

Why Deep Learning’s Performance Data Are Misleading

Juyang Weng^{*†}

^{*}Brain-Mind Institute

[§]GENISAMA LLC, Okemos, MI 48864 USA

Abstract—This is a theoretical paper, as a companion paper of the keynote talk at the same conference. In contrast to conscious learning, many projects in AI have employed deep learning many of which seem to give impressive performance data. This paper explains that such performance data are probably misleadingly inflated due to two possible misconducts: data deletion and test on training set. This paper clarifies what is data deletion in deep learning and what is test on training set in deep learning and why they are misconducts. A simple classification method is defined, called nearest neighbor with threshold (NNWT). A theorem is established that the NNWT method reaches a zero error on any validation set and any test set using Post-Selections, as long as the test set is in the possession of the author and both the amount of storage space and the time of training are finite but unbounded like with many deep learning methods. However, like many deep learning methods, the NNWT method has little generalization power. The evidence that misconducts actually took place in many deep learning projects is beyond the scope of this paper. Without a transparent account about freedom from Post-Selections, deep learning data are misleading.

I. INTRODUCTION

Since 2012, AI has attracted much attention from public and media. A large number of projects in AI have published [8], [1], including AlphaGo, AlphaGoZero, AlphaZero, AlphaFold, and IBM Debater. If the authors of these projects understand the principles in this report, they could benefit much for reducing the time and manpower to reach their target systems as well as improving the generalization powers of their target systems.

Post-Selections [14] mean first training multiple systems each starting with random weights (and searched hyper-parameters) and then selecting a trained system based on validation set, test set, or both.

This paper, based on the analysis of Post-Selections [14], raises two flaws that seem to widely exist in machine learning projects: Data deletion and test on training data. The latter is applicable when test data are in the possession of the programmer, which is the case with all published paper where the authors possess the test data.

An open-competition is different, such as Deep Blue versus Garry Kasparov Feb. 10, 1996 - May 11, 1997, AlphaGo versus Lee Sedol March 9, 2016 - March 15, 2016, and AlphaGo versus Ke Jie May 23, 2017 - May 27, 2017, because the test data arrives on the fly. This paper defines Post-Selections Using Test Set (PSUTS) On The Fly (OTF) conducted by humans behind the scene. The author does not claim that PSUTS OTF indeed took place during any of these

three events (although Kasparov did). The PSUTS OTF mode is for the academic community to be aware of, to be alert about, and to investigate in the future.

In the remainder of the paper, we will discuss four learning conditions in Sec. II from which we can see that we cannot just look at superficial “errors” without limiting resources. Sec. III discusses four types of mappings for a learner, which gives spaces on which we can discuss errors. Post-Selections are discussed in Sec IV. Section V provides concluding remarks.

II. THE FOUR LEARNING CONDITIONS

First, let us consider four learning conditions that any fair comparisons of AI methods should take into account.

Many AI methods were evaluated without considering how much computational resources are necessary for the development of a reported system. Thus, comparisons about the performance of the system have been tilted toward competitions about how much resources a group has at its disposal, regardless how many networks have been trained and discarded, and how much time the training takes.

Here we explicitly define the Four Learning Conditions for development of an AI system:

Definition 1 (The Four Learning Conditions): The Four Learning Conditions for developing an AI system are: (1) A body including sensors and effectors, (2) a set of restrictions of learning framework, including whether task-specific or task-nonspecific, batch learning or incremental learning; (3) a training experience and (4) a limited amount of computational resources including the number of hidden neurons.

Weng 2021 [14] discussed the conditions (2) to (4) without condition (1). It further discussed why any Big Data set violates what is called the sensorimotor recurrence principle, namely, any learning process that uses a static Big Data set is physically flawed.

III. FOUR DIFFERENT MAPPINGS

Traditionally, a neural network is meant to establish a mapping f from the space of input X to the space of class labels L ,

$$f : X \mapsto L \quad (1)$$

[4], [9]. X may contain a few time frames.

Many temporal problems, such as video analysis problems, speech recognition problems, and computer game-play problems, can include context labels in the input space, so as to

learn

$$f : X \times L \mapsto L. \quad (2)$$

where \times denotes the Cartesian product of sets.

