



HAL
open science

Simulation of virtual cohorts increases predictive accuracy of cognitive decline in MCI subjects

Igor Koval, Stéphanie Allassonnière, Stanley Durrleman

► To cite this version:

Igor Koval, Stéphanie Allassonnière, Stanley Durrleman. Simulation of virtual cohorts increases predictive accuracy of cognitive decline in MCI subjects. 2019. hal-02089881

HAL Id: hal-02089881

<https://hal.science/hal-02089881>

Preprint submitted on 4 Apr 2019

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Simulation of virtual cohorts increases predictive accuracy of cognitive decline in MCI subjects

Igor Koval¹, Stéphanie Allassonnière², and Stanley Durrleman¹, for the Alzheimer’s Disease Neuroimaging Initiative

¹ Inria Paris-Rocquencourt, Inserm U1127, CNRS UMR 7225, Sorbonne Universités, UPMC Univ Paris 06 UMRS 1127, Institut du Cerveau et de la Moelle épinière, ICM, F-75013, Paris, France

² INSERM UMRS 1138, Centre de Recherche des Cordeliers, Université Paris Descartes, Paris, France

Abstract. The ability to predict the progression of biomarkers, notably in NDD, is limited by the size of the longitudinal data sets, in terms of number of patients, number of visits per patients and total follow-up time. To this end, we introduce a data augmentation technique that is able to reproduce the variability seen in a longitudinal training data set and simulate continuous biomarkers trajectories for any number of virtual patients. Thanks to this simulation framework, we propose to transform the training set into a simulated data set with more patients, more time-points per patient and longer follow-up duration. We illustrate this approach on the prediction of the MMSE of MCI subjects of the ADNI data set. We show that it allows to reach predictions with errors comparable to the noise in the data, estimated in test/retest studies, achieving an improvement of 37% of the mean absolute error compared to the same non-augmented model.

Keywords: Data augmentation · Sequence data · Virtual cohort · Disease stage prediction · Alzheimer’s Disease

1 Introduction

Predicting the future progression of patients with neurodegenerative diseases (NDD) is a key challenge to treat patients at an earlier stage than today or to better evaluate drug efficacy in clinical trials. Longitudinal data sets, consisting of repeated observations of the same patients over time, play a central role to describe and predict disease progression. Machine Learning techniques, trained on sequence data (multiple observations per patient), have seized this challenge. However, the databases often lack patients with sufficient follow-up visits, a future visit to predict and a sufficiently large delay in between, leading to poor generalization on test data. More importantly, the different experimental settings - time to prediction or patients/visits presenting the considered feature(s) - are difficult to compare as they involve a subset of the initial cohort that has specific characteristics in terms of number of patients, number of follow-up visits

and duration. It is often impossible to balance the training and test set for all these characteristics. This problem prevents from a reliable comparison of the algorithms, some being better due to the size or characteristics of the training set, rather than the intrinsic performance of the algorithm or choice of the features.

To increase the size of real data sets, data augmentation techniques have been developed : virtual data are drawn with the intention to reproduce the characteristics of the data in the initial cohort. Most of the literature focuses on techniques for independent and identically distributed observations such as image classification [7], text categorization [6] or speech recognition [3]. Due to the unrealistic hypothesis for sequence data, some techniques have been proposed for uni-dimensional time-series. They rely on a continuous transformation of the time domain by warping, slicing or sliding the time window [5]. Such techniques do not apply to NDD as the temporal pattern is key in the disease progression. Recently, Generative Adversarial Networks (GAN) have received interest due to the characteristic of the generative part of the model: it can sample virtual realistic data. It is however non-trivial to generate sequence data as there is no straightforward way to propagate the gradient updates from the discriminator to the generator [10]. Furthermore, these models rely on large training data sets, which are typically inaccessible in the targeted medical applications.

