

AI Reloaded: ACAT2017 Conference Summary

Pushpalatha C. Bhat

Fermi National Accelerator Laboratory (Fermilab), Batavia, IL 60510

pushpa@fnal.gov

Abstract. The first edition of the Advanced Computing and Analysis Techniques in Physics Research (ACAT) workshop series was held at Fermilab in the year 2000, which was in fact the 7th meeting of its parent the original AIHENP workshop series that started in 1990 in Lyon, France. ACAT 2017 was the 18th workshop in the series, which is growing stronger and increasingly successful. It covered the areas of computing in physics research, computations in theoretical physics, and data analysis algorithms and tools. In this summary article, I will give some historical perspective on the workshop series and provide highlights of this workshop. Since Artificial Intelligence (AI) was the featured theme, I will briefly discuss machine learning (ML) and comment on the latest developments in ML applications in physics research covered at this workshop.

1. Introduction

The 18th edition of the ACAT (formerly AIHENP) workshop returned to its roots and featured Artificial Intelligence with the theme “AI reloaded.” Coinciding with the first session of the conference, a spectacular celestial phenomenon, a total solar eclipse that came to be called “The Great American Eclipse”, swept across the United States from Northwest to the Southeast, bringing beautiful sights of the eclipse and an exciting atmosphere to the start of the conference.

2. History of the Workshop Series

The workshop series started as the *Artificial Intelligence in High Energy and Nuclear Physics* (AIHENP) workshop [1] held in March 1990 in Lyon, France. The series was renamed as *Advanced Computing and Analysis Techniques* (ACAT) [2] when it was held at Fermilab in October 2000. ACAT 2017 would be the 18th workshop in the series, growing ever stronger and successful. The meeting is held every 18 months, alternating between spring and late summer/autumn and rotating through various regions of the world – Asia, Europe and Americas. Among the pioneers of the early series were Denis Peret-Gallix, Rene Brun, Fred James, Bruce Denby, Slava Ilyin, Andrei Kataev, Christian Kiesling, Jos Vermaseren and Monique Werlen.

The “Raison d’être” of the conference series has been to bring together theory, experimental physics and computing communities to “bridge” the disciplines and communities. The general strategy for the program has been to have three parallel “tracks” of sessions at each meeting: Track 1 dealing with Computing Technology for Physics Research, Track 2 on topics of Data Analysis Algorithms and

Tools, and Track 3 on Computations in Theoretical Physics, Techniques and Methods. Occasionally, a fourth track has been added to the meeting, depending on the venue and relevance to the ongoing events. The Fermilab workshop, for example, had a fourth track on Very Large Scale Computing. That meeting brought many trailblazers and pioneers in the advanced computing field: Bjarne Stroustrup (father of C++), Ian Foster (father of Grid Computing), Stephen Wolfram (creator of Mathematica), Alex Szalay and Robert Ryne (Large Scale computing).

The ACAT workshops have also highlighted a major theme for each of the meetings. Just to note some examples: World Wide Web was the theme for the 1992 meeting where Tim Berners-Lee gave a plenary talk, Parallel Distributed Computing was the theme for the 1996 meeting, Large Scale Simulations in 2000, Quantum Computing in 2007, Cloud Computing in 2011. The meetings organizers have always tried to sense major paradigm shifts that are coming our way and highlighted them. The theme for ACAT 2017 has been, very aptly, “AI Reloaded”.

3. This Meeting

This meeting amply met and surpassed the expectations on many fronts. As mentioned earlier, the major theme for the meeting was chosen to be “AI Reloaded”, given the recognition by the high energy physics (HEP) community of the enormous positive impact the machine learning (ML) methods have already had in our field’s major discoveries and studies, and the enormous potential of the ML/AI methods in the future [3,4].

The meeting drew quite a diverse gathering of about 200 participants. As has been the practice, the meeting was organized as a combination of common plenary sessions and three parallel sessions each following one of the three tracks of the meeting. Twenty-four invited plenary talks and eighty-six parallel session talks were given. About seventy posters were presented to a very keen audience in two successful poster sessions.

Another important feature at each ACAT meeting has been a number of round table discussions that are organized on topics of high relevance for the field and the times. There were four round table discussions on: Machine Learning, Using Heterogeneous Resources for HEP Computing, Analytical vs. Numerical Methods for NNLO+ Calculations for the Large Hadron Collider (LHC), and Diversity.

