

Meta-Weighted Gaussian Process Experts for Personalized Forecasting of AD Cognitive Changes

Ognjen (Oggi) Rudovic¹

ORUDOVIC@MIT.EDU

Yuria Utsumi¹

YUTSUMI@MIT.EDU

Ricardo Guerrero²

REG09@IMPERIAL.AC.UK

Kelly Peterson¹

KELLYPET@MIT.EDU

Daniel Rueckert²

D.RUECKERT@IMPERIAL.AC.UK

Rosalind W. Picard¹

ROZ@MIT.EDU

¹*MIT Media Lab, Massachusetts Institute of Technology, USA*

²*Department of Computing, Imperial College London, UK*

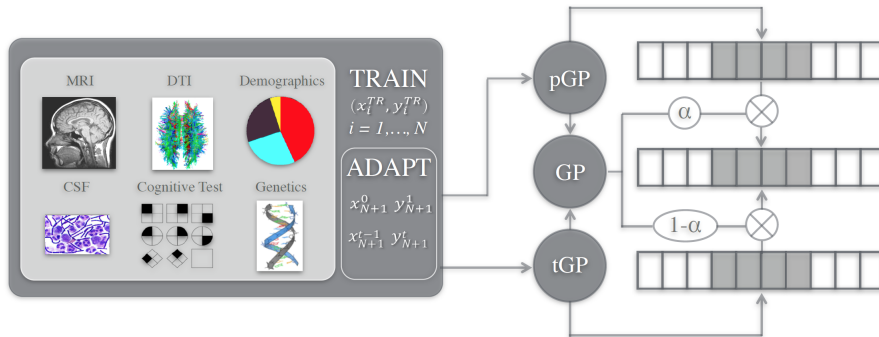
Abstract

We introduce a novel personalized Gaussian Process Experts (pGPE) model for predicting per-subject ADAS-Cog13 cognitive scores – a significant predictor of Alzheimer’s Disease (AD) in the cognitive domain – over the future 6, 12, 18, and 24 months. We start by training a population-level model using multi-modal data from previously seen subjects using a base Gaussian Process (GP) regression. Then, we personalize this model by adapting the base GP sequentially over time to a new (target) subject using domain adaptive GPs, and also by training subject-specific GP. While we show that these models achieve improved performance when selectively applied to the forecasting task (one performs better than the other on different subjects/visits), the average performance per model is suboptimal. To this end, we used the notion of meta learning in the proposed pGPE to design a regression-based weighting of these expert models, where the expert weights are optimized for each subject and his/her future visit. The results on a cohort of subjects from the ADNI dataset show that this newly introduced personalized weighting of the expert models leads to large improvements in accurately forecasting future ADAS-Cog13 scores and their fine-grained changes associated with the AD progression. This approach has potential to help identify at-risk patients early and improve the construction of clinical trials for AD.

1. Introduction

Alzheimer’s disease (AD) is the most common form of dementia, usually associated with the elderly population (over 65 years of age). AD had a worldwide prevalence of around 33.9 million cases by 2011, with predictions suggesting an increase to about 100 million by 2050 (Barnes and Yaffe, 2011). Recently, the view on AD diagnosis has shifted towards a more dynamic process in which clinical and pathological markers evolve gradually before diagnostic criteria are met. The AD Assessment Scale-cognition sub-scale (ADAS-Cog) (Mohs et al., 1997) is the most widely used general cognitive measure in clinical trials of AD (Skinner

Figure 1: **Personalized GP Experts (pGPE)**: The population model is first trained using all visits data of N training subjects (x^{TR} are the input features and y^{TR} are the corresponding ADAS-Cog13 scores), where the time difference between two visits is 6 months. The model personalization to the target (previously unseen) subject ($N + 1$) is then achieved by sequentially adapting the model predictions of the future ADAS-Cog13 scores $y_{t+1:t+4}$ (using the posterior distribution of GPs - f_{GP}), informed by the visits data up to time step t , in the personalized GP (pGP) model. We also train the subject-specific model - target GP (tGP) - using only the past data of the target subject. The shaded fields in the output vector represent the time points for which we aim to predict ADAS-Cog13 scores. The proposed pGPE performs a regression-based weighting of the experts' pGP and tGP using meta-weights (α), estimated using another GP regression model (GP).



et al., 2012). While it was developed as an outcome measure for dementia interventions, its primary purpose was to be an index of global cognition in response to antidementia therapies. The ADAS-Cog assesses multiple cognitive domains including memory, language, praxis, and orientation (Skinner et al., 2012). Because ADAS-Cog has proven important for target clinical assessments, in this paper we focus on a machine learning method that can successfully forecast the future values of this score for target subjects. Specifically, we use the modified ADAS-Cog 13-item scale (Mohs et al., 1997), which is scored on a 85 point scale, where higher scores indicate greater severity.

Clinical Significance One of the main challenges in the clinical assessment of subjects at risk of developing AD is the ability to accurately predict subjects' future cognitive scores. Such predictions can play an important role in subject selection during the design of AD clinical trials. For this, an automated approach that could forecast future cognitive scores, such as ADAS-Cog13, would be of great value during assessment procedures, and could potentially improve clinical trial design and early detection of at-risk subjects. For example, out of hundreds of clinical trials, costing billions of dollars, fewer than 1% have proceeded to the regulatory approval stage and none have managed to prove a disease-modifying effect (Marinescu et al., 2017; Cummings, 2006). Previously, it has been suggested that the low success rate of pharmaceutical AD clinical trials could be due to the inclusion of study

populations that were too heterogeneous (Falahati et al., 2014). More successes depend on an improved ability to accurately identify subjects at early ages of the disease where treatments are most likely to be effective. Thus, developing models with improved ability to automatically predict subjects’ future AD-related metrics indicating disease progression – and to do so as early as possible, especially before the emergence of clinical symptoms – is an important step towards this end. Furthermore, accurate prediction of symptom onset in the time window of 6 to 24 months is critical to participant selection and formation of clinical trials. Thus, having access to accurate future estimates of the progression of cognitive scores such as ADAS-Cog13 within this time frame is of great importance.