In this paper, we propose to use a simulation framework to perform data augmentation for sequence data in the presence of small training samples, and to evaluate to which extent it increases the performance of a predictive algorithm. The model introduced in [9] recombines short-term individual observations at different disease stages to estimate a long-term scenario of disease progression. The model estimates also parameters that change the pattern of progression, the pace of progression and the age at onset, so that it can reconstruct a continuous trajectory by fitting individual data, or simulate entirely synthetic trajectories by sampling the empirical distribution of the parameters. This model, once trained on a small and unbalanced training data set, may be used therefore to re-sample the trajectories of the training subjects by adding new visits and covering larger time span, and even increase the number of subjects by adding data sampled from simulated trajectories. This augmented virtual cohort may then be used in lieu of the training samples to train a predictive algorithm. We propose to evaluate the performance of the algorithm trained with this augmented and virtualized data, as compared to the original training set, for the prediction of the cognitive decline in subjects with mild cognitive impairments.

2 Sequence-based prediction

In the following, we consider a longitudinal data set $\mathbf{y} = (\mathbf{y}_{ij}, t_{ij})_{\substack{1 \leq i \leq n \\ 1 \leq j \leq n_i}}$ where the i -th subject has been observed n_i times at ages $t_{1i} < \dots < t_{in_i}$ with $\mathbf{y}_{ij} \in R^d$ a set of biomarkers. Each observation corresponds to a snapshot of the individual spatiotemporal trajectory. We aim to predict the value of one biomarker in ΔT years after a given visit, knowing the values of the biomarkers at the previous visits.

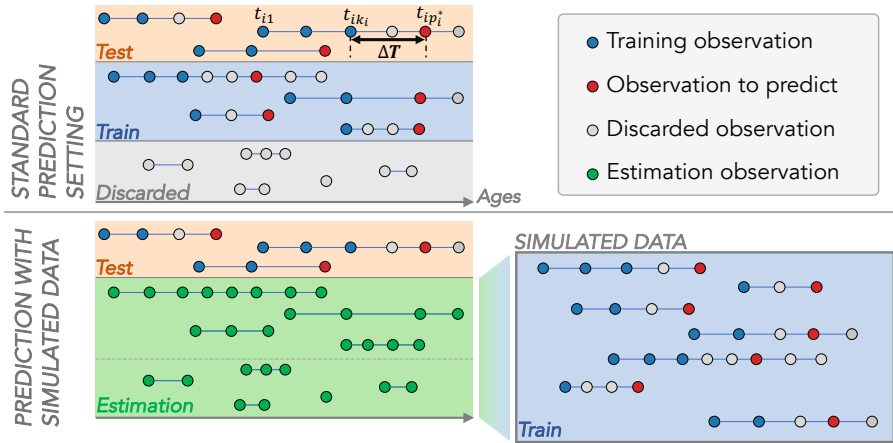


Fig. 1. The top row describes the standard prediction setting where the data set is split in a train and test set. A fixed time-delay ΔT is set between the input visits (blue dots) and the target prediction (red dot). It leads to discarding visits in between (grey dots) or entire subjects that do not present sufficient follow up visits (discarded set). The bottom row corresponds to the procedure with simulated data : the training set is composed of virtual patients that are simulated thanks to the estimation set.

2.1 Standard prediction setting

In a standard setting, one needs to discard patients that do not have sufficient follow-up visits to cover the required temporal span, i.e. $t_{in_i} - t_{i1} < \Delta T$, as shown on the top row of Figure 1. For the remaining patients, the input biomarkers $(\mathbf{y}_{ij})_{1 \leq j \leq k_i}$ at some early visits $(t_{ij})_{1 \leq j \leq k_i}$ (blue dots) are used to predict the biomarker $\mathbf{y}_{ip_i^*}$ at age $t_{ip_i^*}$ (red dot) such that $t_{ip_i^*} = t_{ik_i} + \Delta T$. This task may be achieved by any machine learning algorithm, for instance a neural network. In this setting, possible intermediate visits t_{ij} such that $t_{ik_i} < t_{ij} < t_{ip_i^*}$ are discarded (grey dots). If multiple splits of input/output visits $((\mathbf{y}_{ij}, t_{ij})_{1 \leq j \leq k_i}, (\mathbf{y}_{ip_i^*}, t_{ip_i^*}))$ are possible, one is selected at random. Once the input and target visits have been selected for each patient, they are split into the train and test set.