A developmental approach deals with space and time in a unified fashion using a neural network such as Developmental Networks (DNs) whose experimental embodiments range from WWN-1 to WWN-9. The DNs went beyond vision problems to attack general AI problems including vision, audition, and natural language acquisition as emergent Turing machines [12]. DNs overcame the limitations of the framewise mapping in Eq. (2) by dealing with lifetime mapping without using any symbolic labels:

$$f : X(t-1) \times Z(t-1) \mapsto Z(t), t = 1, 2, \dots \quad (3)$$

where $X(t)$ and $Z(t)$ are the sensory input space and motor input-output space, respectively.

We need to consider two factors: (A) Space: Because X and Z are vector spaces of sensory images and muscle neurons, we need internal neuronal feature space Y to deal with sub-vectors in X and Y and their spatial hierarchical features. (B) Time: Furthermore, considering symbolic Markov models, we also need further to model how Y -to- Y connections enable something similar to higher and dynamic order of time in Markov models. With the two considerations (A) Space and (B) Time, the above lifetime mapping in Eq. (3) is extended to:

$$f : X(t-1) \times Y(t-1) \times Z(t-1) \mapsto Y(t) \times Z(t), t = 1, 2, \dots \quad (4)$$

in DN-2. It is worth noting that the Y space is inside a closed “skull” so it cannot be directly supervised. $Z(t-1)$ here is extremely important since it corresponds to the state of an emergent Turing machine.

In terms of performance evaluation, all the errors occurred during any time in Eq. (4) of each life is recorded and taken into account in the performance evaluation. This is in sharp contrast with Post-Selection.

IV. POST-SELECTIONS

Before we discuss Post-Selections, we need to discuss three types of errors.

A. Fitting, Validation and Test Errors

Given an available data set D , D is divided by a partition into three mutually disjoint sets, a fitting set F , a validation set V , and a test set T so that

$$D = F \cup V \cup T. \quad (5)$$

Two sets are disjoint if they do not share any elements. The validation set is possessed by the trainer, the test set should not be possessed by the trainer since the test should be conducted by an independent agency. Otherwise, V and T become equivalent.

Given any hyper-parameter vector \mathbf{a}_i (e.g., including receptive fields of neurons), it is unlikely that a single network initialized by a set of random weight vectors can result

in an acceptable error rate on the fitting set, called fitting error, that the error-backprop training intends to minimize locally. That is how the multiple sets of random weight hyper-parameter vectors come in. For k hyper-parameter vectors \mathbf{a}_i , $i = 1, 2, \dots, k$ and n sets of random initial weight vectors \mathbf{w}_j , the error back-prop training results in kn networks

$$\{N(\mathbf{a}_i, \mathbf{w}_j) \mid i = 1, 2, \dots, k, j = 1, 2, \dots, n\}.$$

Error-backprop locally and numerically minimizes the fitting error $f_{i,j}$ on the fitting set F .




Fig. 1. A 2D-terrain illustration for the global minima problems in a high-dimensional terrain (e.g., 200B-dimensional space of hyper parameters and weights). TM: Turing machine.

Fig. 1 gives a 2D illustration for the limitations of Post-Selection as well as the difference between the sensor-only mapping in Eq. (1) and the sensorimotor mapping in Eq. (3). Suppose X represents the space of sensory input and Z represents the space of motor input, but in general should be any initial pair $(\mathbf{a}_i, \mathbf{w}_j)$. Each location on the 2D plane of Eq. (1) corresponds to the initial random weights of a neural network (e.g., 200B-dimensional). Of course, each network has many neurons not just two input values but Fig. 1 can only schematically represent the initial weights of two values as a 2D-terrain illustration. The height of a curve at is the system fitting error $e_{i,j}$ of a particular trained network $N(\mathbf{a}_i, \mathbf{w}_j)$. We use the biological term “life” to indicate the entire process of a system’s learning.

