National Cancer Institute
Frederick National Laboratory Advisory Committee
***Ad Hoc* National Cancer Institute–Department of Energy Collaborations Working Group Meeting**

National Institutes of Health (NIH) Campus
9000 Rockville Pike, Bethesda, MD
Conference Room 10
March 8–9, 2018

SUMMARY OF WORKING GROUP MEETING

PARTICIPANTS

Working Group Members

Dr. Piermaria Oddone (Chair)
Dr. Melissa Smith (Co-Chair)
Dr. Emily Greenspan (Executive Secretary)
Dr. Michael J. Becich
Dr. Martin Berzins (absent)
Dr. Kenneth Flurchick (absent)
Dr. David Galas (remote)
Dr. Elizabeth Jaffee
Dr. Warren Kibbe
Dr. Nilsa Ramirez-Milan
Dr. Robert Moser (remote)
Dr. Sarah Richardson
Dr. Kevin White
Dr. Cheryl Willman (remote)

Other Attendees

Dr. Tanmoy Bhattacharya, LANL
Mr. Jim Brase, LLNL
Dr. Michael Cooke, DOE
Dr. Yvonne Evrard, FNLCR
Dr. Ethan Dmitrovsky, FNLCR
Dr. James Doroshow, NCI
Dr. Paul Fearn, NCI
Dr. Amy Gryshuk, LLNL
Dr. Toby Hecht, NCI
Dr. Barbara Helland, DOE
Dr. Thuc Hoang, DOE
Dr. Tony Kerlavage, NCI
Dr. Dimitri Kusnezov, DOE
Dr. Joe Lake, ORNL
Dr. Carolyn Lauzon, DOE
Dr. Douglas Lowy, NCI
Dr. Christopher Miller, DOE
Dr. Dwight Nissley, FNLCR
Dr. Lynne Penberthy, NCI (remote)
Dr. Andrea Peterson, DOE
Dr. Eytan Ruppin, NCI

Dr. Eric Stahlberg, FNLCR
Dr. Rick Stevens, ANL
Dr. Frederick Streitz, LLNL
Dr. Georgia Tourassi, ORNL
Dr. Glendie Marcelin, SCG (Rapporteur)

THURSDAY, MARCH 8, 2018

WELCOME AND INTRODUCTIONS
Drs. Emily Greenspan and Carolyn Lauzon

Dr. Emily Greenspan welcomed the participants to the inaugural meeting of the Frederick National Laboratory Advisory Committee (FNLAC) National Cancer Institute (NCI)–Department of Energy (DOE) Collaborations Working Group (WG). The goal of the NCI-DOE Collaborations is to develop new and leverage existing exascale and high-performance computing (HPC) technologies developed by the DOE to advance NCI-supported cancer research. The mission benefit of the collaboration to both agencies is bi-directional. For the NCI, the development and use of exascale and high-performance computing (HPC) technologies can drive advances and innovation in cancer research. For the DOE, computing advances can be advanced through the application to real-world, complex biological systems. The responsibility of the WG includes participation with the biannual JDACS4C Governance Review Committee meetings and the scientific evaluation of projects that support or are relevant to the collaborations. The WG is tasked with providing technical recommendations on current pilot efforts under the Joint Design of Advanced Computing Solutions for Cancer (JDACS4C) collaboration and other relevant partnership opportunities. The WG also is responsible for advising the FNLAC on these recommendations. Dr. Greenspan added that two other committees have been established to advise the NCI-DOE collaboration: the JDACS4C Governance Review Committee (GRC) and the internal Trans-NCI-DOE Collaboration Program Team. The JDACS4C GRC has met two or three times yearly since the collaboration's inception. The GRC, which is comprised of senior federal leadership and program leads from NCI and DOE, is charged with providing guidance on strategy, policy, program and project management, budget and communication within federal government and executive branch. The goal of the trans-NCI-DOE Collaboration Program Team is to bring awareness of collaboration activities across NCI Divisions and Centers as well as create opportunities for collaboration with other existing NCI programs and NCI's extramural research community. It is comprised of NCI extramural program directors and intramural scientists who are involved in research programs that align with the goals of the NCI-DOE Collaboration. Dr. Carolyn Lauzon encouraged the WG to consider additional collaborative opportunities.

Dr. Greenspan reminded participants that the meeting is closed and confidential, and the discussions should not be disseminated broadly until a presentation has been made to the FNLAC or one of its subcommittees.

NCI-DOE COLLABORATIONS OVERVIEW
Drs. Amy Gryshuk and Eric Stahlberg

Drs. Amy Gryshuk and Eric Stahlberg presented an overview of the JDACS4C and described other current collaborative projects. Describing the purpose of the collaborations, Dr. Gryshuk indicated that the DOE is partnering with the NCI to develop exascale-ready tools, algorithms, and capabilities to meet NCI needs and advance DOE's exascale mission. The intended result of this partnership is to better understand cancer biology and develop more effective cancer therapies. The JDACS4C overall goal is to advance predictive oncology as part of the President's Precision Medicine Initiative and Cancer Moonshot[SM]. Dr. Gryshuk presented a timeline of activities that led to the creation of the JDACS4C.

Preliminary discussions between Dr. Douglas Lowy (then Acting Director of NCI) and Dr. Ernest Moniz (then Secretary of Energy) in December 2015 led to monthly meetings with staff from the NCI and the DOE. Through a Memorandum of Understanding between the NCI and the DOE, the formal collaboration for a period of five years was established on June 27, 2016.