The longer ΔT , the fewer patients remain in the train and test set and the fewer visits per remaining patients there are. Therefore, when ΔT is varied, the size and composition of the train and test set may vary dramatically, and thus the performance of the predictive algorithm. This problem is even more critical if several biomarkers that are not observed at every visit are used.

2.2 Sequence data simulation

We consider a generative mixed-effect model such that the individual observations at time-point t_{ij} is $\mathbf{y}_{ij} = f(\theta, \mathbf{z}_i, t_{ij}) + \epsilon_{ij}$, $\epsilon_{ij} \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(\mathbf{0}, \sigma^2 \text{Id}_d)$ where θ

corresponds to the fixed-effects and \mathbf{z}_i the random effects associated to the subject i . Assuming that it is possible to estimate θ and $(\mathbf{z}_i)_{1 \leq i \leq n}$ given a training data set, we can draw a new sample $\mathbf{z}_{i'}$ from the empirical distribution of the $(\mathbf{z}_i)_{1 \leq i \leq n}$. It corresponds to a new individual for which is it possible to simulate new observations at arbitrary time-point $(t_{i'j})_{1 \leq j \leq n_{i'}}$.

An example of such model is introduced in [9], where the authors consider that the fixed-effects θ define the group-average disease progression and its variability in the population. This fixed-effects, of small dimension relatively to the feature space, can be estimated, thanks to the MCMC-SAEM [1, 4], with short-term observations for a relatively small number of subjects. As for the i -th individual trajectory, the authors consider that it derives from the group-average trajectory according to the random-effects $\mathbf{z}_i = (\alpha_i, \tau_i, (s_{ij})_{1 \leq j \leq N_s})$ where α_i corresponds to the pace of progression and τ_i relates about the time delay between the group-average and the individual scenario. On top of these temporal parameters that impact the observation coordinates similarly, the space-shifts $(s_{ij})_{1 \leq j \leq N_s}$ characterize the inter-coordinates variations ($N_s \leq N$). These random-effects are learnt by optimizing the individual complete likelihood $p((\mathbf{y}_{ij})_{1 \leq j \leq n_i}, \mathbf{z}_i; \theta) = p((\mathbf{y}_{ij})_{1 \leq j \leq n_i} | \mathbf{z}_i; \theta) p(\mathbf{z}_i; \theta)$ thanks to the L-BFGS-B method [11].

It is possible to draw a set of individual parameters $\mathbf{z}_{i'}$ by, first, simulating the temporal parameters $(\alpha_{i'}, \tau_{i'})$ with a kernel density estimation on the empirical distribution $(\alpha_i, \tau_i)_{1 \leq i \leq n}$. Then, considering the multivariate Gaussian distribution $\mathcal{N}(\mu, \Sigma)$ estimated on the whole learnt random effects $(\mathbf{z}_i)_{1 \leq i \leq n}$, it is possible to draw $((s_{i'j})_{1 \leq j \leq N_s} | \alpha_{i'}, \tau_{i'}) \sim \mathcal{N}(\tilde{\mu}, \tilde{\Sigma})$ where $\tilde{\mu}$ and $\tilde{\Sigma}$ are functions of μ and Σ [8].

2.3 Prediction setting with simulated patients

From the whole data set, we first select an estimation subset that includes the subjects that were discarded in the standard procedure along with some subjects that have more follow-up duration, as shown on the bottom row of Figure 1. The remaining subjects form the test set, with same split constraints on the input and target visits as in the standard prediction setting.

Once θ and $(\mathbf{z}_i)_{1 \leq i \leq n}$ are learnt on the estimation set, we can simulate an arbitrary number of patients and visits per individual that will be used as the training set. To have similar characteristics in the training and test set, the time between two simulated visits is kept similar to the one of the real patient (e.g. one year), and, the input and target visits are split as before. It is important to mention that the estimation procedure now uses the visits t_{ij} such that $t_{ik_i} < t_{ij} < t_{ij^*}$ that were previously discarded (grey dots).