4. Highlights from the Three Parallel Tracks

4.1. Computing Technology for Physics Research

The main topic areas covered in talks in this track included (a) Grids and Cloud-federation - the continuing transition to virtual, standardized, commodity-based computing infrastructure (b) Heterogeneous Architectures, particularly on the use of Graphics Processing Units (GPUs) in HEP experiment triggers and promising applications of Field-Programmable Gate Arrays (FPGAs) in lattice QCD computing; (c) Online Systems and Triggers; (d) Modernizing tools for Software builds; (e) Containers, and their rising importance for software deployment and issues of their portability and scalability on supercomputers; (f) Machine Learning Tools for fast simulation; (g) Challenges in HEP Visualization in events at the LHC and HL-LHC (High-Luminosity LHC) and (h) Data Preservation Projects at the Tevatron, LHC and INFN.

4.2. Computations in Theoretical Physics

Several new developments in methods for symbolic calculations and loop integrals calculations were presented. Developments in event generators and their optimizations, and higher order radiative

corrections were discussed. Specific tools such as Loopedia, a new database for loop integrals, and Go_HEP, a new language for concurrent programming, open source codes such as pySecDec which provides tools for numerical evaluation of multiscale integrals, were discussed. It was clear that Machine Learning has been making its way into theoretical physics from the presentations on ML usage in PDFs to MC tools [5], Active Learning, and discussion on possible venues to use ML to accelerate lattice QCD.

4.3. Data Analysis: Algorithms and Tools

Several applications of ML algorithms such as Neural Networks (NN) and Boosted Decision Trees (BDT) were presented. In fact, many flavors of Deep Neural Networks (DNN) are being explored and exploited in HEP applications, especially by large experiments [6]. The use of ML is now ubiquitous in triggers and every aspect of event reconstruction and physics analysis – tracking, object ID, energy corrections, signal/background discrimination, simulations and so on. There were presentations on implementing them in end-to-end event reconstruction. Some examples of ML applications presented were (a) NN and BDT algorithms used in triggers in Belle-II, LHC, (b) event pile-up mitigation at the LHC, (c) event reconstruction in neutrino experiments, (d) use of DNN for online and offline tracking, (e) ML in theoretical physics as alluded to in the previous section. There were also talks on use of expert systems and cellular automata.

5. Machine Learning

Since this meeting's featured theme was AI, and that Machine Learning is becoming ubiquitous in HEP and elsewhere, I think it is worthwhile to muse over what it is, how it is still a burgeoning field and what to expect in terms of future applications and meetings.

In the context of physics data analysis, Machine Learning is the paradigm for automated learning of input-output mapping (or stimuli-response mapping) or “functional fitting” from data, using computer algorithms (or in hardware), requiring little *a priori* information about the mapping or function to be learned. A method that can approximate a continuous non-linear function to arbitrary accuracy is called a *universal approximator* and neural networks are an example of such a universal approximator. In the past nearly three decades since the AIHENP/ACAT workshop series started in 1990, when some of us who were the proponents and early adopters struggled to gain acceptance in the mainstream, these methods have gained gradual acceptance in HEP and other sciences, and have now become the “state of the art” in HEP data analyses. Some of the most important physics results in HEP, in the past over two decades, have come from the use of multivariate and/or machine learning methods. While in the 1990s, at the beginning of my talks, I had to proclaim that “we are riding the wave of the future”, that future is here and the use of MVA/ML methods is expected to only grow in relevance and importance.

5.1. Deep Learning and all that Jazz

Recent explosion in interest and applications of ML come from the availability of lots of computing power including the GPUs as well as algorithmic advances. Multi-scale feature learning with multiple hidden layers as in deep neural networks, where each higher-level layer of neurons learns increasingly higher-level features in the data, has revolutionized image recognition capabilities. Deep learning makes it possible to use raw data as inputs without having to derive “intelligent” discriminating variables to input to the classifier algorithm. One could use both raw inputs as well as derived variables in a DNN to achieve the best possible learning. When raw variables are inputs, it has been shown that the DNN performs better than a shallow NN. The so-called “Drop-out” algorithm has been

shown to help avoid over-fitting of the NNs. Unsupervised pre-training with stacked auto-encoders to build up the representation of the input space and then using supervised learning of the input-output mapping as a fine-tuning of optimization is shown to outperform supervised learning in a DNN where all hidden layers at once start out with randomized initial weights. A lot of processing power and lots of training data are needed indeed, and DNNs are generally implemented in GPUs and large amounts of training data can be generated with small perturbations in the initial training datasets.