Technical Significance In this paper, data from the Alzheimer’s Disease Neuroimaging Initiative (ADNI) (Weiner et al., 2017) were used, and, specifically, the dataset processed for the TADPOLE Challenge (Marinescu et al., 2017). These data are highly heterogeneous and multi-modal, and include imaging (MRI, PET), cognitive scores, CSF biomarkers, genetics, and demographics (e.g. age, gender, race). Although the heterogeneous nature of this dataset lends itself to building powerful, informative multi-modal models, the dataset itself is very sparse, with different combinations of features missing for different subjects. Partial records are also missing for roughly 80% of the subjects (Campos et al., 2015). Moreover, given the wide variability in available data per subject, inherent per-person differences, and the slowly changing nature of the disease, accurate forecasting of cognitive decline and related measures of disease progression is a significant and difficult challenge. To tackle these challenges, we focus on machine learning models that can easily adapt to each subject when forecasting his/her cognitive scores. More specifically, we investigate the effects of the model personalization for the forecasting task using the framework of Gaussian Processes (GP) (Rasmussen and Williams, 2006). This non-parametric probabilistic framework offers great modeling flexibility for not only predicting future cognitive scores but also uncertainty in their predictions. This, in turn, provides a principled approach for model adaptation to target subjects via posterior distribution of GP (Peterson et al., 2017). Moreover, GP are well-suited for forecasting tasks where data are noisy time series with many missing values (Futoma, 2018), as in ADNI.

To this end, we used a cohort of subjects from the ADNI dataset (see Sec. 4) to predict future ADAS-Cog13 scores of target subjects using data from each subject’s previous visits. A subject’s visit is defined as data collected at a single time point (a subject visit) during ADNI. We start by training a population-level model using multi-modal data of previously-seen (source) subjects using the base GP regression (sGP). Then, this model is adapted sequentially over time to a new (target) subject using the notion of domain adaptive GPs (Liu and Vasconcelos, 2015; Eleftheriadis et al., 2017). We extend this personalization approach, denoted as pGP (Peterson et al., 2017), for the forecasting task of predicting the ADAS-Cog13 scores at 6, 12, 18 and 24 months in the future. We also compare this approach with the subject-specific GP model, trained only on the data of the target subject (tGP). While we show that these models achieve good performance when selectively applied to the forecasting task (one performs better than the other for certain subjects/visits), their individual (average) performance is suboptimal. For this, we propose a novel weighting scheme for the GP experts based on the notion of meta-learning (Vanschoren, 2018). The main idea of meta-learning is to observe how different machine learning approaches perform on a wide

range of learning tasks, and then learn from this experience, or meta-data, to learn new tasks more effectively. Following this approach, we designed a regression-based weighting of these expert models in the proposed personalized GP experts (pGPE) model, where the expert weights are optimized for each subject and his/her future visit using data-derived meta-features (see Sec. 3.3). Our results show that this newly introduced personalized weighting of the expert models leads to large improvements in accurately forecasting future ADAS-Cog13 scores for each target subject, while also outperforming traditional weighting schemes for expert models (e.g. by simply averaging the model predictions or using the GP-variance-based weighting), and the standard models for time-series forecasting tasks.

2. Related Work

We briefly review here the related work on forecasting of cognitive scores and clinical status for AD assessment, with the focus on the ADNI dataset. Most existing approaches (e.g., (Schmidt-Richberg et al., 2016; Guerrero et al., 2016; Gavidia-Bovadilla et al., 2017)) focus on modeling subjects based on their clinical status (CS): cognitively normal (CN), mild cognitive impairment (MCI) and Alzheimer’s Disease (AD). Additionally, the majority of these model biomarkers at the population level; for instance, estimating typical trajectories of markers over the full course of the disease to estimate current disease progress and progression rate (Schmidt-Richberg et al., 2016; Guerrero et al., 2016). Guerrero et al. (Guerrero et al., 2016) used mixed effects modeling to derive global and individual biomarker trajectories for a training population, which was later used to instantiate subject-specific models for unseen subjects. Some of the modeling techniques (Guerrero et al., 2016; Schmidt-Richberg et al., 2015, 2016) require cohorts with known disease onset and are prone to bias due to the uncertainty of the conversion time.

When it comes to forecasting clinical status, several authors attempted predicting target scores for longer time windows. For example, several recent works have explored the use of multiple, multi-modal predictors in combination with various machine learning techniques (e.g. SVMs, neural nets, GPs, etc.) to predict conversion from MCI to AD for various future time periods, e.g. from 1-5 years after baseline assessment (Long et al., 2017; Minhas et al., 2017). Likewise, the BrainAGE framework has been proposed for predicting MCI-AD conversion within 3 years of follow-up (Gaser et al., 2013). In addition, to predict conversion within time windows of up to 2 years (short-term converter) and 2-4 years (long-term converter), (Pereira et al., 2017) proposed a stepwise learning approach, where the learned model first predicts whether a subject converts to dementia, or remains stable, and then predicts the more likely progression window (short-term or long-term conversion). More recently, (Grassi et al., 2018) evaluated different machine learning models in order to develop an algorithm for a 3-year prediction of conversion to AD in MCI and PreMCI subjects based only on non-invasive and effectively collectible predictors.

Most of these works attempted forecasting the changes in subjects’ CS, which deals with a limited number of future outcomes (i.e. either binary or a three-class). By contrast, we aim to forecast ADAS-Cog13, defined on a more fine-grained scale (85 levels), which poses a more challenging machine learning problem. While recent work investigated forecasting of ADAS-Cog13 (Peterson et al., 2017), along with other cognitive scores from the TADPOLE challenge, using the GP framework, their forecasting horizon was limited as it only predicted

one time step ahead (6 months). In this work, we attempt forecasting of ADAS-Cog13 up to 24 months ahead. While recent work in (Utsumi et al., 2018) has shown that a simple combination of personalized GP models leads to better prediction accuracy of the target cognitive scores, to the best of our knowledge, no prior work has proposed a principled approach for combining this expert knowledge. In this paper, we do so using the notion of meta learning to design an optimal weighting scheme for personalized GP models.