Through JDACS4C, world-leading HPC experts at DOE national laboratories, including Argonne National Laboratory (ANL), Lawrence Livermore National Laboratory (LLNL), Los Alamos National Laboratory (LANL), and Oak Ridge National Laboratory (ORNL), and multiple cancer research scientists from NCI and Frederick National Laboratory for Cancer Research (FNLCR) work together on teams to extend the frontiers of precision oncology and exascale computing tools. JDACS4C is scoped as a 3-year program consisting of three pilot projects and two crosscutting DOE-driven programs: CANcer Distributed Learning Environment (CANDLE) and Uncertainty Quantification (UQ). Dr. Stahlberg described the three pilots and project deliverables:

- *Pilot 1: Predictive Modeling for Pre-Clinical Screening.* Framework for predictive models for preclinical screening. Integrate machine learning functionality into the Collaboration of Oak Ridge, Argonne, and Livermore (CORAL) systems.
- *Pilot 2: RAS Biology in Membranes.* Develop machine learning for dynamic validation of the RAS complex interaction model. Adaptive time and length scaling in dynamic multi-scale simulations.
- *Pilot 3: Population Information Integration, Analysis, and Modeling for Comprehensive Cancer Surveillance*. Modeling framework for predictive simulations of patient health trajectories. Integration of big data analytics with data-driven modeling and simulation for CORAL architectures.

Discussion

Dr. Stahlberg commented that the use of exascale computing is justifiable for the pilot projects because of the system complexity and the large data set that is expected to be generated from deep learning modeling.

Dr. Gryshuk summarized the Year 1 (June 2016–July 2017) progress for each pilot. Pilot 1 used existing NCI data sources to create model frameworks for machine learning. Pilot 2 ensured that experiments performed at FNLCR were driven by what was needed for computational models. Data generated from these models informed new experiments. Pilot 3 developed natural language processing (NLP) tools for automated identification in pathology reports. The optimization of millions of parameters per model was a crosscutting capability achieved. Dr. Stahlberg mentioned that the outcome of Year 1 includes evolving roles for deep learning and UQ, as well as new computational capabilities and data foundations.

In response to a question from Dr. Piermaria Oddone, Dr. Greenspan answered that the NCI Center for Biomedical Informatics and Information Technology serves as an "anchor" to centralize all the JDACS4C collaborations.

GOVERANCE REVIEW COMMITTEE AND WG
Drs. Lauzon, Greenspan and Lowy

Dr. Lauzon welcomed the members of the JDACS4C Governance Review Committee. The purpose of the session was for the WG to meet with the GRC, receive guidance, and ask questions of the NCI and DOE leadership. The JDACS4C GRC reviews the policy, strategy, and programs of partnership activities; provides advice on how to move pilot projects forward; and provides support for necessary gov-gov interactions and reach-back into both agencies.

Dr. Lowy acknowledged the efforts of Drs. Warren Kibbe and Dmitri Kusnezov to initiate the NCI-DOE Collaborations concept and expressed his gratitude for the DOE project execution and the bidirectional interactions between the DOE and the NCI.

In response to a question from Dr. Oddone, Dr. Lowy confirmed that one of the JDACS4C GRC's responsibilities is to make recommendations on the direction, scope, and approaches the pilot projects are taking and provide feedback on their value for continuation in the next fiscal years.

PRESENTATIONS

Pilot 1—Predictive Modeling for Pre-Clinical Screening

Drs. Rick Stevens and Yvonne Evrard

Dr. Rick Stevens updated participants on the goals, technical approaches, and status of the Pilot 1 project, Predictive Modeling for Pre-Clinical Screening. The Pilot 1 team is building an integrative, data-driven, deep learning model to predict responses to cancer therapeutic drugs, referred to as Combo, which has two initial versions, Mark 2 and Mark 3. Combo Mark 2 performs dose independent single drug and drug pair growth predictions based on training data from the ALMANAC study. It can predict on any datasets for which there is existing RNAseq data, including cell lines, organoids and patient derived xenografts (PDXs). Combo Mark 3 performs dose dependent single drug and drug pair predictions based on training data from the ALMANAC study. It can also predict on any datasets for which there is existing RNAseq data. The approach, a Drug Pair Synergy and Uncertainly Landscape map, computationally predicts responses across multiple tumor types using different drug combinations. The map combines the prediction of drug synergy and uncertainty. *In vitro* experiments are used to validate the plot results that show synergy, which could represent therapeutically useful drug combinations. The expected deliverables from Pilot 1 include training and inference with UQ to inform experimental design and developing a pathway for integration with other modeling approaches.

Drs. Stevens and Yvonne Evrard described *in vitro* and *in vivo* modeling of computational predictions. The Pilot 1 predictions were made using standard-of-care drugs in all three *in vitro* models: organoids, PDXs, and cell lines. Highlighting the diversity of the *in vitro* models, the three-dimensionality of organoids may simulate a more representative model for diffusion of drugs in human tumors. The heterogeneity of PDXs may mirror that of tumors. In contrast, cell lines are less heterogeneous and as they are transformed, they grow more effectively than PDXs and organoids. As part of the NCI-sponsored trial Molecular Profiling–Based Assignment of Cancer Therapy (MPACT), primary human tumor biopsy tissues were also analyzed for responses to various drug combinations. Near term next steps include using Combo to accurately predict *in vivo* combination drug results based on *in vitro* modeling.

Dr. Stevens remarked on the plan to use relevant data sets in a training database. For the model, the primary mode of molecular coding is via RNA sequencing. An important model feature is unsupervised machine learning used to determine how to normalize data filtering to achieve tight clustering between cell lines, patient tumors, and PDXs. Highlighting the capabilities of the model, tumors can be characterized by expression via microRNAs, single-nucleotide polymorphisms, or proteomics. The model can be "trained" to output growth curves (e.g., growth inhibition fraction).