The error-backprop learning method is a greedy method. Starting from any initial pair $(\mathbf{a}_i, \mathbf{w}_j)$ it steps along the direction that descends the quickest without knowing where the global minimum is in the 200B-dimensional space. In Fig. 1, we can see that $(\mathbf{a}_i, \mathbf{w}_j)$ leads to a local minimum a , b , or c if we use the sensory mapping in Eq. (1). The point a is the lowest point, but there is no guarantee to reach it, depending on where the network starts from in the 200B-dimensional space. The more networks have been trained, the more likely the luckiest network finds the global minimum. In Fig. 1, the luckiest network starts from the valley where a is located, but this is much harder for the 200B-dimensional space. Typically, the more networks a project has trained, the



Fig. 2. Two annotated windows for an object class labeled as “steel drum” for single object localization. Figure courtesy of [10].

more likely for the Post-Selection stage to find a network with a smaller e_{i^*,j^*} .

Graves et al. [5] seems to have mentioned that the number of trained systems is at least $n = 20$. Saggio et al. [11] reported that n is at least 10,000. Krizhevsky & Hinton [7] did not give n but seems to have mentioned 60 million parameters which probably means each w_i and each a_j combined to be of 60 million dimensional. Consider a small example: The number of tried value l for each hyper-parameter: $l = 3$, and the number of hyper-parameters $d = 10$, the total number of hyper-parameter vectors is $k = l^d = 3^{10} = 59049$. Letting $n = 20$, $kn = l^d n \approx 1\text{M}$ networks must be trained, a huge number that requires a lot of computational resources to do number crunching and a lot of manpower to manually tune the range of hyper-parameters!

Definition 2 (Distribution of fitting, validation and test errors):

The distributions of all kn trained networks’ fitting errors $\{f_{ij}\}$, validation errors $\{e_{ij}\}$, and test errors $\{e'_{ij}\}$, $i = 1, 2, \dots, k$, $j = 1, 2, \dots, n$ are random distributions depending on a specific data set D and its partition $D = F \cup V \cup T$. The difference between a validation error and a test error is that the former is computed from the same author using an author-possessed validation set V but the latter is computed by an independent agency using an author-unknown test set T .

We define a simple system that is easy to understand for our discussion to follow. Consider a highly specific task of recognizing patterns inside the annotated windows in Fig. 2. This is a simplified case of the three tasks—recognition (yes or no, learned patterns at varied locations and scales), detection (presence of, or not, learned patterns) and segmentation (of recognized patterns from input). These three tasks of natural cluttered scenes were dealt with by the first deep learning networks for 3D—Cresceptron [15]. Later data sets like ImageNet [10] contain many more image samples.

B. Post-Selection Illusion

To understand Post-Selection in Deep Learning (using CNN, LSTM, etc.), including supervised learning, reinforcement learning and adversary learning, let us consider a simple classifier that shares the same principle of Post-Selection but

does not have the distraction of details in more commonly used networks as well as learning modes.

Definition 3 (Nearest neighbor with threshold, NNWT):

Define a network that stores the entire fitting set F where each image in F may contain multiple annotated windows for matching (see Fig. 2). For each input image q , a scan subwindow x searches across the input image q for a range of locations and scales, normalizes the scale, and compares with each annotated window in the fitting set F . Suppose an input window x matches the nearest sample (annotated) window s in F . If the distance between x and s is not larger than a distance threshold d (a hyper-parameter), then the network outputs the associated label of the nearest sample s . Otherwise, the system outputs a label randomly and uniformly sampled from the label set L .

Let us give some formality of the Post-Selection with NNWT. Suppose that there are three data sets, fitting set F , verification set V , and test set T , a desired verification error $e_v \geq 0$, and a desired test error $e_t \geq 0$, run the following program $P(s)$ that starts from random seed s .

- 1) Store all data from the fitting set F .
- 2) For each receptive field (e.g., rectangular) \mathbf{q} inside of a cluttered image or a cluttered video $\mathbf{x} \in V$, or $\mathbf{x} \in T$, compute the distance $d_i = d(\mathbf{q}, \mathbf{x}_i)$ between receptive field \mathbf{q} and each annotated receptive field \mathbf{x}_i in F for all values of i (each sample in F have multiple annotated receptive field \mathbf{x}_i 's).
- 3) Produce the class label l_j of the nearest neighbor \mathbf{x}_j that produces the smallest distance: $j = \arg \min_i \{d_i\}$.
- 4) If $d_j \leq d(s)$ where $d(s)$ is a threshold depending on random seed s , output the label l_j of the nearest neighbor. Otherwise, output a label that is randomly and uniformly sampled from the label set L .