3. Methodology

Notation. We consider a supervised setting, where $\mathbf{X} = \{\mathbf{X}^{(s)}, \mathbf{X}^{(t)}\}$ represent the input features to our model, and superscripts $\mathbf{l} = \{\mathbf{s}, \mathbf{t}\}$ refer to the data of source (i.e., training) (s) and target (t) subjects, respectively. Similarly, the ADAS-Cog13 scores we aim to forecast are stored in $\mathbf{Y} = \{\mathbf{Y}^{(s)}, \mathbf{Y}^{(t)}\}$. More specifically, $\mathbf{X}^{(l)} = \{\mathbf{x}_n^{(l)}\}_{n=1}^{N_l}$, where N_s and N_t are the numbers of the source and target subjects, respectively. Furthermore, $\mathbf{x} = [\mathbf{x}_1, \dots, \mathbf{x}_t, \dots, \mathbf{x}_T]$ are the input features (predictors) for each time step of the subject’s data¹, and $T \leq 13$ is the maximum number of visits per subject over the period of 10 years, which varies largely per subject. Each \mathbf{x}_t is a vector containing the features as $\mathbf{x}_t = [x_t^{m_1}, \dots, x_t^{m_6}, y_t]^T$, where T is the transpose operation, and $m_{i=1, \dots, 6}$ denotes the features from each data modality. We also add the current cognitive score $y_t \in (0-85)$ as a predictor since we found that it is a strong predictor of the future scores. The scores that we forecast are stored in $\mathbf{Y}^{(l)} = \{\mathbf{y}_n^{(l)}\}_{n=1}^{N_l}$, and as $\mathbf{y} = [\mathbf{y}_1, \dots, \mathbf{y}_t, \dots, \mathbf{y}_T]$. Here, each \mathbf{y}_t contains the scores in the forecasting window of four steps ahead, and is given by $\mathbf{y}_t = [y_{t+1}, \dots, y_{t+4}]^T$. Thus, the goal of the forecasting models proposed here is to learn an efficient mapping: $\mathbf{X} \rightarrow \mathbf{Y}$. As many subjects in ADNI missed certain visits and not all biomarkers were recorded at every visit, we fill in subjects’ missing values using their nearest available past visit; however, no data of future visits are used.²

3.1. Population-level GPs

We first build the population-level forecasting model using data from the source subjects $\{\mathbf{X}^{(s)}, \mathbf{Y}^{(s)}\}$. To this end, we use the GP framework (Rasmussen and Williams, 2006) to train the following forecasting function:

$$\mathbf{y}_t^{(s)} = f(\mathbf{x}_t^{(s)}) + \epsilon_t^{(s)}, \quad (1)$$

where $\epsilon_t^{(s)} \sim \mathcal{N}(0, \sigma_s^2)$ is i.i.d. additive Gaussian noise, which variance σ_s^2 is estimated from the training data. In this framework, central to learning of the forecasting function $f(\cdot)$ is the GP prior that is placed over the function space, leading to the marginal likelihood of our data $p(\mathbf{Y}^{(s)} | \mathbf{X}^{(s)}) = \mathcal{N}(\mathbf{Y}^{(s)} | \mathbf{0}, \mathbf{K}^{(s)} + \sigma_s^2 \mathbf{I})$, where the elements of the kernel matrix are given by $\mathbf{K}^{(s)} = \mathbf{k}(\mathbf{X}^{(s)}, \mathbf{X}^{(s)})$ and \mathbf{I} is the identity matrix. The choice of the kernel function $k(\cdot, \cdot)$ is at the heart of GPs as it encodes the relationships in our data. To this end, we use the standard radial basis function (RBF) kernel. Specifically, we investigate two types of the RBF kernel: isotropic (iso) and automatic relevance determination (ard) (Rasmussen

1. The subscript t refers to time and is different from superscript (t) that refers to the target subject.
 2. To tackle this, more advanced approach based on auto-encoders can be used, e.g., see (Campos et al., 2015), and we leave this for the future work.

and Williams, 2006). The latter assigns different weights (length-scales) to each feature (in our case, each data modality), effectively doing feature selection.

The kernel parameters θ are optimized by minimizing the negative log-marginal likelihood: $-\log p(\mathbf{Y}^{(s)}|\mathbf{X}^{(s)}, \theta)$ using conjugate gradient method. Then, given the data from the visit at time t of a *new* subject, $\mathbf{x}_* = \{\mathbf{x}_t, \mathbf{y}_t\}$, the GP predictive distribution provides the mean and variance forecasts of the cognitive scores \mathbf{y}_* as:

$$\boldsymbol{\mu}_*^{(s)} = \mathbf{k}_*^T (\mathbf{K}^{(s)} + \sigma_s^2 \mathbf{I})^{-1} \mathbf{Y}^{(s)} \quad (2)$$

$$\mathbf{V}_*^{(s)} = k_{**} - \mathbf{k}_*^T (\mathbf{K}^{(s)} + \sigma_s^2 \mathbf{I})^{-1} \mathbf{k}_*, \quad (3)$$

where $\mathbf{k}_* = k(\mathbf{X}^{(s)}, \mathbf{x}_*)$ and $k_{**} = k(\mathbf{x}_*, \mathbf{x}_*)$. We use the mean of this predictive distribution for the point estimate of target outputs, denoted as $\boldsymbol{\mu}_*^{(s)} = \boldsymbol{\mu}_*^{(s)}(\mathbf{x}_*)$ for notational convenience. Note that we use the shared covariance function for simultaneous forecasting of the four future scores in \mathbf{y}_* . Consequently, the model assigns the same variance (\mathbf{V}_*) to the forecasting window. We refer to this setting as the source GP (sGP).

3.2. Personalized GPs (pGP and tGP)

We use the notion of domain adaptive GPs (DA-GP) (Liu and Vasconcelos, 2015; Eleftheriadis et al., 2017) to personalize the population GP model to target subjects. This is achieved by sequentially adapting the GP posterior for the test subject using the data of his/her past visits, to forecast the future ADAS-Cog13 scores \mathbf{y} . This is achieved by using the obtained posterior distribution of the source (population) data as a GP prior for the GP of the ADAS-Cog13 scores of the target subject at time t , given by $p(\mathbf{y}_t^{(t)}|\mathbf{x}_t^{(t)}, \{\mathbf{X}^{(s)}, \mathbf{Y}^{(s)}\}, \theta)$. Assuming that we have already observed the data of the target subject, up to time $t-1$, we use the data pairs $\{\mathbf{x}_{1:t-1}^{(t)}, \mathbf{y}_{1:t-1}^{(t)}\}$ to correct the posterior distribution for the target subject's data at time t . More formally, by applying Eqs. (2-3) to $\{\mathbf{x}_{1:t-1}^{(t)}, \mathbf{y}_{1:t-1}^{(t)}\}$, we obtain the conditional prior on the target subject data, $\mathcal{N}(\boldsymbol{\mu}_{1:t-1}^{(t|s)}, \mathbf{V}_{1:t-1}^{(t|s)})$, where $(t|s)$ denotes the conditioning order. For exact derivation, see (Liu and Vasconcelos, 2015). Given this prior and a test input $\mathbf{x}_* = \mathbf{x}_t^{(t)}$, the correct form of the adapted posterior after observing the target subject data at visit t represents the predictive distribution of the personalized GP (pGP):

