujo Bauer, CMU
- Kathleen Carley, CMU
- Nicolas Christin, CMU
- Fei Fang, CMU
- Giulia Fanti, CMU
- Matt Fredrikson, CMU
- Daniel Fremont, University of California, Berkeley
- Roxana Geambasu, Columbia University
- Virgil Gligor, CMU
- Engin Kirda, Northeastern University
- Zico Kolter, CMU and Bosch Center for AI
- Srijan Kumar, Georgia Tech
- Jason Li, Siege Technologies
- Changliu Liu, CMU
- Patrick McDaniel, Pennsylvania State University
- Corina Pasareanu, CMU CyLab and NASA Ames
- André Platzer, CMU
- Mike Reiter, UNC Chapel Hill
- Tom Ristenpart, Cornell Tech
- Tuomas Sandholm, Strategy Robot, Inc. and CMU
- Yan Shoshitaishvili, Arizona State University
- Yevgeniy Vorobeychik, Washington University in St. Louis

- Michael Wellman, University of Michigan

Government:

- Jeff Alstott, Intelligence Advanced Research Projects Activity (IARPA)
- Tracy Bauer, ARL
- Kevin Chan, ARL
- Michael De Lucia, ARL
- Frank Geck, Army
- Kelly Ann Giraud, Army Futures Command (AFC)
- Charles Kamhoua, ARL
- Frederica Nelson, ARL
- Stephen Raio, ARL
- Brian Rivera, ARL
- Ted Senator, Defense Advanced Research Projects Agency (DARPA)
- Paul Yu, ARL

Organizers:

- Lujo Bauer and Matt Fredrikson, CMU
- Cliff Wang, Army Research Office (ARO)

5. Conclusion

Undoubtedly, rapid growth of new, powerful research results at the intersection of AI and other fields will continue. This report illustrates irresistible opportunities in applying AI in fields as diverse as SB, cyber defense, and the design of aeromechanical systems. Although the findings of the three ASPSM workshops described in the report are as different as their respective research fields, several themes emerge as fairly common. Two of them are particularly salient:

1) AI approaches must be applied to the right problems. The use of AI for problems that can be solved well by other means is not going to yield significant results. The greatest opportunities lie in new perspectives and novel paradigms in a given field, paradigms that truly require the strengths of AI.

2) Today's AI approaches rely primarily on large volumes of data. Although research on learning from small samples will continue, near-term opportunities require capabilities and facilities for the generation, collection, storage, sharing, and processing of massive amounts of data.

List of Symbols, Abbreviations, and Acronyms

AFC	Army Futures Command
AI	artificial intelligence
ANN	artificial neural network
ARL	Army Research Laboratory
ARO	Army Research Office
ASPSM	Army Science Planning and Strategy Meetings
CBC	Chemical Biological Center
CCDC	US Army Combat Capabilities Development Command
CMU	Carnegie Mellon University
DARPA	Defense Advanced Research Projects Agency
DOD	Department of Defense
EDD	Experimental Data Depot
ERP	Essential Research Program
GMO	genetically modified organism
IARPA	Intelligence Advanced Research Projects Activity
JBEI	Joint BioEnergy Institute
LIMS	Laboratory Information Management System
MDO	Multi-Domain Operations
ML	machine learning
NASA	National Aeronautics and Space Administration
NITRD	Networking and Information Technology Research and Development
OGA	other government agency
R&D	research and development
SB	synthetic biology
SC	Soldier Center
SME	subject-matter expert

UAS	unmanned aerial system
UNC	University of North Carolina

1 (PDF) DEFENSE TECHNICAL INFORMATION CTR DTIC OCA

1 (PDF) CCDC ARL FCDD RLD CL TECH LIB

14 (PDF) CCDC ARL FCDD RLC IA D SUMMERS-STAY FCDD RLD A KOTT FCDD RLS W BENARD FCDD RLS DE J J CARROLL C J CHIARA FCDD RLR EN S STANTON FCDD RLR PC D POREE FCDD RLV B GLAZ B H PIEKARSKI B M SADLER FCDD RLV A J L SHUMAKER FCDD RLV V F GARDEA FCDD RLW MB R WILDMAN FCDD RLW MC J SNYDER