Post-Selection stage: If the measured verification error or the test error is larger than required, do the above procedure $P(s)$ for a new seed s . Repeat until the measured verification error and test error are both not larger than the required e_v and e_t , respectively.

In order words, the NNWT classifier with Post-Selections uses a lot of space and time resources for over-fitting. It has a perfect fitting error (zero) but it randomly guesses an output label if the distance is larger than the threshold d .

To understand why Post-Selections give misleading results, let us derive the following important theorem.

Theorem 1 (Post-Selection Illusion): Given any validation error rate $e_v \geq 0$ and test error rate $e_t \geq 0$, using Post-Selections the NNWT classifier using Post-Selections satisfies the e_v and e_t , if the author is in the possession of the test set and both the storage space and the time spent on the Post-Selections are finite but unbounded.

Proof: Because the number of seeds to be tried during the Post-Selection is unlimited, we can prove that there is a finite time at which a lucky seed $d(s)$ will produce the good enough verification error and test error. Although the waiting time is long, the time is finite because V and T are finite. Let us formally prove this. Suppose l is the number of labels in the

output set L and for the set of queries in V and T , there are k outputs that must be guessed. The probability for a single guess to be correct is $1/l$ due to the uniform sampling in L . The probability for k guesses to be all correct is $(1/l)^k = 1/l^k$ because guesses are independent. For an independently and randomly initialized network to guess less than k cases correct is $1 - 1/l^k$, with $0 < 1 - 1/l^k < 1$. The probability for as many as n independently and randomly initialized networks, all of which do not satisfy the e_v and e_t , is

$$p(n) = (1 - 1/l^k)^n \rightarrow 0, \text{ as } n \text{ approaches infinity}$$

because $0 < 1 - 1/l^k < 1$. Therefore, within finite time span, a process of trying incrementally more networks will get a lucky network that satisfies both the computable e_v and e_t . ■

As we can see from the proof, the smaller the threshold, the more guesses must be made and, thus, typically the longer time one needs to spend during Post-Selections.

Let us imagine that the threshold $d(s)$ gradually increases from zero toward infinity. The nearest neighbor classifier changes from a total-lottery scheme (when $d(s) = 0$) to a traditional nearest neighbor classifier (without label guesses) when all query inputs are beyond the threshold.

Theorem 1 has established that Post-Selections can even produce a perfect classifier that gives a zero validation error and a zero test error! Yes, while the test sets are in the possession of authors, the authors could show any superficially impressive validation error rates and test error rates (including zeros!) because they used Post-Selections without a limit on resources (to store all data sets and to search for a lucky network).

The above theorem means any comparisons without an explicit limit on storage and time spent on Post-Selections are meaningless, like ImageNet and many other competitions. It is of course time consuming for a program to search for a network whose guessed labels are good enough. But such a lucky network will eventually come within a finite time!

Corollary 1 (Misleading AI papers): If a paper trains more than one system and the author is in the possession of test set, the performance data from the paper are misleading if the paper does not report the number of systems trained in Post-Selections, the amount of computational resources (e.g., the amount of storage, the number of computations), the amount of waiting time, along with the fitting errors, the validation errors, and the test errors of *all* trained systems. The generalization power of the reported system is still unknown.

Proof: From Theorem 1, if Post-Selections are allowed, a NNWT can satisfy any nonzero validation error and any nonzero testing error using Post-Selections, since the training set, validation set and the test set are all in the possession of the authors. The generalization power of all trained systems is unknown, so is the reported luckiest system. ■

This corollary is a scientific basis for the author to raise a violation of protocols and a lack of transparency about the Post-Selection stage in almost all machine learning papers appeared in *Nature*, *Science*, *Communication of ACM* [2] and other publication venues since around 2015. Since all the

fitting sets, validation sets and test sets are in the possession of the authors, many papers have claimed misleading results using a flawed protocol.

Following the terminology of [3], we use the term “Test on Training Set” below.

Theorem 2 (Test on Training Set): Post Selections that used test set amounts to tests on training set even though trained networks did not “see” the training set during their error backprop training.