Many flavors of deep learning are possible and have been implemented and used. Many of these were presented at this meeting. Apart from the deep feed-forward fully connected (FF FC) neural networks, Convolutional Neural Networks (CNN), Recurrent Neural Networks (RNN) and Generative Adversarial Networks (GAN) are some of the popular architectures being explored and applied in particle physics research. The CNNs have found novel applications in event classification in neutrino experiments. The applications of deep learning at the LHC have ranged from simple object-ID and signal/background discrimination and regression to fast simulation of events [7], pileup mitigation (in CMS and ATLAS) [8], and end-to-end reconstruction [9]. In some of the simpler applications of the first category, it is not clear that the resources applied to using the complicated architecture is justified. However, the current explorations will provide us with the experience and higher level of understanding leading to better heuristics.

Now that the Standard Model (SM) of particle physics is complete, in the sense that all the expected matter and force particles and the Higgs boson which provides for electroweak symmetry breaking have been discovered, new physics beyond the SM is vigorously being pursued by the HEP experimental community. Also very interesting for data analysis and theoretical interpretations are the ongoing efforts of the community to extract ML internal parameters to glean new insights on the process under study. We have the sophisticated, very powerful tools of ML in our hands. But the stakes are very high with regards to any claims of discovery of new physics. It is critical to establish robust methods of validation as well as proper handling of systematic uncertainties with the use of ML. As the vanguard of these powerful methods and tools, the ML community also has serious responsibilities in this regard.

6. Thoughts for Future Meetings

The ACAT2017 meeting covered an amazing array of diverse topics and applications and succeeded in bridging the theory, experiment and computing communities in high energy physics and their persuasions. Expansion into related areas for future meetings could include attracting contributions from astrophysics where advanced computing and ML applications have been in vogue for some time, and accelerator physics where ML is being used in controls applications and simulations. It is needless to say that it would be very fruitful to interact with statisticians and build strong connections. Since AI has been not only transforming scientific research but also our society and our daily lives, it would also be of great value to have invited plenary talks on other major advances in computing and AI/ML in fields outside the realm of physics. These issues as well as the interplay of High Performance Computing and Machine Learning will be on the agenda at the next ACAT (March 11-15, 2019) in Saas-Fee (Switzerland) [10].

7. Acknowledgments

I would like to thank the local organizing committee chaired by Gordon Watts and co-chair Shih-Chieh Hsu, the Program Committee chaired by Federico Carminati, and track coordinators N. Neufeld, M. Girone, G. Stewart, S. Gleyzer, G. Golovanov, A. Haas, T. Burnett, A. Freitas, S. Jones, F. Yuasa, and the International Advisory Committee. I thank D. Perret-Gallix and S. Gleyzer for reading this

paper and making useful suggestions. I also thank all the excellent speakers, presenters and participants for a very delightful and productive meeting.

My research is supported by Fermilab which is managed by Fermi Research Alliance, LLC (FRA), under the contract number DE-AC02-07CH11359 with the U.S. Department of Energy.

References

- [1] New Computing Techniques in Physics Research, Proc. 1st International Workshop on Software Engineering, Artificial Intelligence and Expert Systems in High Energy and Nuclear Physics, Lyon (March 19-24, 1990), ed. By D. Perret-Gallix and W.Wojcik, CNRS, Paris 1990.
- [2] Advanced Computing and Analysis Techniques in Physics Research, Proc. VII International Workshop, Fermilab, Batavia, IL, (Oct. 16-20, 2000), ed. By Pushpalatha C. Bhat and Matthias Kasemann, AIP Conference Proceedings Vol. 583, NY, USA 2001.
- [3] P.C. Bhat, "Multivariate Analysis Methods in Particle Physics," Annu. Rev. Nucl. Part. Sci., Vol. 61 (2011) 281-309.
- [4] <http://ipc.fnal.gov/computing/HEPMachineLearningCommunityWhitePaper.pdf>
- [5] S. Carrazza, "Machine Learning Challenges in Theoretical HEP," these proceedings.
- [6] B. Nachtmann, "Deep Learning usage by Large Experiments," and references therein, these proceedings
- [7] S. Vallecorsa, "Generative Models for fast simulation," these proceedings
- [8] P.T. Komiske, E.M. Metodiev, B. Nachman and M.D. Schwartz, "Learning to remove Pileup at the LHC with Jet Images," these proceedings.
- [9] M. Andrews, M. Paulini, S. Gleyzer, and B. Poczos, "End-to end Event Classification of High-Energy Physics Data," these proceedings.
- [10] <https://indico.cern.ch/event/708041/>