Dr. Stevens reported on the planned activities for Pilot 1. The focus is to improve machine learning technology to enhance Combo's accuracy, performance, and "explainability." He added that CANDLE, the crosscutting project to build the deep learning infrastructure for cancer, is making improvements (i.e., data parallelism; model parallelism) that will affect all pilots. Recently, a novel capability has been incorporated into the model: two methods for UQ (bootstrapping and dropout UQ). A next step is to

develop a hybrid model that incorporates mechanistic modeling. Combo may be used to prioritize experiments and optimize preclinical screening of novel drug combinations. The new model called Uno Mark 1 (aka Combo Mark 4) includes the features of Combo Mark 3 but also includes the ability to train with data from multiple experiments including ALMANAC, GDSC, CCLE, gCSI, etc. Uno Mark 2 (aka Combo Mark 5) is under development and will include multi-task learning and better support for transfer learning. An exascale inference run is planned for spring 2018 using Summit, an ORNL high-performance supercomputer. Dr. Stevens added that Combo results are validated experimentally by commercial companies contracted to screen cell lines. This experimental data will then be fed back into the model.

Discussion

Dr. Elizabeth Jaffee inquired whether one can determine the right amount of samples that will generate sufficient data and produce accurate predictions. Dr. Stevens replied that good models are 60 - 80% predictive. The model learns (mean absolute error prediction <1%) and creates predictions, but we do not yet know its generalizability from cell lines to PDXs, and ultimately from these model systems to predictions in primary human tumors and in humans in vivo.

Dr. David Galas asked Dr. Stevens to speculate on the reasons for the observed lack of predictive power in the model's results. Dr. Stevens guessed that instances when the model made inaccurate predictions could be due to predicting outside of the training distribution. Deep-learning models normally have redundancy.

In response to a question from Dr. Michael Becich regarding CORAL and Summit computing systems, Dr. Stevens mentioned that CORAL is a collaboration between ORNL, ANL, and LLNL to process procurements and acquire diverse HPC systems. CORAL includes the pre-exascale systems Summit at ORNL and Sierra at LLNL which are currently being installed, and the planned A21 exascale computer at ANL in 2021. CORAL-2 is the planned upgrade to these systems. All of the CORAL systems are being designed to fully support deep learning as well as data analytics and complex simulations.

Dr. Kevin White requested clarification about the planned experimental cycle. Dr. Stevens reiterated that Combo is validated experimentally by the high-throughput screening of cell lines. The data are processed, then incorporated into the training; iterations are then run to determine how the data affect the model's predictive ability. Dr. James Doroshow added that FNLCR runs 20 complex spheroid cell lines with 20 combinations each month.

Dr. White questioned how well synchronized the predictions and experiments planned in the Patient-Derived Model (PDM) program are. Dr. Stevens indicated that once the model is trained, the turnaround time to make predictions is fast. Tumor characterization is the rate-limiting step for this process. Dr. Doroshow added that receiving model predictions is necessary for the experiments to commence; however, new experiments take 5 months of preparation.

Dr. Stevens mentioned several outreach activities for Pilot 1: The Pilot 1 team delivered more than 20 presentations to external organizations, and Drs. Stevens and Robert Grossman taught a course in machine learning at the University of Chicago. Dr. Stevens noted the general lack of advanced computing capability at the university level.

In reply to an inquiry from Dr. Oddone, Dr. Stevens estimated that for ANL, $4.5 million has been allocated for Pilot 1 and CANDLE. The funding sources are such diverse entities as the Exascale Computing Project and the NCI.

Dr. Elizabeth Jaffee wondered about the relevance of the model's input data, given that the field of precision oncology is focused on biological pathway and targeted therapeutics. Dr. Stevens replied that the model uses investigational drugs that have an established mode of action and target. It is possible to use information on the target to improve the model; however, this hinders the model's ability to generalize.

In reply to an inquiry from Dr. Kibbe, Drs. Stevens, Evrard, and Doroshow specified that the "cancer type" in the model refers to histological type. Dr. Stevens speculated that bringing histological data into the model is an alternative modality for characterizing tumors.

Pilot 2—RAS Biology in Membranes
Drs. Dwight Nissley and Fred Streitz

Drs. Dwight Nissley and Fred Streitz summarized the goals of the Pilot 2 project, RAS Biology in Membranes. Dr. Streitz presented an overview of RAS biology. Mutations in RAS are seen in approximately 30 percent of all human cancers; thus RAS has been considered an important therapeutic target for which drugs have not been successfully developed. The RAS signaling pathway is crucial to understanding its oncogenic properties. The activation pathway for RAS is critically dependent on its movement to and localization in the plasma membrane. Membrane-associated RAS binds to the RAF effector protein; thus, this protein complex could be a target for cancer therapeutics. The goal of Pilot 2 is to use molecular dynamic simulations to model this complex and understand the molecular mechanisms by which the complex forms at the membrane. The NCI RAS Initiative Program performs biophysical analysis of the complex to achieve this goal. Experimental unknowns include the order of RAS and RAF protein engagement at the membrane, the structure of membrane-bound RAS, and the various dynamic states of RAS.

Dr. Streitz described the approach for Pilot 2. The plan involves RAS-activation experiments, predictive simulation, adaptive sampling of molecular dynamics simulation codes, and machine learning–guided dynamic validation. The project is being executed across several laboratories: LANL, ANL, LLNL, ORNL, and FNLCR.

The project aims of Pilot 2 are—

- Aim 1: Develop spatial hierarchical multiscale modeling.
  - Capture the entire plasma membrane topology (lipid bilayer, water, and protein) and particle degrees of freedom.
- Aim 2: Understand the activation of the extended RAS complex.
  - Extend the findings from Aim 1 to the CORAL architecture.
- Aim 3: Create a machine learning–enabled dynamic validation approach to high-fidelity simulation.
  - Use machine learning as a framework to help guide simulation campaigns of exploration to help develop candidate therapeutics.

