A comparison of distributed machine learning methods for the support of "Many Labs" collaborations in computational modelling of decision making

Lili Zhang^{1,2} Himanshu Vashisht¹, Andrey Totev¹, Nam Trinh¹, Tomas Ward^{1,2}

1 School of Computing, Dublin City University, Dublin, Ireland

2 Insight SFI Research Centre for Data Analytics

¶Current Address: McNulty Building, Dublin City University, Glasnevin Campus, Dublin 9

* lili.zhang27@mail.dcu.ie

Abstract

Deep learning models, especially RNN models, are potentially powerful tools for representing the complex learning processes and decision-making strategies used by humans. Such neural network models make fewer assumptions about the underlying mechanisms thus providing experimental flexibility in terms of applicability. However this comes at the cost of requiring a larger number of tunable parameters requiring significantly more training and representative data for effective learning. This presents practical challenges given that most computational modelling experiments involve relatively small numbers of subjects, which while adequate for conventional modelling using low dimensional parameter spaces, leads to sub-optimal model training when adopting deeper neural network approaches. Laboratory collaboration is a natural way of increasing data availability however, data sharing barriers among laboratories as necessitated by data protection regulations encourage us to seek alternative methods for collaborative data science. Distributed learning, especially federated learning, which supports the preservation of data privacy, is a promising method for addressing this issue. To verify the reliability and feasibility of applying federated learning to train neural networks models used in the characterisation of human decision making, we conducted experiments on a real-world, many-labs data pool including experimentally significant data-sets from ten independent studies. The performance of single models that were trained on single laboratory data-sets was poor, especially those with small numbers of subjects. This unsurprising finding supports the need for larger and more diverse data-sets to train more generalised and reliable models. To that end we evaluated four collaborative approaches for comparison purposes. The first approach represents conventional centralized data sharing (CL-based) and is the optimal approach but requires complete sharing of data which we wish to avoid. The results however establish a benchmark for the other three distributed approaches; federated learning (FL-based), incremental learning (IL-based), and cyclic incremental learning (CIL-based). We evaluate these approaches in terms of prediction accuracy and capacity to characterise human decision-making strategies in the context of the computational modelling experiments considered here. The results demonstrate that the FL-based model achieves performance most comparable to that of a centralized data sharing approach. This demonstrate that federated learning has value in scaling data science methods to data collected in computational modelling contexts in circumstances where data sharing is not convenient, practical or permissible.

Author summary

Introduction

One of the most important tasks in cognitive psychology and behavioural science is to understand and explain how people think, infer and process information. This is generally achieved through observing their behaviour under experimental settings. In the field of human decision making, where we evaluate how people make decisions based on learning from experience, we can either analyse the process by which information is itself processed [1, 2] or investigate how the expectation values of options are learned and updated through experience [3, 4]. In such approaches, we use cognitive models that specify a set of assumptions about the underlying learning processes used by the subjects, thus generally requiring manual engineering with an iterative process to examine the consistency between the hypothesis and the empirical data [5]. The approach may fail if subjects adopt a completely different strategy during decision making. Deep Learning (DL) models, especially Recurrent Neural Network (RNN) models, present an alternative approach to the characterisation of human decision making behavior. They can automatically capture behavioural trends exhibited by subjects without strong assumptions about the mathematical structure of the underlying process through taking advantage of their data-driven design and higher capacity for representing complex computational processes [5–7].

Unfortunately, this does not come for free since modern DL models involve the learning of a large number of parameters which in turn requires large amounts of diverse data for effective performance. For example, in a recent medical study [8], it was found that DL models overfit when trained on small and biased institutional data and generalized poorly on data from institutions whose data were not seen during training. Although the data used in that study comprised of medical images rather than behavioural data, it is a good example and a warning that DL models need to be carefully trained. Efforts must be made to minimize the introduction of confounding factors associated with experimental biases which can dominate the training process and obscure the relationship of predictions to the underlying targeted pathology. Such models may perform with good accuracy when testing against held-out data from the same experiment, but may not generalize well to the same experiments conducted in other studies. A natural way to increase both data size and diversity is through collaborative learning, in which multiple laboratories cooperate and contribute data to train a global model together. It is an effective approach that has the potential to solve the limitation in most cognitive studies where the number of subjects is adequate for the modelling approach originally intended but then becomes inadequate when the same data is used to train new models involving a larger number of parameters. Furthermore, one of the significant challenges in computational psychiatry [9], i.e. translating advances in understanding the cognitive biases of people with mental illness into improvements in clinical practice, will also benefit from the pattern of multi-experiment collaboration, given that the scale of current cognitive studies is not sufficient for developing robust models that satisfy clinical needs.

One of the most conventional options for conducting multi-experimental analysis requires different laboratories to share subjects' data to a centralized location for model training (i.e. centralized training). However, in practice, data sharing is not easy, especially when involving larger number of laboratories from different legal jurisdictions, due to privacy, ethical, and data regulation barriers [10]. Consequently, information coming from various populations worldwide remains distributed in isolation across multiple laboratories, highlighting the need to seek alternative approaches. The solution for addressing this problem is distributed model training, in other words, sharing model

parameters instead of data, as the model itself does not contain any individually-identifiable information. One approach for distributed training is unparallel, or as we will refer to it here - incremental distributed learning (IL) [11]. With this approach each laboratory trains the model and passes the learned model parameters to the next laboratory for training on its data, until all have trained the model once. Another approach is an extension of incremental distributed learning, called cyclical incremental distributed learning (CIL). This involves the repetition of the incremental learning process multiple times, i.e. fixing the number of training iterations at each laboratory and cycling repeatedly through the laboratories. Chang et al. [12] explored these two distributed training approaches using medical imaging data and found the cyclical approach performed comparably to centralized training, suggesting that sharing data may not always be necessary to build deep learning models for patient imaging data. The incremental distributed learning approaches, however, have been criticized for introducing the problem of 'catastrophic forgetting' [13], where the trained model strongly favors the data it has most recently seen. Nevertheless the repeated cycles and limited iterations per institution performed in cyclic incremental learning enable it to make gradual progress, despite the forgetting issue, resulting in better models than the first distributed training approach.