Proof: Originally, the term “Test on Training Set” means the training stage involves the test set. Although all trained networks did not “see” the test set, the Post-Selection is part of the training stage that produces the reported network and its error. Since the Post-Selection uses the test set to pick up the luckiest network among all trained networks using their errors on the test set, all trained networks see the test set during the post-selection stage. ■

A typical neural network architecture has a set of hyper-parameters represented by a vector \mathbf{a} , where each component corresponds a scalar parameter, such as convolution kernel sizes and stride values at each level of a deep hierarchy, the neuronal learning rate, and the neuronal learning momentum value, etc. Let k be a finite number of grid points along which such hyper-parameter vectors need to be tried, $A = \{\mathbf{a}_i \mid i = 1, 2, \dots, k\}$. Let’s give more detail to the above example. Suppose there are 10 scalar parameters in each vector \mathbf{a}_i . For each scalar parameter x of the 10 hyper parameters, we need to validate the sensitivity of the system error to x . With uncertainty of x , we estimate its initial value as the mean \bar{x} , positively perturbed estimate $\bar{x} + \sigma_x$ (σ is the estimated standard deviation of x), and negatively perturbed estimate $\bar{x} - \sigma_x$. If each scalar hyper parameter has three values to try in this way, there are a total of $k = 3^{10} = 59049$ hyper-parameter vectors to try, a very large number. For example, in NNWT, the threshold \bar{d} can be estimated by the average of nearest distance between a sample in V and the nearest neighbor in F and the σ_d be estimated by the standard deviation of these nearest distances.

C. Test Data Available in Batch

Ideally, test sets should not be in the possession of authors, e.g., during a blind test. However, this is often not true since the authors may use publically available data sets that include test sets.

Definition 4 (Post selection in batch): A human programmer trains multiple systems using the fitting set F . After these systems have been trained, the experimenter post-selects a system by searching, manually or assisted by computers, among trained systems based on the batch validation set V (or the batch test set T). This is called Post-Selection in batch—selection of one network from multiple trained and verified (or tested) networks.

Definition 5 (Data deletion in post-selection): Data deletion is a misconduct during which an author deletes the performance data of some bad-performing networks

during post-selection, not reported in the corresponding project report.

Data deletions are violations of the well-known statistical protocol of cross-validation. In a minimally acceptable form, the distribution of the performances of all trained networks should be reported by the worst error, the average error and the best error among all the trained networks. Only reporting the best error amounts to misconduct known as data deletion, deleting the performance data of all networks that the author does not like.

Obviously, a post-selection wastes (deletes) all trained systems except the selected one. As we can predict [14], a system from the post-selection tends to have a weaker generalization power than the reported luckiest error indicates, as illustrated in Fig. 1.

A Post-Selection in batch can use the validation set V or the test set T . However, if both sets are in the possession of the human programmer, the difference between V and T almost totally vanishes under Post-Selections.

Definition 6 (PSUTS and PSUVS): A Machine PSUTS is defined as follows: If the test set T is available to the author, suppose the test error of $N(\mathbf{a}_i, \mathbf{w}_j)$ is $e_{i,j}$ on the test set T , find the luckiest network $N(\mathbf{a}_{i^*}, \mathbf{w}_{j^*})$ so that it reaches the error of the luckiest hyper-parameters and the luckiest initial weights from Post-Selection Using Test Set (PSUTS):

$$e_{i^*,j^*} = \min_{1 \leq i \leq k} \min_{1 \leq j \leq n} e_{i,j} \quad (6)$$

and report only the performance e_{i^*,j^*} but not the performances of other remaining $kn - 1$ trained neural networks. PSUVS, V for validation, is similarly defined.

Set T is like set V since it is available. Similarly, a human PSUTS is a procedure wherein a human selects a system from multiple trained systems according to $\{e_{i,j}\}$ using also human visual inspection of internal representations of the system and their test errors.

D. Cross-Validation

The above PSUTS is an absence of cross-validation [6]. Originally, the cross-validation is meant to mitigate an unfair luck in a partition of the dataset D into a fitting set F and a test set T (empty validation set). For example, an unfair luck is such that every point in the test set T is well surrounded by points in the fitting set F . But such a luck is hardly true in reality.

To reduce the bias of such a luck, an n -fold cross-validation protocol [3] is suggested.