In the following, to assess the quality of the simulated data only, we prevent ourselves from using patients from the estimation set that are eligible in term of number of visits to be used in the training set (upper part of the estimation set of Figure 1) : that the training procedure relies on the simulated patients only. In other cases, it is indeed possible - and recommended - to add them to the training set.

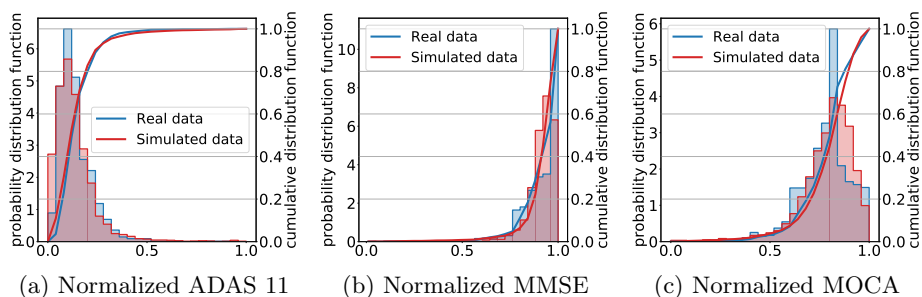


Fig. 2. Empirical histograms and cumulative distributions of original cohort (blue) and of the simulated data (red) based on a model that estimate the evolution of the normalized ADAS Cog with 11 and 13 items and the normalized MOCA.

3 Experimental results

3.1 Data and experimental setting

The experiments focus on the prediction of the mini-mental state examination (MMSE) for subjects with mild cognitive impairment (MCI) from the ADNI database. Therefore, our cohort includes early and late MCI, but also MCI who converted to Alzheimer’s Disease and stable MCI. We considered predictions at 1, 2, 3 and 4 years, based on a set of features among the MMSE, Alzheimer’s disease assessment scale - cognitive subscale (ADAS-Cog) with 11 or 13 items, clinical dementia rating sum of boxes (CDRSB), the Montreal cognitive assessment (MOCA) and the functional assessment questionnaire (FAQ). One estimation set was defined per subset of features used in the predictive model, e.g. an estimation set that simulate patients with MMSE, ADAS-11 and ADAS-13 was defined for the predictions that are based on these features.

To assess the performance of the simulation framework, we create a virtual cohort with the same characteristics (number of patients and time-points) as the estimation set. The empirical data distribution for the real and virtual cohort are nearly identical as shown on Figure 2 for three different features.

We choose a long short-term memory (LSTM) neural network, with 10 hidden dimensions, stacked with a linear layer, as our algorithm to predict future feature values. The Mean Squared Error (L2-norm) loss is optimized thanks to the ADAM optimizer (learning rate of 10^{-3} and weighted decay of 10^{-5}). To prevent the model from overfitting, a subset of the real patients, namely the validation set, is used to apply the early stopping criterion procedure : it stops the training if no loss improvement is detected from a given number of epoch on the validation set. The code is available at [shown after paper acceptance].

The results report the Mean Absolute Error (MAE). To estimate the variance of the estimation procedure, the results are presented with error-bars corresponding to the mean and standard deviation of 10 independent runs with different

test splits. The results are compared to, first, the constant prediction, i.e. the hypothesis that there is no change of MMSE within the time interval, and, on the other side, the noise in the data. For the MMSE, [2] reports two noise values : a standard deviation of 1.3 and 2.8 (out of 30) for respectively cognitively normal and MCI patients. Once normalized and converted to absolute values, it corresponds to MAE errors of 0.035 and 0.074, represented by a pale orange interval on Figures 3 and 4.