To implement these aims, Pilot 2 employs a multiscale approach that couples a diverse set of models, such as the Martini coarse-grained force field. A hyper-coarse protein macro and phase-field scale model has been used to computationally construct the cell-protein environment. The observed conditions that are most favorable and localized to the membrane are then fed back into the membrane model (feedback loop). These approaches have enabled the dynamic following of the RAS protein at a micro scale. An interesting feature of the model is that it learns the environment and corrects itself as the simulation is run. Demonstrating the biological relevance of this approach, a bilayer cell membrane spanning a micron in size (μm) and 100 RAS proteins will be simulated. As the first of its kind, this simulation approach will be submitted to the Association for Computing Machinery Gordon Bell Prize competition.

Dr. Nissley discussed the results of testable hypotheses that were developed from these initial simulations. A change in the dynamics of KRAS (a RAS isoform) and RAF on the membrane of live cells was tested by mutating the hypervariable region (HVR) of RAS. The mutations altered membrane potential; specific mutations in HVR caused RAS repositioning further away from the membrane in a charge-dependent manner. Nuclear magnetic resonance spectroscopy was used to further test the biological relevance of the simulation. By incorporating RAF kinase, membrane association of isolated RAF-cysteine-rich domains was observed in the simulations.

Dr. Streitz highlighted the outreach activities of Pilot 2. Since 2017, several journal publications and news releases have been published. Drs. Streitz and Nissley also have presented this work at conferences.

Discussion

Dr. David Galas wondered what effect pH (ions) and water changes have on membrane dynamics. Dr. Streitz replied that global variations in pH require changing the parameters of the free-energy state for the lipid membrane. Dr. Galas asked whether ions may mediate dimerization of RAS and RAF. Dr. Nissley cautioned against referring to the RAS-RAF interaction as dimerization. Dr. Nissley added that changes in ion concentrations may impact membrane dynamics and protein interactions, but the co-localization of RAS and RAF is driven mainly by the HVR-lipid interaction.

Dr. Kibbe asked how to predict when RAS would serve as a functional switch based on the lipid membrane composition. Dr. Nissley said that the lipid microenvironment is important for clustering RAS/RAF, but RAF localized to the membrane surface is sufficient for RAS activation. Dr. Streitz announced that Pilot 2 soon will investigate lipid membrane behavior, and the simulation capability is expected to be complete 4 weeks from March 8, 2018.

Dr. Kibbe wondered what types of simulation outcomes permit the development of new testable questions. Dr. Streitz responded that the behavior of RAS in the absence of RAF on the membrane will help elucidate whether RAS-RAF dimerizes and if it does what is the composition of the membrane.

Dr. White recommended adding small molecules (e.g., inhibitors) to the model; Dr. Streitz replied that this is possible computationally but requires prior understanding of the molecule going through the functional switch. Dr. Nissley added that ongoing experiments are using proteins that bind to the HVR and hinder RAS localization.

In response to a question from Dr. Sarah Richardson, Dr. Streitz affirmed the importance of gauging RAS affinities in different environments. Dr. Nissley added that performing structural analysis is imperative.

Dr. Oddone commented about experimental limitations; Dr. Streitz replied that the force field used in the model is an estimation. Addressing more mechanistic questions will require increased computing capability, but the framework to answer such questions exists.

Dr. Robert Moser asked about the dimensionality of the membrane environments. Dr. Streitz responded that it is in the μm range.

Pilot 3—Population Information Integration, Analysis, and Modeling for Comprehensive Cancer Surveillance
Drs. Georgia Tourassi and Paul Fearn

Drs. Georgia Tourassi and Joe Lake presented an overview of the Pilot 3 project, Population Information Integration, Analysis, and Modeling for Comprehensive Cancer Surveillance. Pilot 3 leverages HPC and advanced machine learning to support comprehensive, scalable, and cost-effective cancer surveillance. By using the NCI Surveillance, Epidemiology, and End Results (SEER) cancer database, patient data collected geographically by cancer registries are more robust. The aims of the project include (1) developing deep NLP for information capture; (2) crafting novel data analytic techniques for patient information integration; and (3) creating data-driven integrated modeling and simulation for precision oncology.

Highlighting recent DOE and NCI progress, the Pilot 3 team partnered with Information Management Services (IMS) to develop an Application Programming Interface (API) to test across the registries. As of December 2017, several machine learning algorithms were developed; the multitask Convolutional Neural Network (MT-CNN) algorithm was deployed to SEER via IMS. UQ approaches for Louisiana registry data were received and tested in an API. The NCI plans to develop an NLP infrastructure to annotate distance recurrence and biomarker information from pathology reports. As a long-term initiative, the team plans for improved integration of NLP and the formation of informatics research teams that are affiliated with SEER registries.

Dr. Tourassi elaborated on the specific approaches and challenges of Pilot 3, noting that the NLP is rule-based and requires intense domain expert involvement. The manual effort for this rule-based approach is not easily scalable. Conventional machine learning is scalable but requires intense feature engineering. Deep learning also is scalable, provided that sufficient computing power and data are available. The Louisiana registry data sets were imbalanced: 85 percent of patient cases represented 20 cancer sites and 60 histology codes. To address the lack of scalability, two new approaches were attempted to implement MT-CNN: hard parameter sharing and cross-stitch networks. Experiments were performed on Titan and Summit supercomputers at Oak Ridge Leadership Computing Facility. Dr. Tourassi noted the crosscutting activities with CANDLE. Data parallel training of MT-CNN was set up using Louisiana data. Softmax function and Dropout were two methods used to perform UQ in the neural networks. The Louisiana data sets are being entered into a graph framework to enable large-scale queries. This will help identify patients who are eligible for clinical trials.

Another goal is to complete a graph analytics method for registry data, which involves implementing a semantic web-based approach to registry data. The framework is in place to receive and analyze new data sets. Since December 2017, the Pilot 3 team has made presentations at two conferences and published a peer-reviewed paper.