$$\boldsymbol{\mu}_*^{(p)} = \boldsymbol{\mu}_*^{(s)} + \mathbf{V}_*^{(t|s)T} (\mathbf{V}_{1:t-1}^{(t|s)} + \sigma_s^2 \mathbf{I})^{-1} (\mathbf{Y}_{1:t-1}^{(t)} - \boldsymbol{\mu}_{1:t-1}^{(t|s)}) \quad (4)$$

$$\mathbf{V}_*^{(p)} = \mathbf{V}_*^{(s)} - \mathbf{V}_*^{(t|s)T} (\mathbf{V}_{1:t-1}^{(t|s)} + \sigma_s^2 \mathbf{I})^{-1} \mathbf{V}_*^{(t|s)}, \quad (5)$$

where $\mathbf{V}_*^{(t|s)} = k(\mathbf{x}_{1:t-1}^{(t)}, \mathbf{x}_*) - k(\mathbf{X}^{(s)}, \mathbf{x}_{1:t-1}^{(t)})^T (\mathbf{K}^{(s)} + \sigma_s^2 \mathbf{I})^{-1} k(\mathbf{X}^{(s)}, \mathbf{x}_*)$. Eqs. (4-5) show that final forecast by the pGP is the combination of the population-model forecast, plus a correction term.³ The latter shifts the mean toward the distribution of the target subject and improves the model's confidence by reducing its predictive variance as more data of the target subject is observed. The inference in pGP is efficient as it uses the kernel parameters of the sGP, thus, no further training is employed.

3. For $t=1$, the population model is used instead, and from $t=T-4:T$ the future target scores are repeated to fill in the window of target scores.

We also form the target-subject-specific GP (tGP) model using only the observed data of the target subject up to time t , i.e., the same data we used for the adaption in pGP. However, since this data set is of insufficient size to train the GP parameters – the GP would easily overfit – we use the kernel parameters of sGP. Unlike sGP, which has a fixed covariance matrix (trained using the source subjects), tGP continually updates its covariance matrix as more past data (up to t) of a target subject become available. Nevertheless, the inference procedure is the same, and for $\mathbf{x}_* = \mathbf{x}_t^{(t)}$, the predictive distribution of tGP is given by:

$$\boldsymbol{\mu}_*^{(t)} = \mathbf{k}_*^T (\mathbf{K}_{1:t-1}^{(t)} + \sigma_s^2 \mathbf{I})^{-1} \mathbf{y}_{1:t-1}^{(t)} \quad (6)$$

$$\mathbf{V}_*^{(t)} = k_{**} - \mathbf{k}_*^T (\mathbf{K}_{1:t-1}^{(t)} + \sigma_s^2 \mathbf{I})^{-1} \mathbf{k}_*. \quad (7)$$

where $\mathbf{k}_* = k(\mathbf{x}_{1:t-1}^{(t)}, \mathbf{x}_*)$, $k_{**} = k(\mathbf{x}_*, \mathbf{x}_*)$, and $\mathbf{K}_{1:t-1}^{(t)} = k(\mathbf{x}_{1:t-1}^{(t)}, \mathbf{x}_{1:t-1}^{(t)})$. Note that the kernel matrix here increases in size after each visit of a target subject.

3.3. Personalized GP Experts (pGPE)

After obtaining the pGP and tGP expert models, an optimal weighting scheme is learned to combine these two models within the proposed pGPE approach. We introduce a novel weighting scheme based on meta-learning that assigns the optimal weights to the expert models for each subject and his/her visit. More formally, recall that the predictive distribution of each expert’s forecasts for the time window ahead is given by:

$$\text{pGP: } \mathbf{y}^{(p)} \sim \mathcal{N}(\boldsymbol{\mu}^{(p)}, \mathbf{V}^{(p)}), \quad \text{and tGP: } \mathbf{y}^{(t)} \sim \mathcal{N}(\boldsymbol{\mu}^{(t)}, \mathbf{V}^{(t)}). \quad (8)$$

Using the expert distributions, we obtain the forecasts $\mathbf{y}^{(g)}$ by the pGPE model as:

$$\mathbf{y}^{(g)} = \alpha \mathbf{y}^{(p)} + (1 - \alpha) \mathbf{y}^{(t)} \sim \mathcal{N}(\alpha \boldsymbol{\mu}^{(p)} + (1 - \alpha) \boldsymbol{\mu}^{(t)}, \alpha^2 \mathbf{V}^{(p)} + (1 - \alpha)^2 \mathbf{V}^{(t)}) \quad (9)$$

In this work, we consider only the point predictions, defined as the mean of the predictive distribution defined above. To find an optimal weight α for each subject and his/her visit, we solve the following optimization problem:

$$\boldsymbol{\alpha}^{opt} = \arg \min_{\alpha_i} \frac{1}{4N} \sum_{i=1}^N \sum_{k=1}^4 \left\| \mathbf{y}_{i,k}^{(g)} - \mathbf{y}_{i,k} \right\|^2, \quad (10)$$

where $\boldsymbol{\alpha}^{opt} = [\alpha_1^{opt}, \dots, \alpha_N^{opt}]$, and $N = N_s \times \mathbf{T}$ is the total number of training samples from the training subjects and their visits. Since we seek to optimize one α for each data sample, Eq.10 is a convex optimization problem with a closed form solution given by:

$$\alpha_i^{opt} = \sum_{k=1}^4 \frac{\mathbf{y}_{i,k}^{(g)} - \mathbf{y}_{i,k}}{\mathbf{y}_{i,k}^{(p)} - \mathbf{y}_{i,k}^{(t)}}, \quad i = 1 \dots N. \quad (11)$$

These optimal weights are used for further learning in the proposed approach. While we are able to learn these weights for the training data, it is unclear how to generate those weights for the target subjects and their visits since the future scores are unknown. To solve this, we first derive the following feature vector for each visit of the training subjects:

$$\mathbf{m}_t = [\mathbf{y}_t^{(p)} \mathbf{y}_t^{(t)} y_t], \quad \dim(\mathbf{m}_t) = 1 \times 9, \quad (12)$$

where \mathbf{m}_t contains the four predictions per expert, pGP ($\mathbf{y}^{(p)}$) and tGP ($\mathbf{y}^{(t)}$), obtained by evaluating the pre-trained expert models on the data of the training subjects’ visits $t = 1, \dots, \mathbf{T}$. We also include their ground truth score for the current visit, \mathbf{y}_t . However, we do not use the data of the target subjects. Finally, the expert weights for the new data point at which we aim to forecast the future ADAS-Cog13 are obtained as:

$$\alpha_* = f(\mathbf{m}_*) + \epsilon_*, \quad (13)$$

where the function $f(\cdot)$ is trained using a GP regression model with the RBF-ard kernel function and training data pairs $\{\mathbf{M}, \boldsymbol{\alpha}^{opt}\}$, where $\mathbf{M} = \{\mathbf{m}_{i=1, \dots, N}\}^T$. Note that the feature vector was centered by subtracting the mean of its elements. Also, we applied the quadratic expansion of such feature vector followed by the element-wise normalization with the norm of the expanded feature vector, resulting in a 55-dim feature vectors. This allowed the GP to capture the underlying relationships between these meta-features and optimal weights for the target experts. We also experimented with other feature types (e.g., using the input features, the prediction uncertainty, and the optimal weights from previous visits of the target subject), but the ones reported here achieved the best performance in the task.