Federated learning (FL) can be considered another distributed machine learning paradigm seeking to directly address the problem of data governance and privacy. It was first introduced by Google AI [14] to allow mobile devices collaboratively learn machine learning models without sharing data from the devices. It was applied very successfully in training Google's autocomplete keyboard application [15]. Institutions or laboratories can also be viewed as 'devices' in the context of federated learning and therefore FL has emerged as a promising strategy in scenarios where, for example, hospitals operate under strict privacy practices and may face legal, administrative, or ethical constraints that require data to remain local. Different from the unparallel approaches, FL is a data parallel training process, in which multiple collaborators train a DL model simultaneously (each on their own data in parallel) and then send their model parameters to a central sever where these are aggregated into a global model. The central sever then sends the global model to all collaborators for further training. Each iteration of this process, i.e. parallel training, parameter update, and distribution of global parameters, is referred to as a federated round. Micah et al. [11] compared federated learning and incremental learning approaches on imaging datasets and unsurprisingly, the basic incremental learning performed the worst compared to federated learning and cyclic incremental learning. While cyclic incremental learning may seem a simpler alternative, given that implementation of federated learning depends on a set of key challenges [16], e.g., communication efficiency, which is outside the scope of this article, it required additional validation steps at the end of each cycle, which are basically as complex as the synchronization logic of FL, to achieve comparable results to federated learning paradigm. Critically, cyclical incremental learning was less stable than federated learning, resulting in an inferior alternative.

The degree to which the laboratory data-sets used for distributed learning are independent and identically distributed (IID) can have a significant influence on the learning performance compared to centralized learning [17] as deep learning models rely on Stochastic Gradient Decent (SGD) algorithm and the IID samples can ensure that the stochastic gradient is an unbiased estimate of the full gradient [18]. However, in practice, it is unrealistic to assume that the local data of each laboratory is always IID, which is also a statistical challenge for applying federated learning. It was found in [19] that the accuracy of a federated learning model trained for image classification was reduced by up to 55% depending on how much institutional bias or degree of non-IID they introduced while dividing a single dataset into hypothetical institutions. The

institutions and institutional biases in their paper were created artificially by defining hypothetical clients and assigning various number of samples to the clients. However, the results of applying the distributed approach on artificially partitioned hypothetical institutional datasets may fail to account for how real-world clients biases impact distributed learning paradigm, a point that has been argued in [17]. Consequently this also motivated this paper in which we experiment with real-world institutional datasets.

In this study, we examined the reliability and feasibility of applying federated learning to RNN models for predicting human behaviour and characterising their decision making processes, importantly, using a diverse data pool coming from 10 independent real studies. Subjects in all studies were healthy participants and completed a computerized version of the Iowa Gambling Task (IGT), which is one of the most widely used tasks measuring human decision making under uncertainty in an experimental context. Although previous papers have demonstrated distributed learning applications in various scenarios, e.g. institutional distributed medical image classification [11, 17], none of them has ever applied distributed learning on trial-to-trial decision making data originating from different laboratories. We first trained single laboratory models for each laboratory in the data pool and then evaluated each of these models against held-out testing datasets from each laboratory defined prior to model training. This experiment is used to demonstrate the need for numerous and diverse data for training a robust deep learning model. Second, we illustrated the benefits of multi-laboratory collaboration and the superiority of federated learning compared to incremental learning and cyclic incremental learning. Thus, one centralized model and several distributed models were trained for this purpose. The centralized model is expected to be the best collaborative model as it is the only model that would not lose any information during collaboration. Although all the distributed learning models are expected to sacrifice accuracy and robustness in return for data privacy, the federated learning model was considered to be the most competitive paradigm to train RNN models for predicting human decision actions in the circumstance where data sharing is not possible and thus distributed learning has to be employed.

Materials and methods

The Iowa Gambling Task

The Iowa Gambling Task was originally developed to study the decision making deficits of ventromedial prefrontal cortex patients [20]. Over two decades, it has been one of the most widely used neuropsychological paradigms for simulating complex and experience-based decision-making [21]. Participants in the IGT are initially given €2000 virtual money and presented with four decks of cards labelled A, B, C, and D. Each card in these decks can generate rewards, and sometimes cause losses. Participants have to choose one card from these four decks consecutively, until the task terminates automatically after a fixed number of trials have been reached. In each trial, feedback on rewards and losses for their choice and the running tally over all trials so far are provided to the participants, but no information is given regarding how many trials they will play and how many trials they have completed during the task. Participants are instructed that they can choose cards from any deck and they can switch decks at any time. They are also told to make as much money as possible by minimising losses.

The IGT dataset

The data pool ($N = 617$) we used in this study derives from 10 studies assessing performance of healthy participants (i.e. without any known neurological impairments)

on the Iowa Gambling Task. It involves a broad range of healthy populations aging from 10 to 88 with various education backgrounds and social status. Participants completed a computerized version of the IGT consisting of 95, 100 or 150 trials. All included studies used (a variant of) the traditional IGT payoff scheme or the payoff scheme introduced by Bechara & Damasio [22]. Since the RNN model we used for predicting subjects' actions required us to have inputs with the same size, the trial-by-trial decision sequences of all subjects were truncated according to the length of the smallest sequence, i.e. 95 trials. In other words, for those subjects who completed 100 or 150 trials of the IGT, only the first 95 trials were used for training or testing the model. Data from different studies were collected in different environments. The number of subjects in each laboratory, which we will refer to *laboratory1-10*, are 15, 162, 19, 40, 70, 25, 153, 35, 57, and 41 respectively. The proportion of females varied approximately from 0.32 to 0.74.

This dataset is reused here because, 1) first of all, participants in all studies were healthy subjects and they all completed the same decision making task - IGT, which makes it feasible to train a collaborative model for all laboratories; 2) the subjects from ten studies were coming from different backgrounds with various ages and female proportions, thus, conducting multi-laboratory analysis is a natural operation to increase data size and diversity and improve the performance of the model; 3) it created a perfect real-world scenario where the local data of each lab is potentially non-IID and biased, thus it is suitable for testing out the reliability and effectiveness of distributed learning paradigms, especially federated learning, that are sensitive to the data distribution. Each laboratory can be seen as a 'device' in the context of federated learning and the data assignments for the clients in FL can be matched with the real-world data distributions, such that all subjects from the same study are assigned to the same client; 4) last but not least, we can more thoroughly investigate how healthy subjects learn and behave on the IGT using deep learning models, which as stated previously, are promising in the area of decision making modelling because of their higher capability of learning from data in a less model-constrained manner.