Discussion

Drs. Kibbe and Becich recommended developing an automated pipeline for data across all registries. Dr. Becich speculated that open data sharing of outputted data (i.e., histology data) could be valuable for promoting translational cancer research. Dr. Tourassi replied that the DOE possesses the infrastructure for such an activity.

Dr. Becich wondered how many index patients are represented in all four registries. Dr. Lynne Penberthy replied that 120,000 cases are added per year, totaling, as of now, millions of cases. She shared that an important feature of the system is the incorporation of UQ to enhance data quality and accuracy.

Dr. Tourassi mentioned that a future direction is to implement UQ and develop comprehensive data sets on patient treatments, biomarkers, and outcomes at the population level from participating SEER registries. Dr. Lake mentioned that an iterative feedback loop will be implemented.

In reply to a question from Dr. Becich, Dr. Penberthy described the goals for distributing information to the SEER community. One goal is to automate the five variables (site, histology, laterality, grade, and behavior) with 70 percent of all pathology reports, which will provide real-time case accession. Another goal is the development of automated hypotheses. Dr. Penberthy announced that the Pilot 3 team is receiving digital images for pilot testing for a biorepository project (Year 3).

CANDLE
Dr. Stevens

Dr. Stevens presented an overview of CANDLE and described several project milestones. He acknowledged the many collaborators and principal investigators across various Institutes. The initiative was created upon recognizing the challenge of achieving deep-learning infrastructures across all three pilots. CANDLE uses the DOE's most capable supercomputers to answer the most challenging deep-learning problems in cancer research. It is a supported application under the Exascale Project with multiple HPC vendors (i.e., IBM, Intel, and the Exascale Computing Project). The functional goals of CANDLE include (1) using deep learning to help others increase productivity; (2) supporting established deep-learning frameworks (e.g., Google) to run on DOE supercomputers; and (3) managing CANDLE training data. Project deliverables comprise open-source tutorials and documentation available through the GitHub repository.

Expounding on the success of CANDLE, Dr. Stevens described completed milestones. In August 2017, the project created a prototype deep neural network (DNN) for information extraction from clinical reports (Pilot 3). In October 2017, the first version of Combo in CANDLE was delivered (Pilot 1). For Pilot 2, CANDLE created a prototype DNN that performs unsupervised feature learning. The ability to detect lipid clusters in computational simulations is complete. A proposed milestone for the remainder of Fiscal Year 2018 is to integrate the interface of Livermore Big Artificial Neural Network into CANDLE. The CANDLE team also plans to release CANDLE version 1.0 by the end of April 2018.

Dr. Stevens reported on the 2017 accomplishments and outreach activities for CANDLE. The team won the HPCwire Readers' Choice Award. A deep-learning workshop to discuss computing and cancer applications is scheduled for May 2018.

Discussion

Dr. Becich questioned how the CANDLE resources are allocated to the general public. Dr. Stevens responded that the use of DOE machines is merit based and requires a request.

Dr. Smith wondered about the feasibility of deploying CANDLE information on Amazon's computers. Dr. Stevens replied that the team has experts working on robust deployment engineering for CANDLE.

Dr. Becich asked what Institutes expressed interest in running CANDLE on their supercomputers. Dr. Stevens answered that many crosscutting activities involve various groups. The team is interested in integrating with other groups, provided that such efforts align with the goals of CANDLE.

UQ
Dr. Tanmoy Bhattacharya

Dr. Tanmoy Bhattacharya provided a conceptual overview of UQ and acknowledged the project members across the different laboratories (LANL, ORNL, ANL, and LLNL). UQ is the process to measure uncertainty in predictions and is central to all scientific results and pilot projects. For UQ, one must consider both interpolation and extrapolation errors. The UQ team is particularly interested in analyzing extrapolation errors. Generalization errors arise from overfitting; therefore, there is a large capability for the system to learn. To address generalization errors, the team is employing a dropout technique to assess the robustness of a given model. Upon completion of this assessment, decisions are made about whether to perform a more intensive analysis for output that produces uncertainty. The task of determining with certainty whether a model prediction is "good" or "bad" is called certainty distillation.

Dr. Bhattacharya discussed current UQ efforts related to Pilot 2. The RAS project sought to identify the dynamic movement of RAS protein in a two-dimensional (2-D) space. Part of the UQ assessment was to determine the certainty that the full motion of the protein is captured by the simulations. Dimension reduction analysis was performed to identify protein states in the 2-D space and to study escape dynamics. Other UQ efforts are to perform coarse graining analysis in the RAS simulation model and to implement UQ into the CANDLE framework.
Discussion

Drs. Becich and Kibbe recommended improving the messaging about the UQ project to broader communities to better convey the initiative's importance. Dr. Kibbe suggested that when explaining UQ to the cancer research community, the project team should provide examples of an inability to train the system when there are insufficient data. Dr. Bhattacharya responded that explaining UQ would not require a lot of examples if the illustrations are sufficiently distinct from each other. He asserted that the goal of UQ is to quantify correctness.

In response to a question from Dr. Moser, Dr. Bhattacharya replied that model form uncertainty is being addressed.

Accelerating Therapeutics for Opportunities in Medicine (ATOM)
Drs. Gryshuk and Stahlberg

Drs. Gryshuk and Stahlberg detailed efforts made by the ATOM team and introduced project leadership. Showcasing that ATOM is a public and private partnership, Dr. Gryshuk indicated that the consortium partners include Lawrence Livermore National Laboratory (LLNL); the University of California, San Francisco (UCSF); GlaxoSmithKline (GSK); and FNLCR. In January 2017, the DOE, GSK, and NCI signed a Memorandum of Agreement to establish ATOM, which UCSF later joined. ATOM's vision is to transform drug discovery from a slow, sequential, high-failure process into a rapid, integrated, patient-centric model. The mission is to accelerate the development of more effective cancer therapies for patients. To accomplish this, it is important to develop an integrated and concurrent precompetitive platform consisting of HPC, preclinical data sets (mainly from GSK), and emerging biotech capabilities. Dr. Gryshuk described the current ATOM moonshot goals, which involve accurate computational data prediction and showcasing *in silico* and *in vitro* models to reduce reliance on animal and cell-line models. Dr. Stahlberg outlined the data-driven drug discovery process for the ATOM workflow: first identifying patient-specific data and samples, relevant predictions, and gaps, and then determining the appropriate mechanistic experiments. Mr. Jim Brase emphasized important principles for ATOM, including taking a systematic approach to using "dark data" (unused data) from pharmaceutical companies to develop better learning-based models for predictive drug design.