E. Types of Lucks in a Neural Network

In a neural network, there are at least three types of lacks:

Type-1 order lacks: The luck in a partition P_i into a fitting set F_i and a test set T_i from a data set D resulting in test error e_i , $i = 1, 2, \dots, n$.

Type-2 weights lacks: As discussed in [14], weights specify the role assignment for all the neurons in the neural network. A random seed value determines the initialization of a pseudo-random number generator, which gives initial weights \mathbf{w}_i

for a neural network $N(\mathbf{w}_i)$, resulting in a test error e_i , $i = 1, 2, \dots, n$, after training of these n networks and testing on T .

Type-3 hyper-parameter lacks: The initial hyper-parameter vector \mathbf{a}_j of the neural network gives an error e_j , $j = 1, 2, \dots, k$. Because such a luck of \mathbf{a}_j might not capture the internal rules of the fitting set F_j , this paper argues that a statistical validation of the reported error estimate should be performed and the distribution of $\{e_j\}$ be reported.

In contract, a batch-trained neural network typically uses a Post-Selection to pick the luckiest network without cross-validation for either of the above three types of lacks, e.g., in ImageNet Contest [10]. Namely, errors occurred during batch training of the network before the network is finalized and how long the training takes are not reported. Many researchers have claimed error-backprop “works” without providing much-needed three types of validations. This seems not true since [16] shows a huge difference between the luckiest CNN with error-backprop and the optimal DN.

We also need to be aware of another protocol flaw: Random seeds and hyper parameters are all coupled. Under such a coupling, Type 2 validation seems unnecessary with $n = 1$ but the search of the luckiest weights is embedded into the search for the luckiest hyper-parameter vector. where each hyper parameter vector uses a different seed.

Since a PSUTS procedure picks the best system based on the errors on the test set (like a validation set), the resulting system does not do well on new test sets because doing well on a validation set does not guarantee doing well on an open test set. See Fig. 1. Typically, due to a very large number of samples, availability of test sets and unavailability of open test sets in a properly managed contest, Post-Selections cause the reported error to be smaller than an open test error rate. (However, in Table 2 of [7], the test error rate is smaller than the validation error for 7CNNs, causing a reasonable suspicion that PSUTS could be used in addition to PSUVS.)

F. The Luckiest Network from a Validation Set: Data Deletion

Many people may ask: Are there any technical flaws in at least PSUVS, since it does not use the test sets? Any post-selection is technically flawed and results in misleading results, including both PSUVS and PSUTS, because of the data deletion misconduct.

In general, Type-1 cross-validation is to filter out lacks in data partition that a typical user does not have during a deployment of the method. Namely, it is a severe technical and protocol flaw in reporting only the luckiest network, regardless the post-selection uses validation sets or test sets. At least the average error over Post-Selections must be reported.

This conclusion has a great impact on evolutionary methods that often report only the luckiest network, instead of those of all networks in a population. Namely, the performances of all individual networks in an evolutionary generation should be reported. Furthermore, a reasonably disjoint test set must be used to evaluate the generalization of the luckiest network.

If the test set T is in the possession of the author, which seems to be true for almost all so-called “deep learning” publications other than open competitions, we define machine PSUTS as a Post-Selection process where test set is used, in addition to the validation set.

Some researchers have claimed that the test data was “unseen” by trained systems when they were tested, but the network selector has seen their performances when they see the test set. In other words, the Post-Selection stage belongs to the training phase—training the Post-Selector.

Weng [14] discussed also human PSUTS where it is a human that does Post-Selection. The subsection below discusses human PSUTS on the fly, not in batch.

Weng [14] also discussed why error-backprop needs machine PSUTS and human PSUTS.

G. Open Test Data Arriving on the Fly

Definition 7 (PSUTS On The Fly): During an open competition, the participating $m > 1$ luckiest networks meet new open test data that arrive on the fly. One or more humans conduct an additional round of Post-Selections on the fly during the competition using human PSUTS from the m machine outputs for some actions. Here, all human selected actions across the entire game do not even have to be consistently from a single network.

For example, $\mathbf{x}(t)$ is the current board configuration in a Go game, treated as a game state at player time t . Human experts (behind the scene or on the scene), based on the output actions from the m luckiest networks, manually select a network’s output as the next action $\mathbf{z}(t+1)$. Of course, this is not a fair competition, because the so-called “human player side” has only a single human who is not allowed to use any computer but the so-called “machine player side” has one or more humans who have assistance from computers.