4. Results

Data used in this paper come from the ADNI database (adni.loni.usc.edu). We downloaded the standard dataset processed for the TADPOLE Challenge (Marinescu et al., 2017); this dataset represents 1,737 unique subjects and was created from the **ADNIMERGE** spreadsheet, to which regional MRI (volumes, cortical thickness, surface area), PET (FDG, AV45, AV1451), DTI (regional means of standard indices) and cerebrospinal fluid (CSF) biomarkers were added. From this data, we excluded: the PET data, because of their sparsity, the cognitive scores other than ADAS-Cog13, as well as the clinical status (CS) and normalized ventricle volumes (ICVn)⁴, to construct a multi-modal feature set containing: demographics (6 features), genetics (3 features), CSF (3 features), MRI (365 features), and DTI (229 features), thus 606 input features in total. Furthermore, we selected a cohort of subjects with at least 10 visits, and whose records were not missing more than 82.5% of the subject data (from all modalities/visits), resulting in 100 subjects, 48 of whom were diagnosed with AD. We performed a 10-fold subject-independent cross-validation (i.e., each fold contained data of 10 subjects).⁵ The input features were z-normalized (zero mean, unit variance). To measure the models’ performance, we report $\mu \pm \sigma$ of mean absolute error (MAE) for the 10-folds. To form the forecasting window of four consecutive visits (6, 12, 18 and 24 months), we imputed the ADAS-Cog13 scores for the missing visits using the scores of the past visits. However, the reported evaluation metrics were computed using the scores for the existing visits only. We evaluated three types of GP models: (i) population-level (sGP), (ii) personalized (pGP and tGP), and (iii) personalized experts (pGPE). The settings for these models are as follows:

4. In the TADPOLE Challenge, ADAS-Cog13, CS and ICVn are treated as target outputs, so in this work we excluded the latter two from the predictors in our model.

5. The models were trained using data from 9 folds, and tested on data of the remaining fold. This was repeated for each of 10 folds.

Table 1: Comparison of different forecasting approaches. All the differences between pGPE and the compared models are statistically significant ($p < 0.001$) based on a paired t-test with equal variances. The last row shows a lower bound on MAE, obtained when the ADAS-Cog13 scores from the *target* subjects are used to determine the optimal weights (W_{opt}) for the expert models (Sec. 3.3). In pGPE, the meta-weights (W_{reg}) for target subjects are estimated with GP tuned on *training* subjects.

Models	MAE				
	$t + 1$	$t + 2$	$t + 3$	$t + 4$	Avg.
Base (y_t)	3.74±0.17	3.84±0.17	4.72±0.29	5.02±0.33	4.34
LassoR	3.76±0.18	3.85±0.16	4.61±0.26	4.51±0.27	4.20
SVR(RBF-iso)	3.97±0.41	4.49±0.57	4.88±0.49	5.21±0.66	4.63
LSTM	4.87±0.31	5.44±0.36	5.52±0.42	6.48±0.53	5.58
sGP(VIS)	3.62±0.33	3.67±0.31	4.37±0.54	4.41±0.54	4.02
sGP(RBF-iso)	3.86±0.37	4.39±0.52	4.72±0.44	5.04±0.62	4.50
sGP(RBF-ard)	3.50±0.30	3.55±0.38	4.11±0.61	4.27±0.77	3.86
pGP(RBF-ard)	3.50±0.31	3.49±0.35	4.00±0.52	4.04±0.63	3.76
tGP(RBF-ard)	3.70±0.37	4.07±0.54	3.90±0.42	4.38±0.79	4.01
pGPE(W_{prior})	3.56±0.34	3.47±0.40	3.97±0.51	4.04±0.65	3.76
pGPE(W_{freq})	3.70±0.37	3.33±0.34	3.67±0.40	3.82±0.59	3.63
pGPE(W_{ave})	3.47±0.31	3.46±0.38	3.64±0.41	3.85±0.60	3.61
pGPE(W_{var})	3.66±0.36	3.98±0.51	3.85±0.41	4.29±0.75	3.95
pGPE(W_{reg})	2.66±0.28	2.57±0.24	2.48±0.29	2.87±0.32	2.65
pGPE(W_{opt})	1.80±0.28	1.50±0.14	1.49±0.17	1.54±0.37	1.58

- Population-level GPs.** To forecast the future ADAS-Cog13 scores, we trained the sGP models using the input features for the visits up to time t and the corresponding ADAS-Cog13 score (y_t). We employed two kernel functions: sGP(RBF-iso) and sGP(RBF-ard) (Rasmussen and Williams, 2006), where the latter learns a separate length-scale parameter for each data modality used, and y_t , i.e., $\lambda_{i=1:6}$. We also report the performance of the sGP(RBF-iso) model when only y_t from the current visit (VIS) is used as a predictor of the future scores, effectively doing the label smoothing. We denote this model as sGP(VIS).
- Personalized GPs.** We also report the performance by the pGP and tGP models (Sec. 3). The former performs the adaptation of the sGP to the new subject using the subject data up to time t , while the latter builds a new kernel matrix using only the data of the new subject (however, it uses the hyper-parameters of the sGP(RBF-ard) model). We used the sGP(RBF-ard) model described above to derive the pGP and tGP models, since it achieved the best performance among the population-level GPs.
- Personalized GP Experts.** As described in Sec. 3.3, we used the pGP/tGP models as experts in the proposed pGPE approach. Below we describe different weighting schemes, including the proposed meta-weighting, we applied to these experts.
 - W_{prior} : derived by inspecting the *average* performance of the pGP/tGP models and

by applying the best performing model per visit/forecasting step (see Fig.2).