Regarding the specific research and development strategy, Dr. Gryshuk explained that the plan is to execute a matrix approach to integrate experiments and computation. Integrated project teams are tasked with bringing biological, chemical, mathematical, and computational approaches to implement Quarter 1 deliverables. Dr. Stahlberg announced that the consortium is seeking additional partnerships and leadership.

Discussion

Dr. Oddone inquired about the accessibility of the ATOM data and software to the broader community. Dr. Stahlberg said that there is a period of 1 year before these resources are available.

In response to another question from Dr. Oddone, Dr. Gryshuk shared that the four partners are contributing such resources as staffing and financial support.

Dr. White questioned the uniqueness of ATOM, given the existence of several similar partnerships and programs. Similarly, Dr. Greenspan asked how ATOM is different from such initiatives. Dr. Gryshuk responded that the consortium hired a business development representative to reach out to various entities to establish collaborations. Dr. Stahlberg added that ATOM differentiates from other programs because of the computing level, amount of data, and degree of team integration.

Dr. Greenspan clarified that because ATOM is pertinent to the overall NCI-DOE collaborations, it is relevant to the WG's overall evaluation process.

In response to a question from Dr. Becich, Mr. Brase asserted that ATOM will be successful because of the DOE's computing capability for mechanistic and data-driven modeling of dark data. Dr. Becich commended the ATOM team and recommended emphasizing the need to get data to end-users rather than focusing on hardware capability. Mr. Brase agreed.

Dr. White wondered whether there is ongoing data contribution to ATOM, which he considers important for project success; Mr. Brase responded that GSK is doing this. Dr. White suggested turning ATOM into a nonprofit initiative, rather than a federal one through NCI-DOE. Dr. Gryshuk responded that the ATOM team members have ongoing discussions regarding this issue.

FRIDAY, MARCH 9, 2018

WORKING GROUP FEEDBACK ON DAY 1

Dr. Greenspan

Dr. Greenspan asked the WG members for their feedback and comments in response to the technical presentations and discussions with the JDACS4C Governance Review Committee.

Pilot 1—Predictive Modeling for Pre-Clinical Screening

Dr. White indicated that the experimental feedback loop is not moving fast enough and this is hampering feeding data back into the model. He suggested removing the bottleneck of dataflow by creating a target date for completion (i.e., 3 months), performing cell line testing to determine the accuracy of the predictions, and comparing these predictions to existing PDX data. If the predictions are not comparable, then the cell line piece should be dropped. He speculated that existing NCI data sets might replace cell line data. He added that FNLCR is working on testing drug combinations, and some are working, but

right now we cannot predict ahead of time which ones will work. Further development of Combo and additional PDX datasets are therefore required. Dr. Gryshuk responded that discussions have been held among management, Drs. Evrard and Stevens, and members of the extramural community for Pilot 1 to receive additional fully-annotated datasets.

Dr. Galas proposed thinking about how the current approach will evolve over the long term and adapt to future technologies. The volume, quality and nature of molecular data available over the next five years will be fundamentally different from today. He suggested beginning to think about how to grapple with different kinds of data (e.g. detailed single-cell profiling). Dr. Cheryl Willman added that the FNLAC has been excited about Pilot 1, but wants to make sure the focus is on the most compelling problem that will be translatable. Whether the models, novel algorithms, and frameworks generated from using cell line data will be generalizable when applied to PDX and large human data sets, or samples derived from patients who have participated in NCI clinical trials (where response and resistance to specific therapies are known) remains to be determined. While it is understandable that the project started with cell lines, it will be important to move to these types of sample sets for modeling as soon as the project allows. She also recommended that the teams might consider calls for large data sets available for modeling (PDX, human, etc.) from the NCI extramural community. Dr. Kibbe clarified that when Pilot 1 initiated, the NCI PDX data did not exist.

Several WG members also suggested that the Pilot 1 team engage the extramural community. Dr. Kibbe commented that the NCI might create a funding opportunity announcement to facilitate engagement. He also advised the team to make the training components and data sets available to the informatics groups at NCI-supported cancer centers and the broader community. Dr. Willman proposed initiating a forum at the NCI Frederick Laboratory (similar to the highly successful annual meetings of the RAS collaborative that engage large numbers of extramural collaborations and have led to significant intra and extramural collaborations) where experts and informatics/computational teams at NCI Cancer Centers and other NCI-funded investigators and programs can converge for training, sharing of data sets, methods, algorithms, and novel computational approaches. The opportunity for various NCI investigators and NCI programs to meet, train, and collaborate with the Pilot 1 team would be highly desired. Similarly, this effort might lead to a collaborative consortium that would lead to extramural investigators testing and applying various algorithms and novel computational methods on their datasets. Another engagement approach is to develop a process where NCI extramural investigators might submit projects for modeling/collaboration with DOE computational science teams.

Dr. Oddone suggested creating a stepwise process of validation: Perform experiments in cell lines prior to PDXs. Concerning the disadvantage of using cell lines, Dr. Kibbe attributed the loss of heterogeneity and biology of the original tumor samples to the immortalization process. Dr. Galas suggested a workshop to identify problems that current computational approaches can address and working with other existing NCI programs (e.g. Human Tumor Atlas Network).

