

PUBLIC SUBMISSION

As of: 9/16/21 10:53 AM Received: September 15, 2021 Status: Pending_Post Tracking No. ktl-10xi-7thj Comments Due: September 15, 2021 Submission Type: Web

Docket: NIST-2021-0004
Artificial Intelligence Risk Management Framework

Comment On: NIST-2021-0004-0001
Artificial Intelligence Risk Management Framework

Document: NIST-2021-0004-DRAFT-0082
Comment on FR Doc # 2021-16176

Submitter Information

Name: Juyang Weng
Address:
Okemos, MI, 48864
Email:
Phone:

General Comment

As an AI expert who initiated a direction called Autonomous Mental Development (AMD), I inform the US Government with the following three risks that the AI field is facing and two AI bottlenecks:

Risk 1: Shallow fitting of Data. The ImageNet Competition 2012 is laudable to have attracted investments from companies like Google into neural networks, like Convolutional Neural Networks (CNN) and Long Short Time Memory (LSTM), trained by a gradient-based method such as error backprop for a shallow fitting of “big data”. Such methods are both shallow and brute-force because data fitting is easy if a network has a huge number of parameters (see the nearest neighbor classifier in Att. 1). They do not generalize well if the networks are correctly tested (See Risk 2).

<http://www.cse.msu.edu/~weng/research/2021-06-28-Report-to-Nature-specific-PSUTS.pdf>

Risk 2: Protocol flaws. There is a rarely reported stage in report generation called Post-Selections: After many networks have been trained only report the luckiest network. Specifically, Post-Selections are technically flawed protocols because of one or more of the following flaws: (1) a lack of cross-validation for the hand-picked partition between the training set and test set, (2) a flawed use of test sets (see, e.g., PubPeer posts about LeCun et al. Nature 2015, Graves et al. Nature 2016, Li Fei-Fei et al. PAMI 2006, and Fig. 1 of Att. 1), (3) a lack of verification of the random seed luck for initial weights of a network, (4) a lack of verification of the hyper parameter luck of a network architecture, (5) the Post-Selection stage is typically not reported, (6) not reporting many less lucky networks, (7) a use of much human labor to hand-tune architecture parameters to fit a static “big data” set, (8) Some AI Competitions (like ImageNet) were improperly managed thus leaving time for “cracking” test sets and not caring how much human interventions, computational resources and manpower each team used.

Risk 3: “Big data” flaws. Pre-collected “big data” are wrong (e.g., violation of sensorimotor recurrence,

see Theorem 1 in Att. 1) and nonscalable (a lack of invariant rules, see Theorem 2 and observations 1 and 2 in Att. 1). Such machines do not know that they made mistakes nor how to recover from these mistakes. (E.g., for lane keeping, the human driver must periodically apply force to the steering wheel since the lane-keeping machines do not know they miss a plane and how to move back.) Thus, such machines are not trustable for even simple tasks like lane keeping. They still need a human to watch closely and to correct errors. Conscious learning is necessary to deal with these trustability and scalability issues, but conscious learning has not been published until Weng 2020 (See Att. 4) and could be overlooked for 20 more years like AI long overlooked the first Convolutional Networks (CNNs) for 3D (Cresceptron). Conscious learning is necessary to be trustable and scalable.

Bottleneck 1: Lack of Education for Conscious Learning. Universities have updated their curricula considerably, but they are still ill-suited for the required breadth and depth of Consciously Learning. The material in Weng, Natural and Artificial Intelligence (BMI Press) is required for AI experts to catch up with conscious learning. This is the first AI bottleneck.

Bottleneck 2: Lack of brain-scale real-time learning chips. Although brains learn in real time because of the sensorimotor recurrence requirement (see Risk 3), this necessary condition for trustable intelligence has been overlooked. Real-time learning requires brain-scale neural network chips, but chip manufacturers are not aware of this market. They do not know great pay-offs will come from a brain-scale real-time learning Developmental Network (DN). This is the 2nd AI bottleneck.

Attachments:

Att. 1: PostSelection-ArXiv-web.pdf (an archival journal manuscript of Att. 2 and 3, also from arXiv).

Att. 2: PSUTS-IJCNN2021rvsd-cite.pdf

Att. 3: Post-Selection-ICDL-2021rvsd-cite.pdf

Att. 4: ConsciousL-ICDL-2020rvsd-cite.pdf

Attachments

Att1-NNsubm-2021-09-09-web

Att2-PSUTS-IJCNN2021rvsd-cite

Att3-PostSelection-ICDL-2021rvsd-cite

Att4-ConsciousL-ICDL-2020rvsd-cite