The Architecture of the Recurrent Neural Network model

The RNN model we used is composed of a GRU layer and an output softmax layer with four nodes corresponding to the four choices for the different decks in the IGT. The inputs to the GRU layer is the previous choice made by the subject along with the rewards and losses they received after making that choice. The softmax output are the probabilities of choosing each deck option and sum to 1. The architecture of the model is shown in 1.

It is known that the basic RNNs are inefficient in solving problems that require the learning of long-term temporal dependencies because of the gradient vanishing problem [23]. If a sequence is long enough, which is the case of our choice sequences that include at least 95 trials, they'll fail to convey information (the deck preferences) from earlier trials to later ones. LSTM [24] and GRU [25] networks were created as the solution to this short-term memory issue. They have internal mechanisms called gates that can regulate the flow of information. The gates can learn which data in a sequence is important to keep or throw away. By doing this, it can pass relevant information down the long chain of sequences to make predictions. In most cases, the two approaches yield comparable performance. It is often the case that the tuning of hyperparameters may be more important than choosing the networks. We chose the GRU here in this paper since it has fewer parameters and can be trained faster. We believe this to be a suitable choice given that we have a relatively small dataset with moderately long sequences.

The GRU layer is composed of a set of hidden units (N_u). Each unit is associated with a unit output denoted by h_t^k for cell k at time t . Let's define $h_t = [h_t^1, h_t^2, \dots, h_t^{N_u}]^T$

Fig 1. Architecture of the RNN model. The model has a GRU layer (the left part), which receives the previous action (the action is one-hot coded), loss and reward x_t and previous hidden states h_{t-1} and a softmax layer (the right part), which outputs the probability of selecting each deck on the next trial. The GRU layer is composed of two gates, i.e. an update gate and a reset gate, containing a set of weight parameters $\{W_r, U_r, b_r, W_z, U_z, b_z, W, U, b_{hw}, b_{hu}\}$ (b_r, b_z, b_{hw} and b_{hu} are not shown for simplicity). The output of the cells of the GRU layer h_t are connected to the softmax layer using black lines. The weight parameters in the softmax layer are represented as V .

($h_t \in R^{N_u}$) as a vector containing all the cell outputs at time t . They are initialized with the value of zero and are updated after receiving each time step input. $x_t \in R^{N_c+2}$ (N_c is the number of choices, which is 4 for the IGT) is a vector containing inputs to the network at time t , i.e. the action c_t coded using one-hot representation and the reward r_t and loss l_t received after taking action. There are two gates in GRU network, i.e. a reset gate and update gate. The update gate helps the model to determine how much of the past information needs to be passed along to the future. The output of this gate z_t is a linear combination of the input vector of the current time step x_t and the previous cell output going through the sigmoid function.

$$z_t = \sigma(W_z x_t + U_z h_{t-1} + b_z) \quad (1)$$

The reset gate is used to decide how much past information is forgotten. The formula is the same as that used for the update gate. The difference arises from the the weights and the gate's usage, which we will examine later in this section.

$$r_t = \sigma(W_r x_t + U_r h_{t-1} + b_r) \quad (2)$$

The output of the reset gate is then passed through the current memory content, in which the reset gate is used to store the relevant information from the past.

$$h'_t = \tanh(W x_t + b_{hw} + r_t \odot U h_{t-1} + b_{hu}) \quad (3)$$

In the last step, the network needs to calculate the h_t vector which holds information for the current unit and passes it down to the network. In order to do this, the update gate is needed. It determines what to collect from the current memory content h'_t and what to take from previous steps h_{t-1} . The parameters of the GRU layer include $W_z, W_r, W \in R^{N_u \times (N_c+2)}$, $U_z, U_r, U \in R^{N_u \times N_u}$, and $b_z, b_r, b_{hw}, b_{hu} \in R^{N_u}$.

$$h_t = z_t \odot h_{t-1} + (1 - z_t) \odot h'_t \quad (4)$$

The softmax layer takes outputs from the GRU layer (h_t) as its inputs and outputs the probability of choosing each action. The parameter of the softmax layer is $V \in R^{N_u \times N_a}$. The parameters of the RNN model will be $\Theta = \{V, W_z, W_r, W, U_z, U_r, U, b_z, b_r, b_{hW}, b_{hU}\}$

The RNN model was trained using the maximum-likelihood (ML) method and Categorical Cross-Entropy Loss summed over all subjects on all trials:

$$Loss = - \sum_{s=1}^S \sum_{t=1}^T y_{s_t} \log(\hat{y}_{s_t}) \quad (5)$$

The model was implemented in TensorFlow [26] and optimization was based on the use of the Adam optimizer [27].

Experiments

First of all, each laboratory dataset was divided into training and testing datasets with 80% of the data points for the former and the remaining 20% for the latter. Subjects were not mixed across training and testing sets. Two experiments are conducted in this section. The first experiment aimed to establish the need for bigger and more diverse data to train a reliable and robust RNN model that could predict human actions with reasonable accuracy. In this experiment, we trained a single RNN model on the training dataset for each laboratory and evaluated each of these models against testing datasets from each laboratory. In total, we trained 10 independent RNN models, one from each separate laboratory, and evaluated their generalization performance both on the testing dataset of their own laboratory and other data from the other laboratories.

In the second experiment, we seek to demonstrate the improvements of collaborative models compared to single models and also the comparison between the centralized learning model and distributed learning models. Eight collaborative models are trained. The collaborative models were trained directly or indirectly with bigger and more diverse datasets depending on the collaboration methods imposed between the laboratories, i.e. centralized learning (CL-based), incremental learning (IL-based), cyclic incremental learning (CIL-based), and federated learning (FL-based) paradigms. In incremental and cyclic incremental learning, three different orders of training the model were performed, i.e. smallest sample first, largest sample first and random order, yielding three models for these two learning paradigms. We expect CL-based model will behave the best among the four collaborative models, while the other three distributed learning paradigms will sacrifice competitiveness to different extents in exchange for preservation of data privacy. Thus, CL-based model will serve as the baseline model against which comparisons are made with the distributed models. The architectures of the CL-based and FL-based model paradigms are shown in Fig 2.