Dr. Becich mentioned that the Pilot 1 team needs a dedicated NCI or FNLCR scientist to work with it to move the project forward. Drs. Kibbe and Oddone added that there might not be sufficient funding and resources to accomplish this. Drs. Willman and White again highlighted the opportunity to engage the extramural research community in these efforts. Dr. Oddone proposed that each pilot project have a WG of individuals with more in-depth expertise. Dr. Kibbe stated that all three pilots need more management attention and focus from NCI. There is the opportunity to collaborate with computational and systems biologists.

Summarized Proposed Next Steps for Pilot 1

- Assess the bottleneck of dataflow to quicken the experimental feedback loop.
- Develop potential mechanisms to engage the extramural community to generate collaborations.

- Develop potential mechanisms to engage extramural community to gather data that will allow a focus on human samples and PDX models as opposed to cell lines.
- For future studies, consider primary human tissue-derived data sets that might be available or developing from patients entering early and late phase NCTN-sponsored clinical trials, which would be highly annotated and where response/resistance to specific agents was known.
- The working group will consider working with FNLCR to identify an expert to provide additional scientific oversight.
- Down the line, consider convening an *ad hoc* Working Group specifically for pilot 1 and pilot 3 with expertise in Deep Learning applications.
- NCI should consider additional scientific staff at FNLCR to better interface with and support the scientific projects within NIH that could take advantage of exascale computing.

Pilot 2—RAS Biology in Membranes

Drs. Oddone and Willman agreed with the use of HPC to simulate RAS signaling mechanisms, structure, and membrane association. Dr. Willman expressed tremendous enthusiasm for the science and progress from Pilot 2, but encouraged the team to not lose sight of the original goal of the RAS program and project – to therapeutically target RAS. Dr. Richardson was pleased with the integration of experimentalists and the incorporation of simulations and resource utilization.

Dr. Galas suggested identifying other protein targets that can be addressed using the current approaches.

Summarized Proposed Next Steps for Pilot 2
- Determine the feasibility of engaging the private sector for data acquisition, collaboration, and project acceleration.
- Identify potential mechanisms to access additional data sets.
- Down the line, consider convening an *ad hoc* Working Group specifically for pilot 2 with expertise in molecular dynamics simulations.
- Determine if there is a need for NCI (and DOE) to provide more management oversight to ensure that the best capabilities of each agency are brought to the collaboration.

Comments from Dr. Dimitri Kusnezov

Dr. Kusnezov commented on the initiatives and the overall mission of the DOE. The DOE is concerned with integrating cognitive functionality into HPC and developing UQ to evaluate the uncertainty in predictions made by AI. The DOE is formulating partnerships to accelerate its mission while addressing NCI needs. Dr. Kusnezov agreed with the WG advise to perform data validation before addressing more complex questions, emphasizing that this process must inform the next-generation computing technologies and provide value to the NCI. The DOE has the unique capability of predictive modeling, and needs to be pushing HPC technology in areas of deep learning, and artificial intelligence (AI). Dr. Kusnezov predicted that with advances in cognitive approaches, the definition of simulation will change. The DOE is looking to develop scalable and novel machine learning tools (i.e., AI) on their biggest system (CORAL). Dr. Kusnezov would like to see the integration of AI into the software and hardware that are needed for the validation of large and complex experimental data important to DOE's mission.

Dr. White surmised that project acceleration will require more data sets from the private sector.
Dr. Willman added that the WG can identify mechanisms to access these types of data sets; however, those who submit data should be allowed to participate in the scientific process. Concerning additional

data sets, Dr. Kusnezov announced that the DOE has breast and cervical cancer patient records from Norway and is creating a national database of suicide indices.

Pilot 3—Population Information Integration, Analysis, and Modeling for Comprehensive Cancer Surveillance

Dr. Becich commented that the approach for Pilot 3 is flawed and too narrowly focuses on pathology reports. The goals are strong, but the work performed thus far (NLP on a small number of pathology reports) seems underpowered to accomplish the goals. Nevertheless, the machine learning tools are innovative, and the framework is robust. The project should incorporate discovery science and leverage national initiatives (e.g., ACT Network and ITCR for collaborative efforts) for the purpose of data sharing, which will help with the training sets needed for deep learning in NLP. Regarding the partnership with IMS, Dr. Becich expressed concern that IMS' tech platform may not be able to support modern, composable, dockerized software. Dr. White suggested expanding the five variables and number of registries for information extraction. Dr. Becich agreed that it was unclear how innovative the five variables are and how they will help with NCI research goals of Cancer Moonshot and precision oncology

Dr. Willman added that while it is understandable that Pilot 3 focused initially on a limited number of SEER registries, there was great population diversity across the SEER sites and new sites were being added by NCI; thus it would be ideal to expand the project to more SEER sites as applicable and as funding allowed. Dr. Becich agreed that the pilot is behind in addressing some of the more challenging goals presented and needs to implement a timeline for receiving and analyzing datasets from other SEER registries.

It was also not clear to the reviewers whether exascale computing resources, which are the primary focus or rationale for the NCI/DOE collaboration, were needed to achieve the current goals of Pilot 3. One of the challenges with the focus on the NCI SEER registry data is that in most registries, the patient data is usually captured and reported up to 1-2 years after initial diagnosis. Thus, SEER registry data are not currently useful for "real-time clinical decision making" that would be impactful to patients. Dr. Becich mentioned that additional SEER data points, such as geocoding, should be incorporated. Dr. Greenspan commented that the WG should discuss the project scope of Pilot 3 with Dr. Penberthy.