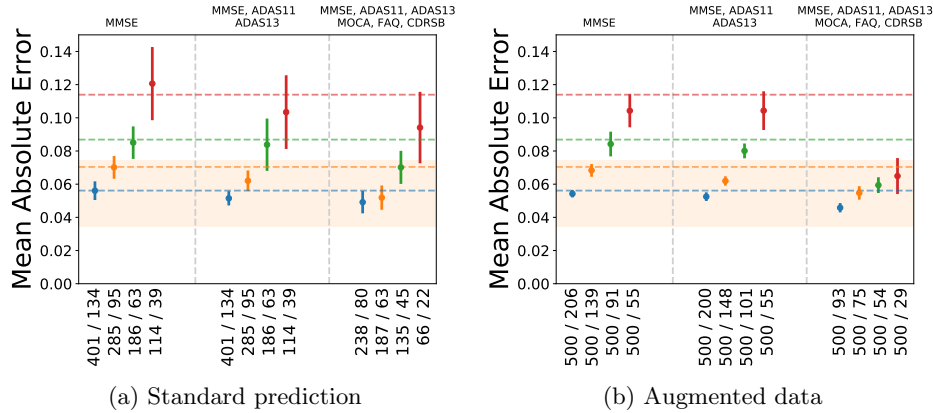


Fig. 3. Prediction of the MMSE in 1 (blue), 2 (orange), 3 (green) and 4 (red) years with different sets of variables (upper part of each column). The colored dashed lines corresponds to the error for the corresponding constant prediction. The pale orange area corresponds to the noise interval in the data. The number at the bottom presents the training and test set sizes.

3.2 Prediction of MMSE

The prediction accuracy in the standard prediction setting, without simulated patients, are presented on Figure 3a. The different columns correspond to different sets of markers used as input. It is possible to reach noise level prediction up to 2 years in advance with the MMSE, ADAS-11, ADAS-13, MOCA, FAQ and CDRSB. At 3 and 4 years, the prediction, although often better than the constant prediction, is still larger than noise level.

Replacing the training set by a simulated cohort leads to a significant improvement of the prediction, as first shown on Figure 4. It corresponds to a decrease of the MAE of 20% (resp. 37%) for prediction 3 years (resp. 4 years) in advance. As the part of the patients used in the estimation set may vary, we tested different scenarios that lead to better results when more patients were used. On the contrary, the number of simulated patients does not seem to have a

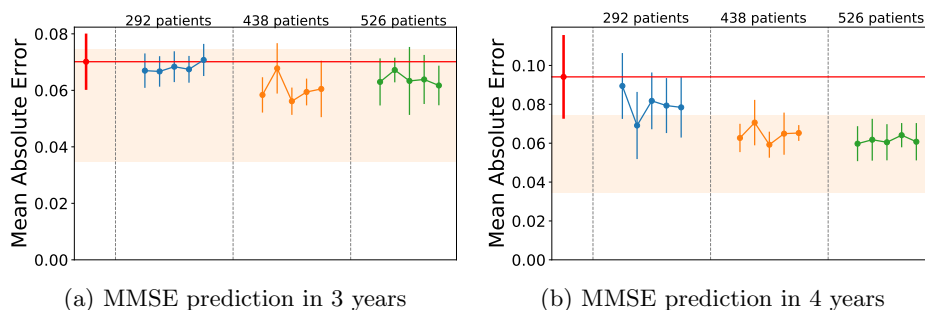


Fig. 4. MMSE prediction based on MMSE, ADAS-11, ADAS-13, MOCA, FAQ and CRDSB. The red value on the left corresponds to the MAE without simulated data. Then, each column corresponds to a different size of the estimation set. Within each column, we simulate, from left to right, 50, 100, 250, 500 and 1000 virtual patients.

great impact on the quality of the prediction. A possible but preliminary explanation lies in the fact that even though there are not a lot of simulated patients, they already incorporate more (simulated) visits than real patients.

3.3 Fair comparison of algorithms

To better exhibit the problem of comparison between different prediction settings, we refer to Figure 3a where the lower figures represent the number of training and test patients in the related model. These numbers decrease for longer time to prediction but also for different sets of features, as not all the examinations have been assessed at each patients' visit.