Weng [13] discussed how the DN algorithm trains only a single network that is always optimal in the sense of maximum likelihood, conditioned on the Four Learning Conditions. In this sense, the DN is free from local minima.

V. CONCLUSIONS

Post-Selections in so called “deep learning” involve “data deletion” and “test on training data”. Many experts consider such practices as misconducts. Performance data from so-called “deep learning” are misleading without explicit exclusion of such flaws in their experimental protocol. Such flawed protocols are tempting to those published papers where the test sets are in the possession of the authors and also to open-competitions where human experts are not explicitly disallowed to interact with the “machine player” on the fly. Limited by scope of paper, the author does not claim here that such misconducts indeed took place.

REFERENCES

- [1] M. G. Bellemare, S. Candido, J. G. P. S. Castro, M. C. Machado, S. Moitra, S. S. Ponda, and Z. Wang. Autonomous navigation of stratospheric balloons using reinforcement learning. *Nature*, 588(7836):77–82, 2020.
- [2] Y. Bengio, Y. LeCun, and G. Hinton. Deep learning for AI. *Communications of ACM*, 64(7):58–65, 2021.
- [3] R. O. Duda, P. E. Hart, and D. G. Stork. *Pattern Classification*. Wiley, New York, 2nd edition, 2001.
- [4] K. I. Funahashi. On the approximate realization of continuous mappings by neural networks. *Neural Network*, 2(2):183–192, March 1989.
- [5] A. Graves, G. Wayne, M. Reynolds, T. Harley, I. Danihelka, A. Grabska-Barwinska, S. G. Colmenarejo, E. Grefenstette, T. Ramalho, J. Agapiou, A. P. Badia, K. M. Hermann, Y. Zwols, G. Ostrovski, A. C., H. King, C. Summerfield, P. Blunsom, K. Kavukcuoglu, and D. Hassabis. Hybrid computing using a neural network with dynamic external memory. *Nature*, 538:471–476, 2016.
- [6] A. K. Jain and R. C. Dubes. *Algorithms for Clustering Data*. Prentice-Hall, New Jersey, 1988.
- [7] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. *Communications of the ACM*, 60(6):84–90, 2017.
- [8] Y. LeCun, L. Bengio, and G. Hinton. Deep learning. *Nature*, 521:436–444, 2015.
- [9] T. Poggio and F. Girosi. Networks for approximation and learning. *Proceedings of The IEEE*, 78(9):1481–1497, 1990.
- [10] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, A. C. Berg, and L. Fei-Fei. ImageNet large scale visual recognition challenge. *Int’l Journal of Computer Vision*, 115:211–252, 2015.
- [11] V. Saggio, B. E. Asenbeck, A. Hamann, T. Stromberg, P. Schiansky, V. Dunjko, N. Friis, N. C. Harris, M. Hochberg, D. Englund, S. Wolk, H. J. Briegel, and P. Walther. Experimental quantum speed-up in reinforcement learning agents. *Nature*, 591(7849):229–233, March 11 2021.
- [12] J. Weng. Brain as an emergent finite automaton: A theory and three theorems. *Int’l Journal of Intelligence Science*, 5(2):112–131, 2015.
- [13] J. Weng. A developmental method that computes optimal networks without post-selections. In *Proc. IEEE Int’l Conference on Development and Learning*, pages 1–6, Beijing, China, August 23-26 2021.
- [14] J. Weng. On post selections using test sets (PSUTS) in AI. In *Proc. Int’l Joint Conference on Neural Networks*, pages 1–8, Shenzhen, China, July 18-22 2021.
- [15] J. Weng, N. Ahuja, and T. S. Huang. Learning recognition and segmentation using the Cresceptron. *Int’l Journal of Computer Vision*, 25(2):109–143, Nov. 1997.
- [16] Z. Zheng and J. Weng. Mobile device based outdoor navigation with on-line learning neural network: a comparison with convolutional neural network. In *Proc. 7th Workshop on Computer Vision in Vehicle Technology (CVVT 2016) at CVPR 2016*, pages 11–18, Las Vegas, June 269 2016.