Summarized Proposed Next Steps for Pilot 3

- Determine feasibility of incorporating different types of patient reports, including whole slide images and radiology reports.
- Determine how SEER data could be more relevant to precision medicine and consider stated project goals. While retrospective analyses of SEER data sets that captured more comprehensive information would be highly desirable, and may well generate novel NLP approaches and methods, it is not clear how the current project is relevant to real time clinical decision making.
- Down the line, consider convening an *ad hoc* Working Group for pilot 1 and pilot 3 with expertise in Deep Learning applications.
- Determine if there is a need for NCI (and DOE) to provide more management oversight to ensure that the best capabilities of each agency are brought to the collaboration.
- Develop potential mechanisms to engage extramural community to ensure that the best capabilities of each agency are brought to provide more data for deep learning applied to NLP and other laudable (but yet addressed) goals of this important pilot.

CANDLE, UQ, and ATOM

Concerning CANDLE, Dr. Becich remarked on the uniqueness of the open-source machine learning approach. He commended the CANDLE outreach efforts with the NCI Informatics Technology for Cancer Research (ITCR) program and through github. He suggested that CANDLE partner with the Informatics Technologies for Cancer Research (ITCR) program to create a federated, open source platform for machine learning, and also commented that his own lab would be happy to do a deeper technical dive. Dr. Stahlberg pointed out that CANDLE has also done outreach through the Frontiers of Predictive Oncology and Computing (FPOC) meetings.

Dr. Moser commented that he is pleased with the broad directives of the UQ project and its approach to model-form uncertainty. Dr. Moser suggested that the sources of uncertainty be specified and that the team address the issue of validation.

Dr. Oddone suggested that the UQ team could improve its messaging by developing a tutorial. Related to messaging, Dr. Richardson added that the team must convey the uniqueness of its approach.

Although the ATOM team has a business development representative, Dr. Greenspan remarked that the WG can provide additional advice on extramural engagement. She added that the ATOM team is considering developing partnerships with the various DOE national laboratories.

Dr. White expressed uncertainty about how ATOM exists in relation to the NCI-DOE collaborations. Dr. Galas requested clarification about how ATOM is managed and interacts with the NCI. Dr. Greenspan indicated that the WG can provide the FNLAC board with more insight into the management of ATOM.

Dr. Becich proposed that other commercial entities share their dark data with the ATOM team to ensure project success.

Summarized Proposed Next Steps for CANDLE, UQ, and ATOM

- CANDLE
  - Consider improving the application of CANDLE by partnering with NCI's ITCR program.
- UQ
  - Determine feasibility of specifying sources of uncertainty.
  - Determine feasibility of addressing validation in relation to UQ.
- ATOM
  - Consider the feasibility of establishing a nonprofit entity.
  - Consider expanding partnership to gain more relevant datasets.

Overview Discussion

Dr. Oddone solicited feedback concerning the NCI-DOE collaborations overall. Dr. White remarked that the collaborations are outstanding; however, certain projects require redirection and the closing of experimental loops.

Dr. Becich commended the initiators of the collaboration and suggested that similar partnerships should be started in the NCI and recommended the engagement of the extramural community seen in Pilot 2 guide the path forward for Pilots 1 and 3.

Dr. Galas commended the choice of topics for the pilots and noted that the biggest project benefits are the organization, outreach, interactions, and methodology.

Dr. Oddone suggested developing a "central hub" (i.e., FNLCR) with sufficient resources and knowledge to support the collaborations. Dr. Greenspan emphasized the importance of receiving knowledge from a wide variety of stakeholders, possibly though a trans-NCI team that might include extramural program directors and intramural scientists. Dr. Oddone proposed an approach to better link the NCI and the DOE: create postdoctoral fellowships in computing and cancer research.

Dr. Becich suggested adding imaging data sets, possibly from the Human Tumor Atlas Network (HTAN) or The Cancer Genome Atlas (TCGA), to open-source machine learning approaches. Dr. Nilsa Ramirez-Milan suggested retrospective imaging (reimaging) from archived TCGA slides.

Dr. White suggested developing a Pilot 4 project focused on radiological and histological image analysis. Dr. Richardson advised contacting the University of California, Berkley for this effort, because it has ongoing biological imaging initiatives.

OVERALL OBSERVATIONS

- Overall, the three pilots are moving well and forcefully in their domains
- The cross cutting elements CANDLE and Uncertainty Quantification are essential components to all projects and are also developing well
- Through CANDLE there exists the foundation for strong participation of the wider research community
- ATOM, possibly organized as a not-for-profit corporation, could achieve major improvements in the time to develop new therapeutics.

OVERALL NEXT STEPS

- Consider strengthening the hub at FNLCR to better connect the DOE efforts to the large number of NCI supported programs that could support and/or profit from the collaboration
- While our working group can provide a broad evaluation of the pilots and cross-cutting efforts, many of the projects, such as Deep Learning, Uncertainty Quantification or Multi-scale Molecular Dynamics are highly specialized. At some time in the future, consider organizing *ad hoc* Working Groups in these subjects.
- Start to plan for proposal driven scientific research using the tools and resources devoted to this program. For example, additional pilots or projects to answer scientific questions using all the machinery developed by the collaboration.

ASSIGNMENT OF WORKING GROUP DUTIES IN CONSIDERATION OF NEXT STEPS

Dr. Greenspan reviewed the names of the WG members assigned to review the proposed next steps. Members will review the meeting summary to help refine their feedback to FNLAC. To apprise the FNLAC of progress to date, an interim WG report will be presented at the FNLAC regular meeting on May 8, 2018. Dr. Greenspan announced that the next WG meeting is scheduled for July 2018.

The following WG members agreed to provide oversight of:

- Pilot 1: Drs. Jaffe, White, and Galas
- Pilot 2: Drs. Oddone and Richardson
- Pilot 3: Drs. Becich, Smith, and Willman
- CANDLE: Dr. Becich
- UQ: Dr. Moser
- ATOM: Open to all WG members

ADJOURNMENT

Dr. Greenspan thanked the participants for their contributions and adjourned the meeting.