The hyperparameters for training the centralized model are set as follows: batch size 50; maximum epochs 250; learning rate 0.02; number of hidden cells 10 (tuned by 10-fold cross-validation, which will be introduced in the following section). To create an equal playing ground for the centralized learning and distributed learning paradigms, we need to guarantee that all models use the same hyperparameter settings or if it is not applicable due to the structure characteristic of the learning framework, we would at least make sure they see exactly the same number of samples in the whole training process. As a result, the learning rate and number of hidden cells are set as the same values for all models, while the batch size and epochs vary for different learning frameworks based on their structures. The batch size for training the local models in all three distributed learning paradigms is set as the size of the training set of the training laboratory. For the IL-based model, the maximum epochs for each laboratory should be set as 250. For the CL-based model, the epochs for each laboratory is considered

(a) Collaborative learning through centralized data sharing: CL-based model training

(b) Federated Learning to maintain data privacy: FL-based model training

Fig 2. Architectures of two collaborative learning methods for multi-laboratory collaboration. The most straightforward method (a) is based on centralized data sharing, i.e. requiring that every laboratory share data to a central node and aggregating to train a model; An alternative method and the primary focus on this work (b) is Federated learning, in which all laboratories train the same model at the same time and update their model parameters to a central sever where it is used to create a global model.

together with the number of cycles so that their multiplication is equal to 250. In our experiment, the frequency of weight transfer is every 50 epochs, thus the number of cycles is 5. For federated learning, we sharded the training datasets across 10 laboratories to match the real-world configuration of the 10 contributing laboratories. Each laboratory was regarded as a client that has different sizes of data samples, ranging from 12 to 129, for training. It is assumed in our experiment that all clients are always powered (as is the case for a typical computer but not for a mobile device), thus instead of selecting clients via an eligibility criterion from multiple client pools, all clients are involved in each round in our implementation. The number of communication rounds is set as 5 (which is the same as the cycle number in cyclic learning) and the epochs for all clients are 50. Thus, the total exposure to the training sets from all laboratories is the same as for the centralized model. Additionally, the aggregation performed in federated learning is a weighted average of laboratory updates and the weightings are calculated as the ratio of the size of each laboratory training set to the size of the aggregated training set.

After model training, the weight parameters of the models were frozen and evaluated under multiple aspects. We first examined the prediction accuracy of the eight models on the testing datasets of each laboratory. We then simulated the collaborative models in an on-policy way in the IGT (with the same payoff schemes for the subjects and the same number of trials we used for training the model). On-policy means that the models completed the task on their own by selecting the decks that they predict a representative subject would take in each simulation. Lastly, we simulated the models off-policy. Off-policy means a model uses previous actions and payoffs to make predictions about the next action. However, the actual next actions used to simulate the models are not derived from these predictions, instead they are derived from human choices. In this way, we can control the inputs the model receives and examine how they affect the predictions. Apart from examining the relative performance of distributed learning compared to centralized learning paradigm, these simulation analysis allow us to reveal the underlying learning mechanisms and decision-making strategies used by the healthy subjects in the meantime.

Tuning the number of hidden units in the GRU layer

The number of hidden units N_u in the GRU layer is a hyperparameter that has significant influence on the performance of the model. Networks with different number of hidden units ($N_u \in 5, 10, 20$) were considered and the best model was selected through 10-fold cross-validation before the single and collaborative models were trained. The total number of free parameters in both the GRU layer and softmax layer were 219, 584, and 1764 for the networks with 5, 10, and 20 hidden units of GRU layer, respectively.

As described in the previous section, each laboratory dataset was split into training and testing dataset with a ratio of 80:20. In the first experiment, 10-fold cross-validation was performed for each laboratory on their training datasets to tune the number of hidden units in the GRU layer of the single models. In the second experiment, 10-fold cross validation was performed on the aggregated training dataset of all laboratories to tune the hyperparameter for the collaborative models. In 10-fold cross validation, the dataset (either the training dataset from a single laboratory or the aggregated training dataset of all laboratories) is randomly split into 10 folds. For each unique fold, we take the fold as a hold out. The remaining folds are then taken as a training dataset. We fit a model with a specific number of hidden cells on the training dataset and evaluate the model on the hold out fold. The evaluation score is retained and the model is discarded then. We repeat this process until every fold served as the hold out set. Finally, we take the average of the evaluation scores and that will be the reported performance metric

Fig 3. 10-fold cross-validation results of models with different numbers of hidden units. Each curve is the average loss of 10 folds on training or testing dataset for each model.

for this model with a particular value of hidden units. The number of hidden units that yields the best performance is applied to train the single models or the collaborative models then. It is worth noting here that the testing datasets from all laboratories were not used to tune the hyperparameter, instead, they were used to estimate the generalization ability of the models on new, previously unseen data.

Results

Cross-validation

10-fold cross validation was conducted before training and testing the single and collaborative models to tune the number of hidden units in the GRU layer. The cross-validation results conducted on the aggregated training dataset from all laboratories are reported in 3, in which the average training and validation losses of 10 folds across epochs for centralized models with different number of hidden units are plotted (the plots for the single models can be found in the supplements). The model with 20 hidden units achieved the lowest training loss among the three models (the purple line), however, the validation loss of this model (the brown line) did not decrease steadily after 250 around epochs, instead it increased indicating overfitting. The lowest mean validation losses was achieved by 10 hidden units in the GRU layer (the green line). This setting was then used for training and simulating the four collaborative models. In a similar way, the number of hidden units for all single models were set as 5 based on their cross-validation results.