\mathbf{W}_{freq} : derived based on the frequency of pGP outperforming tGP per visit and per subject for each forecasting step. The weights range from 0 to 1, and represent the normalized sum over the columns (subjects) depicted in Fig.3.

\mathbf{W}_{ave} : derived by averaging the forecasts by pGP and tGP, as in (Utsumi et al., 2018).

\mathbf{W}_{var} : the variance-based weighting (Deisenroth and Ng, 2015) of pGP/tGP.

\mathbf{W}_{reg} : the proposed regression-based meta-weighting of pGP/tGP (Sec. 3.3)

\mathbf{W}_{opt} : the optimal weights obtained by applying Eq. (11) to the target data.

We also include the baseline comparisons with non-GP models, commonly used for forecasting tasks, and usually applied to the ADNI dataset. Specifically, we compare to Lasso Regression (LassoR), Support Vector Regression (SVR) with RBF(iso), and Long Short-Term Memory (LSTM) (Hochreiter and Schmidhuber, 1997). In relation to the GP models evaluated, LassoR can be seen as our sGP with linear kernel and l_1 regularization, and SVR as a sparse version of our sGP. However, SVR does not provide probabilistic outputs, and, thus, it lacks a principled way for model adaptation, as done in our personalized GP models via their posterior distribution. The optimal regularization parameter and kernel width of LassoR and SVR, respectively, were selected on the held-out subjects from the training set (10/90). As input to LSTM, we tried a sliding window of the size $n = 1, \dots, 5$, containing observations from the past visits. We used again the held-out dataset to select the optimal window size ($n = 2$), and the number of the LSTM output states ($h = 64$). This was followed by a fully-connected linear layer (64×4). During the model training, we applied a dropout (0.3) to prevent the model overfitting. Lastly, we report the results obtained by a model-free approach, denoted as Base (y_t), where the score from the current visit (y_t) was used as a forecast of the four steps ahead.

Table 1 compares the evaluation metrics for the model settings described above. By comparing different population-level GP models, we note that simple smoothing of the previous ADAS-Cog13 score (y_t) using sGP(VIS) is quite effective. We attribute this to the fact that the target scores from the consequent visits are highly correlated. Compared to the model free approach, Base (y_t), which forecasts by "carrying forward" the current score, we note that sGP(VIS) takes advantage of the GP's smoothing property, effectively leveraging the labels (i.e. ADAS-Cog13 scores) from the training population. This results in average error reduction of 0.32 over the four time steps. However, the data-informed sGP(RBF-ard) model (that performs the smoothing of the input features) achieved the best performance overall. By learning the separate length-scale parameters for each modality, this model was able to successfully combine different data types - something that cannot be achieved with RBF-iso, resulting in the adverse performance by the sGP(RBF-iso) model. Next, we compare the personalized models (pGP and tGP) that use the sGP(RBF-ard) as the base model. It is evident that using the domain adaptive GPs to personalize the population-level model to each subject (pGP) helps to reduce MAE, the difference being more pronounced in later time steps. On the other hand, the tGP model outperforms the population-level sGP/pGP only at step $t + 3$. Yet, as can be observed from Fig. 2, tGP outperforms these models at certain visits, even though this is not reflected in its MAE.

The bottom half of Table 1 shows results obtained from different weighting schemes that we devised to compare with our regression-based pGPE approach (pGPE(\mathbf{W}_{reg})). First, note that the weighting of the pGP/tGP models based on the heuristics derived by looking at

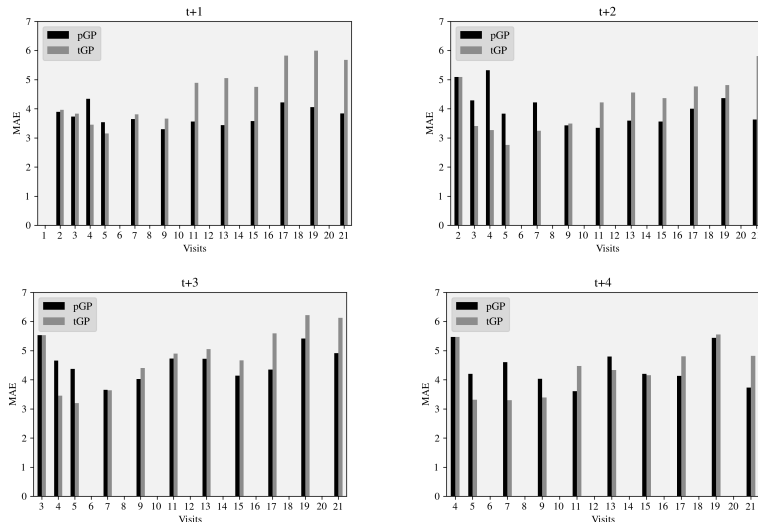


Figure 2: Average MAE for the pGP and tGP models computed per visit, showing specific time-steps where one of the expert models dominates over the other. This is used to derive prior weights (W_{prior}) for the pGPE approach by selecting the optimal model (on average) for each visit and across the four forecasting time steps. Note that after the fifth visit, the subjects in ADNI were seen on a yearly basis.

the models' scores per-visit/subject leads to similar or better performance compared to the pGP model alone. Specifically, frequency-based weighting (W_{freq}) is slightly more effective than W_{prior} as it is not biased by the amplitude of the per-subject MAEs when deriving the weights. This is because some subjects had much higher MAE than the others, which, evidently, did not translate well to test subjects when applying this binary model selection encoded by W_{prior} . Interestingly, averaging the two experts (pGP/tGP) via W_{ave} further reduces MAE. As this model outperforms both the pGP and tGP alone, this also signals that the two models have very different performance per subject/visit, also evidenced by Figs. 2&3. Thus, by simply averaging their predictions, this model reduces the prediction biases of the pGP/tGP models. We also explored the predictive variance of the pGP/tGP experts to assign different weights to the models based on their confidence in the target predictions (Deisenroth and Ng, 2015). While this approach improved the performance of tGP alone, it hurt the overall performance of pGP. By looking into the estimated variances, we found that tGP tended to overestimate the prediction variance, due to the small amount of data (only the target subject visits' data were used to construct the kernel).