To solve this issue, we simulated 500 virtual patients for each scenario and estimated the MAE on real patients, as shown on Figure 3b. The prediction at 3 and 4 years on the left column corresponds to the values of Figure 4. The first result to notice, comparatively to 3a, is that the MAE variance over the 10 runs is reduced, probably due to the increased test set. More interestingly, the predictive power of the ADAS-11, ADAS-13 and MMSE is not better than with the MMSE alone, a result that could not have been stated from the standard prediction. It essentially means that the MMSE alone is a predictor as good as the three variables but needs more patients to train the model on. In the same spirit, FAQ, MOCA and/or CDRSB provide substantial information that allow to reach noise level prediction up to 4 years in advance.

4 Discussion

We proposed a data augmentation technique for small data sets that allow to increase the accuracy of the MMSE prediction for MCI subjects. We believe

this technique to be a milestone in the ability to accurately compare various algorithms and features for different time to prediction, as it helps simulating training cohort that are comparable in terms of number of subjects, number of visits per subject and overall follow-up duration.

In this regard, we need to further evaluate the simulation procedure by measuring, for instance, the impact of the number of visits simulated, the time-interval between them, or the selection of the first visit. This could benefit other studies by providing a more accurate comparison of the predictive quality of models or new biomarkers. Overall, it gives a better idea of the generalisation errors of such predictive algorithms in a real clinical setting.

This work has been partly funded by ERC grant N°678304, H2020 EU grant N°666992, and ANR grant ANR-10-IAIHU-06.

References

1. Allasonniere, S., Kuhn, E.: Convergent stochastic expectation maximization algorithm with efficient sampling in high dimension. application to deformable template model estimation. *Computational Statistics & Data Analysis* **91**, 4–19 (2015)
2. Clark, C.M., Sheppard, L., Fillenbaum, G.G., Galasko, D., Morris, J.C., Koss, E., Mohs, R., Heyman, A.: Variability in annual mini-mental state examination score in patients with probable alzheimer disease: a clinical perspective of data from the consortium to establish a registry for alzheimer’s disease. *Archives of neurology* **56**(7), 857–862 (1999)
3. Cui, X., Goel, V., Kingsbury, B.: Data augmentation for deep neural network acoustic modeling. *IEEE/ACM Transactions on Audio, Speech and Language Processing (TASLP)* **23**(9), 1469–1477 (2015)
4. Delyon, B., Lavielle, M., Moulines, E.: Convergence of a stochastic approximation version of the em algorithm. *Annals of statistics* pp. 94–128 (1999)
5. Le Guennec, A., Malinowski, S., Tavenard, R.: Data augmentation for time series classification using convolutional neural networks. In: *ECML/PKDD workshop on advanced analytics and learning on temporal data* (2016)
6. Lu, X., Zheng, B., Velivelli, A., Zhai, C.: Enhancing text categorization with semantic-enriched representation and training data augmentation. *Journal of the American Medical Informatics Association* **13**(5), 526–535 (2006)
7. Perez, L., Wang, J.: The effectiveness of data augmentation in image classification using deep learning. *arXiv preprint arXiv:1712.04621* (2017)
8. Petersen, K.B., Pedersen, M.S., et al.: *The matrix cookbook*. Technical University of Denmark **7**(15), 40–41 (2008)
9. Schiratti, J.B., Allasonnière, S., Colliot, O., Durrleman, S.: A bayesian mixed-effects model to learn trajectories of changes from repeated manifold-valued observations. *The Journal of Machine Learning Research* **18**(1), 4840–4872 (2017)
10. Yu, L., Zhang, W., Wang, J., Yu, Y.: Seqgan: Sequence generative adversarial nets with policy gradient. In: *Thirty-First AAAI Conference on Artificial Intelligence* (2017)
11. Zhu, C., Byrd, R.H., Lu, P., Nocedal, J.: Algorithm 778: L-bfgs-b: Fortran subroutines for large-scale bound-constrained optimization. *ACM Transactions on Mathematical Software (TOMS)* **23**(4), 550–560 (1997)