The need for more numerous and diverse data

In this experiment, ten single models were trained with the tuned hyperparameters using the training datasets from each laboratory respectively. These models were then evaluated using the testing datasets available from each laboratory. Fig. 4 shows ten

Fig 4. Single lab models tested against testing datasets from each of the laboratories. The vertical axis represents models trained on a single laboratory dataset, and the horizontal-axis represents the testing dataset of each independent laboratory. AVG is the average of each laboratory model performance over all laboratories.

single models' testing results both against their own testing datasets and testing datasets of other laboratories. The chance probability of predicting the next action taken by the subjects correctly is 25% since there are four options in the IGT. However, the prediction accuracy of several single models on some of the testing datasets was lower than the chance probability, which is obviously not satisfactory performance. It can be seen that most single models performed the best on the testing dataset of *laboratory9*, except for the three models that trained on the least size of laboratory samples, i.e. *laboratory1* (12), *laboratory3* (15), and *laboratory6* (20). These laboratories with small number of training samples did not perform well on their own testing sets nor did they generalize well to any other testing sets from other laboratories. The best average model performances were shown on *laboratory2* and *laboratory7*, which have the largest sizes of training samples, 129 and 122, respectively. These results indicate that more numerous and diverse data are needed for a laboratory site in order to train reliable and generalizable models, at least in the contest of training RNN models.

The benefits of training collaborative models

In this section, we aim to demonstrate 1) the benefits of laboratory collaboration and 2) the feasibility of conducting distributed learning when data sharing barriers are present. We compared the prediction accuracy of the eight collaborative models with the single models for the first purpose. The results are illustrated in Fig 5. As a result, apart from random and descending ordered IL and CL models, the prediction accuracy of the collaborative models were much improved when testing on each laboratory's data. Notably, the average performance over all laboratories of the CL-based model were improved to 55%, 6% points higher than the highest accuracy achieved by the single model of *laboratory2*, and the corresponding increase for FL-based, CIL_ascend-based and IL_ascend-based model was 4%, 4% and 3% points, respectively. Since random and

Fig 5. Two collaborative models, i.e. CL-based model and FL-based model, tested against testing datasets from each of the laboratories. The Y-axis represents the two collaborative models trained on the training datasets from all laboratories and the X-axis represents the testing dataset of each independent laboratory. AVG is the average of each collaborative model performance over all laboratories

descending ordered IL and CIL models did not outperform the single models, only the ascending ordered models were selected to conduct further evaluation. We will refer to these as IL-based and CIL-based model instead of CIL_ascend-based and IL_ascend-based in the following analysis for simplicity of reference. The outperformed IL-based and CIL-based models and the other two collaborative models were evaluated from multiple aspects for the second purpose.

On-policy simulation

After training the four collaborative models, we froze the weight parameters of the models and simulated them on-policy in the task. We fed the model with the same payoff schemes as for the subjects and the models selected the decks autonomously based on what they had learned from human behaviours. Since we trained on 491 experimental subjects from training datasets from all studies, we simulated 491 fake agents for each collaborative model as well.

0.0.1 The average probability of selecting each deck over subjects

The SUBJ column in Fig 6 shows the average probability of choosing each deck over all experimental subjects. All subjects selected DeckA the least compared to other Decks, which is not surprising because DeckA is the bad deck that yields negative long-term payoff and with more frequent losses. Consistent with human's choices, the probability of choosing DeckA for both centralized and distributed agents was the lowest. Neither CL-based agents ($\eta = 0.005$, $SE = 0.004$, $p = 0.26$) nor the FL-based and the IL-based agents (**FL-based:** $\eta = 0.001$, $SE = 0.004$, $p = 0.91$, **IL-based:** $\eta = 0.001$, $SE = 0.004$, $p = 0.83$) are significantly different for subjects in terms of the probability of selecting DeckA, although the CIL-based selected slightly significant more DeckA than subjects ($\eta = 0.01$, $SE = 0.005$, $p = 0.03$). FL-based agents had significant aversion to DeckD ($\eta = -0.02$, $SE = 0.010$, $p = 0.03$) and CIL-based and IL-based agents had significant preference to DeckB (**CIL-based:** $\eta = 0.03$, $SE = 0.009$, $p < 0.001$, **IL-based:** $\eta = 0.02$, $SE = 0.008$, $p = 0.010$) and significant aversion to

Fig 6. Probability of selecting each deck (average over subjects). SUBJ refers to the data of the experimental subjects; CL-based, FL-based, CIL-based, and IL-based are the on-policy simulations of the four collaborative models trained on the IGT with the same payoff scheme and same number of trials the subjects completed, respectively.

DeckC (**CIL-based:** $\eta = -0.04$, $SE = 0.009$, $p < 0.001$, **IL-based:** $\eta = -0.03$, $SE = 0.009$, $p < 0.001$) compared to the subjects, while no significant difference was observed between CL-based agents and subjects on all of these decks. This can also be seen from the plot of the overall probabilities of selecting DeckB and DeckC where IL-based and CL-based agents are relatively higher and lower than that of the subjects, respectively and the probabilities of selecting DeckD of FL-based agents are relatively lower.

It is worth noting that subjects prefer decks with infrequent losses, DeckB and DeckD, compared to decks with frequent losses, DeckA and DeckC, although they ended up selecting significantly more good decks (DeckC and DeckD) than bad decks (DeckA and DeckB) ($\eta = 0.078$, $SE = 0.010$, $p < 0.001$). CL-based, FL-based and IL-based agents managed to learn this strategy from the subjects' actions, selecting more good decks than bad decks (**CL-based:** $\eta = 0.059$, $SE = 0.010$, $p < 0.001$, **FL-based:** $\eta = 0.059$, $SE = 0.011$, $p < 0.001$, **IL-based:** $\eta = 0.035$, $SE = 0.008$, $p < 0.001$), whereas CIL-based agents did not learn this characteristic of behavioural pattern, the probability of selecting good decks was not significantly higher than bad decks ($\eta = -0.006$, $SE = 0.010$, $p = 0.51$).