The proposed regression-based weighting outperforms the compared weighting schemes by large margin, as can be seen from the per-time-step and overall results. This is consistent for the whole forecasting horizon. Note that, apart from the variance-based weighting, the other weighting approaches apply the same and fixed expert weights across all subject. This is suboptimal. The benefit of learning the subject- and visit-specific weights, as introduced here, is clear from the results obtained. This further shows that the optimal weights vary largely across subjects/visits (see Fig. 3); however, the proposed meta-weighting using GP

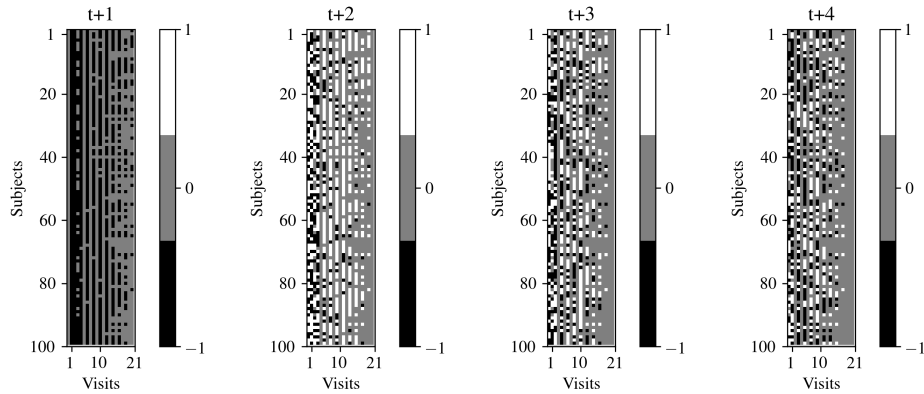


Figure 3: Plots depicting the best performing personalized model for each subject and his/her visit. For pGP, this is shown in white (+1), and for tGP this is shown in black (-1). When there is no visit data available, this is shown in gray (0). Note that the pGP always outperforms tGP at $t + 1$, while for the later time steps, the models exhibit largely heterogeneous performance.

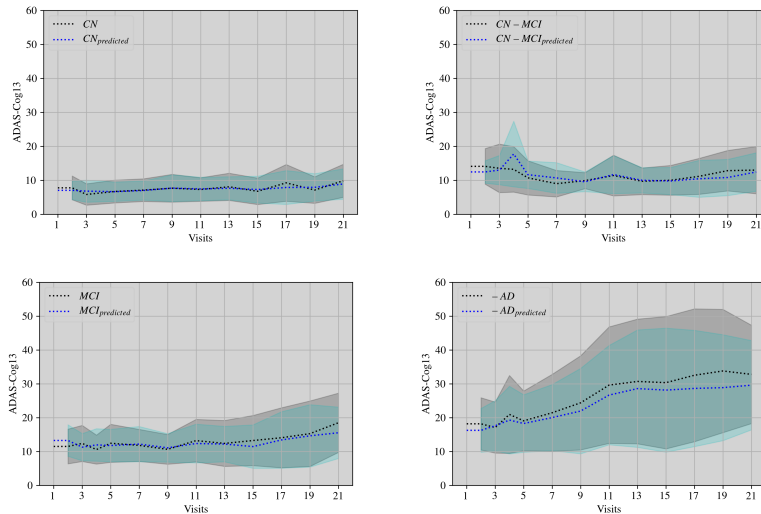


Figure 4: The graphs showing the ADAS-Cog13 scores for each visit (averaged across subjects) by the proposed regression-based pGP expert model ($\text{pGPE}(W_{reg})$) - predicted- and ground-truth scores. Each graph shows these scores along with their standard deviation across the subjects, where the subjects were separated based on their clinical status: cognitively normal (CN), those who converted at one point from CN to mild-cognitive impairment (MCI) (CN→MCI), those who stayed at MCI, and those who were diagnosed initially or converted to AD.

was able to effectively learn those changes. To confirm this interpretation, we applied random weights sampled from the learned GP. This resulted in performance significantly lower than that obtained by the population-based models. We also included the results obtained with the optimal weights (W_{opt} – see Sec. 3.3) that show the lower-bound on MAE that can be achieved with the proposed pGPE approach. By looking at the performance achieved by the baselines, we note that LassoR, typically used in clinical settings because of its interpretability, outperforms the GP and SVR population models with non-linear mappings (RBF-iso). This evidences that a simple linear model can be useful in this context, however, like SVR that is commonly used in prediction tasks in ADNI, it does not provide a probabilistic output. On the other hand, due to the probabilistic nature of GPs, pGP can easily adapt to the target subjects, which result in lower average MAE. While LSTMs have shown great performance in various clinical time series forecasting tasks, in our experiments on the ADNI data, it underperformed. We attribute this to the highly heterogenous and irregularly sampled data, which, evidently, negatively affects the ability of LSTM to leverage its memory properties.

Fig. 4 shows the graphs depicting the predicted ADAS-Cog13 scores for each visit (averaged across subjects) by the proposed pGPE(W_{reg}) that we plot against the ground-truth scores. Each graph shows these scores along with their one standard deviation across the subjects, where the subjects were separated based on their clinical status (not used as input in the models): 25 cognitively normal (CN), 20 who converted at one point from CN to mild-cognitive impairment (MCI) (CN→MCI), 7 who stayed at MCI, and 48 who were diagnosed initially or converted to Alzheimer’s Disease (AD) in one of their visits. Note that the proposed approach can accurately forecast the trend in changes of the target scores. As can be seen, the subjects who were diagnosed with AD or converted to AD at one of the visits have a large deviation from the ADAS-Cog13 scores for the other three types of clinical status, with this difference being more pronounced towards later visits, as expected. While we show here that the proposed approach matches well the distributions of the cognitive scores at the sub-group level, the proposed approach can be used to predict the trend of the cognitive changes for each target subject, and associate it with future changes in his/her clinical status (e.g. to detect subjects that will rapidly progress over a period of 2 years). We plan to investigate the feasibility of such approach in our future work.

5. Conclusions

Accurate forecasting of changes in the biomarkers of AD is challenging, mainly because of the highly heterogeneous, noisy and missing nature of the clinical data available for this task, but also due to the highly pronounced individual differences in such data. As a step forward, this work developed a novel personalized approach for accurate forecasting of ADAS-Cog13, a significant biomarker of AD, up to two years in the future. The proposed personalized GP framework for time series forecasting uses the notion of meta-learning to determine an optimal weighting scheme for individual expert models based on personalized GP models. We showed on a sub-cohort of ADNI subjects that this approach brings significant improvements over population-level GP models. Furthermore, we showed that the proposed personalized experts alone cannot generalize well enough across all the subjects and their visits. By contrast, the introduced personalized meta-weighted GP framework enables accurate fore-

casting of changes in ADAS-Cog13, outperforming the existing personalized GP models, and traditional models commonly used for forecasting of time-series clinical data. In the future, we plan to extend our meta-weighted GPE framework so that it can handle more than two experts and forecast multiple biomarkers of AD simultaneously. For this, we plan to investigate the notion of multi-task GPs and active learning to select the optimal experts for each subject/visit. To evaluate our models on larger cohort of ADNI participants, we plan to employ techniques for filling in the missing data, as the existing data is highly sparse.