0.0.2 The fluctuation of probability of selecting each deck over trials

The analysis in the previous section suggested that CL-based, FL-based, and IL-based agents selected more good decks than bad decks, which is more consistent with the human behavioral data compared to CIL-based agents. There are multiple possible strategies that the models could follow to obtain the final proportions of each deck. Here we aim to examine whether the models were using similar strategies to that used by subjects. We examined the fluctuation of deck preferences of the subjects and model agents over trials. We divided the 95 trials into 10 blocks of 10 trials for the first 9 blocks and 5 trials for the last block. The proportion of each deck selection in each block and the learning scores (i.e. the difference between the number of good deck selections and the number of bad deck selections) for each subject was calculated. Fig 7

Fig 7. Mean proportion of choices from each deck within 10 blocks of experimental subjects and four simulated model agents.

demonstrates the average proportion of choosing each deck as a function of the 10 blocks of the experimental subjects and four model agents. The behaviour characteristics of CL-based and FL-based agents were more similar to experimental subjects. In the beginning of the task, subjects, CL-based and FL-based agents selected DeckB most frequently, but its representation as a proportion of choices decreased after approximately 4 blocks for subjects and CL-based agents, 5 blocks for FL-based agents. At this stage the proportion choices which resulted in the selection of the "good" DeckD exceeded that for DeckB and became the highest. The proportion of choices for selecting DeckC was low in the beginning and gradually increased over trials as subjects learned more from the task. The final proportion of DeckC was a little higher than that of DeckB both for subjects, and CL-based agents, although almost equal for FL-based agents. The proportion of selecting DeckA was always the lowest and decreasing throughout the task. For CIL-based and IL-based agents, however, the proportion of choosing DeckB was almost always the highest. Although consistent with the subjects, the proportion of DeckC for CIL-based agents increased over trials, it never became as high as DeckB and DeckD. It was even worse for IL-based agents because the proportion of choosing DeckC did not vary too much over trials.

Fig 8 shows the learning scores across ten blocks of the IGT. A learning process was apparent both for experimental subjects and CL-based and FL-based agents, in which the learning score progressively improved over blocks, although there was a clear dip in block 10 for subjects and FL-based agents and block 8 for CL-based agents. To quantify the differences between the subjects and the two kinds of model agents over blocks on learning scores, the repeated measures ANOVA tests were conducted in the form of $10 \text{ (blocks)} \times 2 \text{ (groups)}$. The ANOVA test conducted between human and CL-based agents and human and FL-based agents revealed no significant interaction between the two factors (**CL-based:** $F(7.44, 7290) = 1.39, p = 0.20$, **FL-based:** $F(7.1, 6954) = 1.82, p = 0.08$), however, there was a significant main effect of blocks (**CL-based:** $F(7.44, 7290) = 67.63, p < 0.001$, **FL-based:** $F(7.1, 6954) = 60.64, p < 0.001$), with the interpretation that the learning scores of both subjects and CL-based and FL-based model agents varied across blocks and the CL-based and FL-based model did not behave significantly different across subjects in the learning process. The ANOVA test performed between human and the other two collaborative agents showed that the effects of blocks and group were both significant on the learning score (**CIL-based:** $F(7.85, 7693) = 47.11, p = 0.03$, $F(1, 980) = 17.61, p < 0.001$;

Fig 8. The IGT learning scores across ten different blocks of experimental subjects and simulated model agents.

IL-based: $F(7.71, 7558) = 22.66$, $p = 0.02$, $F(1, 980) = 5.69$, $p = 0.02$) and there was a significant interaction between groups and blocks on the learning score (**CIL-based:** $F(7.85, 7693) = 5.98$, $p < 0.001$, **IL-based:** $F(7.71, 7558) = 15.76$, $p < 0.001$). This result suggested that the learning process of CIL-based and IL-based agents were significantly different from experimental subjects, which can also be seen from the figure, in which the learning scores of CIL-based were lower than the subjects almost over all blocks and IL-based agents did not present apparent learning trend in the process.

0.0.3 The switch probability after receiving losses

We finally investigated the immediate effect of loss on choice. The reason why reward effect is not considered here is because subjects can always get rewards no matter which deck they choose, 100 for DeckA and DeckB and 50 for DeckC and DeckD, according to the payoff scheme. Fig 9 shows the effect of receiving a loss in the previous trial on the probability of switching to the other deck in the next trial. For the experimental subjects, receiving a loss significantly increased the probability of switching to other decks ($\eta = 0.218$, $SE = 0.010$, $p < 0.001$). As the figure shows, same pattern was established by all collaborative agents (**CL-based:** $\eta = 0.209$, $SE = 0.007$, $p < 0.001$; **FL-based:** $\eta = 0.201$, $SE = 0.007$, $p < 0.001$; **CIL-based:** $\eta = 0.246$, $SE = 0.007$, $p < 0.001$; **IL-based:** $\eta = 0.179$, $SE = 0.007$, $p < 0.001$), suggesting that the strategy of avoidance to losses used by the four models was similar to that seen in the subjects' behavior. However, the switch probabilities of CL-based agents are more closer to that of the experimental subjects visually and FL-based agents rank second. The switch probability of IL-based agents was obviously higher than the subjects when there was no loss.

Off-policy simulation

In the off-policy simulation analysis, we fed the four collaborative models with a variety of actions and corresponding rewards and losses specified in the payoff schemes, and

Fig 9. Probability of switch to different deck based on whether received losses or not in the previous trial, average over subjects.

recorded the predictions of each model in response to each input set. Simulations of the models are shown in Fig 10. Each panel shows a separate simulation across 30 trials (horizontal axis). For the total of 30 trials, the action that was fed into the model was DeckA (top left), DeckB (top right), DeckC (bottom left), or DeckD (bottom right) for each simulation, respectively. The reward and losses associated with these trials were the same as specified in the payoff scheme of IGT. As mentioned earlier, based on the IGT payoff scheme, the player would always receive rewards no matter which deck they choose and only the rewards magnitude varies depending on which deck they choose, while they only receive losses occasionally. We only marked trials where there were losses with blue dots in the graphs.

The four collaborative models behaved quite similarly in responding to the losses according to the plots, especially the FL-based model and the CL-based model. Receiving a loss from all decks caused a dip in the probability of staying with that deck, which showed a tendency to switch to another deck. This is consistent with the observations we obtained in Fig 9 that the probability of switching is higher when receiving a loss compared to no loss. It seems the CIL-based model was more sensitive to the losses given by DeckC as deeper dips were caused compared to other models. The dips caused by DeckB and DeckD were the deepest, which is not surprising, because DeckB and DeckD are decks with less frequent but larger losses. It is reasonable that the switch probability should be higher if you received a significant loss from that deck choice. The probability of choosing DeckA, DeckB, DeckC, and DeckD was higher than that of other decks in each simulation. This implies sticking with the previously taken deck was an element of a deck selection strategy of the four models. We can also see that the perseverance with DeckA is not as strong as to the other three decks as the probability of selecting DeckA predicted by the two models is almost never more than 80% and this is even more obvious for IL-based model. This result is consistent with the result reported in Fig 6 and Fig 7 where the proportion of selecting DeckA was the lowest both across subjects and trials. Interestingly, the perseverance effect to DeckC was smaller and the sensitivity to the losses of DeckC was higher for CIL-based model and relative to other models.

Off-policy simulations of the four collaborative models

Fig 10. Off-policy simulations of the two collaborative models. Each panel shows a simulation of 30 trials.

Discussion

The results demonstrate, as expected, that models trained on a single laboratory’s dataset only, perform poorly both on their own testing sets and testing sets from other laboratories. The average prediction accuracy of the model of *laboratory1*, which had the smallest number of subjects among studies, on ten testing datasets was 31%, only 6% higher than the chance correct probability. This result highlights the need for larger and more representative datasets and therefore emphasises the relevance of training using collaborative models involving subjects from other studies.

Given the difficulties of sharing data across studies, institutions and jurisdictions, it is attractive to use distributed learning paradigms to replace traditional centralized data sharing to make laboratory collaboration easier. However the decision to adopt such an approach is only justifiable so long as the resultant model yields comparable performance with the data pooling approach. While the work presented here cannot be fully comprehensive, it presents for the first time a representative sampling of eight collaborative models, one of which based on centralized data sharing while the other seven were based on distributed learning including one FL-based, three CIL-based and three IL-based, were trained on the training datasets from all laboratories for comparison. It was found that the order of the laboratories was an important factor that influenced the performance of CIL-based and IL-based models. An ascending order based on the laboratory data-set sizes was the best strategy for both of these two paradigms and the corresponding models were selected, along with the FL-based model, to be further evaluated with the centralized model from multiple aspects. When examining their prediction capability on the unseen testing datasets the prediction accuracy was improved by 24%, 24%, 24%, and 23% points compared to the worst single model on the same testing data-sets for CL-based, FL-based model, CIL-based and IL-based model respectively. This highlights the value of involving larger numbers of data samples for training a deep learning model. It is also clear that FL-based and CIL-based approaches can both achieve competitive performance with CL-based model while retaining the benefit of enhanced data privacy. Another interesting finding is that

all the collaborative models performed clearly better on the testing sets of *laboratory9*, *laboratory5*, and *laboratory4* and the CIL-based and IL-based models did not obviously suffer from the 'catastrophic forgetting' problem as reported in the literature [11, 12, 17]. This is apparent because they did not present performance biased to the the data they had mostly recently seen (i.e. *laboratory2* for the ascending order case, *laboratory1* in the descending order case, and *laboratory10* for the random order case).

We then evaluated the capability of the distributed model in capturing characteristics of human decision-making strategies used on the IGT in comparison with CL-based model through on-policy and off-policy simulations. In the on-policy simulation, 491 fake agents, which is the same as size as the training data-set for training the collaborative models, were generated for the four models. The proportions for choosing each deck averaged over subjects was examined first. The results revealed that CL-based agents behaved more similarly to human subjects exhibiting similar choice proportions on all decks, while the agents generated by the three distributed models had significantly different deck preferences and aversions on one or two particular decks compared to the subjects choices. This characteristic was also reflected in the fluctuation of the probability of selecting each deck over trials except that FL-based agents performed as well as CL-based agents in this comparison. The probability of choosing DeckB was almost always the highest for CIL-based and IL-based agents, whereas the probability of choosing DeckD surpassed DeckB after four blocks for experimental subjects and CL-based agents and five blocks for FL-based agents. The probability of choosing DeckC for CIL-based and IL-based agents never increased as high as that of the subjects and the other two distributed agents by the end of the game. The comparison analysis of the IGT learning scores across trials between the experimental subjects and model agents also suggested CIL-based and IL-based model were weaker in recognizing the good decks as the learning scores of these two agents did not change for IL-based agents or were always lower than the other groups throughout the task, although the overall trend of the learning score was increasing for CIL-based agents. In terms of the strategy of dealing with losses, all model agents adopted the same strategy as human subjects, i.e. switching more often when they received a loss compared to no loss. It is worth noting here that although no statistically significant weakness was identified for CIL-based and IL-based agents, the average switching probability of receiving losses was relatively higher than that of the human subjects visually, while the CL-based and FL-based agents were almost the same as the subjects on this figure. The characteristic of switching more after losses was validated again in the following off-policy simulation, where the actions fed to the model were specified by us and the models were responsible for predicting the next action to be taken for the next trial. According to the model's prediction, the probability of sticking with a deck was decreased whenever there a loss incurred. CL-based and FL-based models reacted almost the same way again in this simulation, but IL-based model demonstrated less perseverance to DeckA and the CIL-based model was more sensitive to losses compared to the other agents.

In summary, consistent with the findings of Sheller et al. [17], where the distributed learning methods were evaluated in the field of medicine, the FL-based model achieved the best prediction accuracy and capability among the distributed learning paradigms in learning and mimicking the decision-making strategies used by the experimental subjects. Although CIL-based and IL-based models also achieved considerable prediction accuracy when compared to the CL-based model, they failed in capturing some features such as subtle deck preferences between DeckB and DeckC and learning curve shapes for the subjects. These findings can help inform researchers when developing multi-laboratory collaborations in the field of cognitive and behavioural science. It suggests simulated datasets may be a useful means for determining to what

extent distributed learning compromises the quality of the models developed. Additionally, emerging improvements to FL [28, 29] hold further promise for the narrowing of the gap in performance between centralized and federated learning.

Conclusion

Deep learning models, especially RNN models, have emerged as useful methods suitable for characterising and even predicting human behaviour in the context of computational modelling tasks. A successfully trained model generalizes well with respect to data it has not been seen before. Successful training in turn requires a large representative dataset. However, typical cognitive and behaviour studies are often relatively small in terms of the numbers of subjects recruited, which limits the application of deep learning models to human behaviour prediction. A natural way of obtaining bigger and more diverse data is sharing data between laboratories. Nevertheless, sharing data is never easy for the purpose of protecting personal privacy, especially when involving data that could potentially identify individuals and institutions that are internationally distributed. Distributed learning is a promising method for addressing this issue through sharing the model parameters instead of data themselves. In this study, the feasibility and reliability of applying several distributed learning paradigms to the field of cognitive studies aiming to characterise human decision-making was examined. We demonstrated that federated learning outperformed the other two distributed learning paradigms in learning the learning features of human subjects in this context. The use of federated learning over centralized learning has the immediate advantage of keeping data confidential. We hope that the development of this techniques and emerging derivations of the approach will allow for data-private collaborative training over various subjects with different ages, education backgrounds, social status collected by different laboratories. Such collaborations may enable the development of more robust models that satisfy clinical needs and thus will be helpful in terms of translating advances in cognitive studies into clinical improvements for patients suffering from psychiatric conditions.

Acknowledgments

References

1. Ratcliff R, McKoon G. The diffusion decision model: theory and data for two-choice decision tasks. *Neural computation*. 2008;20(4):873–922.
2. Summerfield C, Tsetsos K. Building bridges between perceptual and economic decision-making: neural and computational mechanisms. *Frontiers in neuroscience*. 2012;6:70.
3. Dayan P, Niv Y. Reinforcement learning: the good, the bad and the ugly. *Current opinion in neurobiology*. 2008;18(2):185–196.
4. Rescorla RA. A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. *Current research and theory*. 1972; p. 64–99.
5. Dezfouli A, Griffiths K, Ramos F, Dayan P, Balleine BW. Models that learn how humans learn: the case of decision-making and its disorders. *PLoS computational biology*. 2019;15(6):e1006903.
6. Dezfouli A, Ashtiani H, Ghattas O, Nock R, Dayan P, Ong CS. Disentangled behavioral representations. *bioRxiv*. 2019; p. 658252.
7. Fintz M, Osadchy M, Hertz U. Using Deep Learning to Predict Human Decisions and Cognitive Models to Explain Deep Learning Models. *bioRxiv*. 2021;.
8. Zech JR, Badgeley MA, Liu M, Costa AB, Titano JJ, Oermann EK. Variable generalization performance of a deep learning model to detect pneumonia in chest radiographs: a cross-sectional study. *PLoS medicine*. 2018;15(11):e1002683.

9. Huys QJ, Maia TV, Frank MJ. Computational psychiatry as a bridge from neuroscience to clinical applications. *Nature neuroscience*. 2016;19(3):404–413.
10. Voigt P, Von dem Bussche A. The eu general data protection regulation (gdpr). *A Practical Guide*, 1st Ed, Cham: Springer International Publishing. 2017;10:3152676.
11. Sheller MJ, Reina GA, Edwards B, Martin J, Bakas S. Multi-institutional deep learning modeling without sharing patient data: A feasibility study on brain tumor segmentation. In: *International MICCAI Brainlesion Workshop*. Springer; 2018. p. 92–104.
12. Chang K, Balachandar N, Lam C, Yi D, Brown J, Beers A, et al. Distributed deep learning networks among institutions for medical imaging. *Journal of the American Medical Informatics Association*. 2018;25(8):945–954.
13. French RM. Catastrophic forgetting in connectionist networks. *Trends in cognitive sciences*. 1999;3(4):128–135.
14. Konečný J, McMahan HB, Ramage D, Richtárik P. Federated optimization: Distributed machine learning for on-device intelligence. *arXiv preprint arXiv:161002527*. 2016;.
15. Hard A, Rao K, Mathews R, Ramaswamy S, Beaufays F, Augenstein S, et al. Federated learning for mobile keyboard prediction. *arXiv preprint arXiv:181103604*. 2018;.
16. Li T, Sahu AK, Talwalkar A, Smith V. Federated learning: Challenges, methods, and future directions. *IEEE Signal Processing Magazine*. 2020;37(3):50–60.
17. Sheller MJ, Edwards B, Reina GA, Martin J, Pati S, Kotrotsou A, et al. Federated learning in medicine: facilitating multi-institutional collaborations without sharing patient data. *Scientific reports*. 2020;10(1):1–12.
18. Rakhlin A, Shamir O, Sridharan K. Making gradient descent optimal for strongly convex stochastic optimization. *arXiv preprint arXiv:11095647*. 2011;.
19. Zhao Y, Li M, Lai L, Suda N, Civin D, Chandra V. Federated learning with non-iid data. *arXiv preprint arXiv:180600582*. 2018;.
20. Bechara A, Damasio AR, Damasio H, Anderson SW. Insensitivity to future consequences following damage to human prefrontal cortex. *Cognition*. 1994;50(1-3):7–15.
21. Steingrover H, Pachur T, Šmíra M, Lee MD. Bayesian techniques for analyzing group differences in the Iowa Gambling Task: A case study of intuitive and deliberate decision-makers. *Psychonomic bulletin & review*. 2018;25(3):951–970.
22. Bechara A, Damasio H. Decision-making and addiction (part I): impaired activation of somatic states in substance dependent individuals when pondering decisions with negative future consequences. *Neuropsychologia*. 2002;40(10):1675–1689.
23. Gers FA, Schmidhuber J, Cummins F. Learning to forget: Continual prediction with LSTM. *Neural computation*. 2000;12(10):2451–2471.
24. Hochreiter S, Schmidhuber J. Long short-term memory. *Neural computation*. 1997;9(8):1735–1780.
25. Cho K, Van Merriënboer B, Gulcehre C, Bahdanau D, Bougares F, Schwenk H, et al. Learning phrase representations using RNN encoder-decoder for statistical machine translation. *arXiv preprint arXiv:14061078*. 2014;.
26. Abadi M, Agarwal A, Barham P, Brevdo E, Chen Z, Citro C, et al. Tensorflow: Large-scale machine learning on heterogeneous distributed systems. *arXiv preprint arXiv:160304467*. 2016;.
27. Kingma DP, Ba J. Adam: A method for stochastic optimization. *arXiv preprint arXiv:14126980*. 2014;.
28. Ozdayi MS, Kantarcioglu M, Iyer R. Improving Accuracy of Federated Learning in Non-IID Settings. *arXiv preprint arXiv:201015582*. 2020;.
29. Wang H, Kaplan Z, Niu D, Li B. Optimizing federated learning on non-iid data with reinforcement learning. In: *IEEE INFOCOM 2020-IEEE Conference on Computer Communications*. IEEE; 2020. p. 1698–1707.