Acknowledgements

The work of O. Rudovic was funded by EC H2020 Marie Curie Action – Individual Fellowship no. 701236 (EngageMe).

References

- D.E. Barnes and K. Yaffe. The projected effect of risk factor reduction on alzheimer’s disease prevalence. *The Lancet Neurology*, 10(9):819 – 828, 2011. ISSN 1474-4422. doi: [https://doi.org/10.1016/S1474-4422\(11\)70072-2](https://doi.org/10.1016/S1474-4422(11)70072-2).
- S. Campos, L. Pizarro, C. Valle, K.R. Gray, D. Rueckert, and H. Allende. Evaluating imputation techniques for missing data in adni: a patient classification study. In *Iberoamerican Congress on Pattern Recognition*, pages 3–10. Springer, 2015.
- J.L. Cummings. Challenges to demonstrating disease-modifying effects in alzheimer’s disease clinical trials. *Alzheimer’s & Dementia*, 2(4):263–271, 2006.
- M. Deisenroth and J.W. Ng. Distributed gaussian processes. In *Proceedings of the 32nd International Conference on International Conference on Machine Learning-Volume 37*, pages 1481–1490, 2015.
- S. Eleftheriadis, O. Rudovic, M.P. Deisenroth, and M. Pantic. Gaussian process domain experts for modeling of facial affect. *IEEE Transactions on Image Processing*, 26(10):4697–4711, 2017.
- F. Falahati, E. Westman, and A. Simmons. Multivariate Data Analysis and Machine Learning in Alzheimer’s Disease with a Focus on Structural Magnetic Resonance Imaging. *Journal of Alzheimer’s Disease*, 41(3):685–708, jul 2014. ISSN 18758908. doi: 10.3233/JAD-131928.
- J. Futoma. *Gaussian process-based models for clinical time series in healthcare*. PhD thesis, Duke University, 2018.
- C. Gaser, K. Franke, S. Klöppel, N. Koutsouleris, H. Sauer, ADNI, et al. Brainage in mild cognitive impaired patients: predicting the conversion to alzheimer’s disease. *PloS one*, 8(6), 2013.
- G. Gavidia-Bovadilla, S. Kanaan-Izquierdo, M. Mataró-Serrat, A. Perera-Lluna, ADNI, et al. Early prediction of alzheimer’s disease using null longitudinal model-based classifiers. *PloS one*, 12(1): e0168011, 2017.
- M. Grassi, G. Perna, D. Caldirola, et al. A clinically-translatable machine learning algorithm for the prediction of alzheimer’s disease conversion in individuals with mild and premild cognitive impairment. *Journal of Alzheimer’s Disease*, 61(4):1–19, 2018.
- R. Guerrero, A. Schmidt-Richberg, C. Ledig, ADNI, et al. Instantiated mixed effects modeling of alzheimer’s disease markers. *NeuroImage*, 142:113–125, 2016.

- S. Hochreiter and J. Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.
- B. Liu and N. Vasconcelos. Bayesian model adaptation for crowd counts. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 4175–4183, 2015.
- X. Long, L. Chen, C. Jiang, L. Zhang, ADNI, et al. Prediction and classification of alzheimer disease based on quantification of mri deformation. *PloS one*, 12(3):e0173372, 2017.
- R.V. Marinescu, N.P. Oxtoby, A.L. Young, E.E. Bron, A.W. Toga, M.W. Weiner, F. Barkhof, N.C. Fox, S. Klein, and D.C. Alexander. The alzheimer’s disease prediction of longitudinal evolution (tadpole) challenge, May 2017. URL <https://tadpole.grand-challenge.org/home/>.
- S. Minhas, A. Khanum, F. Riaz, A. Alvi, and S.A. Khan. A nonparametric approach for mild cognitive impairment to ad conversion prediction: Results on longitudinal data. *IEEE journal of biomedical and health informatics*, 21(5):1403–1410, 2017.
- R. Mohs, D. Knopman, R. Petersen, S. H Ferris, C. Ernesto, M. Grundman, M. Sano, L. Bieliauskas, D. Geldmacher, C. Clark, et al. Development of cognitive instruments for use in clinical trials of antedementia drugs: additions to the alzheimer’s disease assessment scale that broaden its scope. *Alzheimer disease and associated disorders*, 1997.
- T. Pereira, A. Mendon, F. Ferreira, S. Madeira, M. Guerreiro, et al. Towards a reliable prediction of conversion from mild cognitive impairment to alzheimer’s disease: stepwise learning using time windows. In *Medical Informatics and Healthcare*, pages 19–26, 2017.
- K. Peterson, O. Rudovic, R. Guerrero, and R.W. Picard. Personalized gaussian processes for future prediction of alzheimer’s disease progression. *NIPS Workshop: ML4HC*, 2017.
- C.E. Rasmussen and C.K.I. Williams. *Gaussian processes for machine learning*, volume 1. MIT press Cambridge, MA, 2006.
- A. Schmidt-Richberg, R. Guerrero, C. Ledig, H. Molina-Abril, A.F. Frangi, D. Rueckert, and ADNI. Multi-stage biomarker models for progression estimation in Alzheimer’s disease. In *International conference on Information Processing in Medical Imaging*, volume 9123 of *Lecture Notes in Computer Science*, pages 387–398. Springer International Publishing, 2015.
- A. Schmidt-Richberg, C. Ledig, R. Guerrero, et al. Learning biomarker models for progression estimation of alzheimer’s disease. *PloS one*, 11(4):e0153040, jan 2016.
- J. Skinner, J. Carvalho, G. Potter, ADNI, et al. The alzheimer’s disease assessment scale-cognitive-plus (adas-cog-plus): an expansion of the adas-cog to improve responsiveness in mci. *Brain imaging and behavior*, 6(4):489–501, 2012.
- Y. Utsumi, O. Rudovic, K. Peterson, R. Guerrero, and R.W. Picard. Personalized gaussian processes for forecasting of alzheimer’s disease assessment scale-cognition sub-scale (adas-cog13). In *40th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, pages 4007–4011, 2018.
- J. Vanschoren. Meta-learning: a survey. *arXiv preprint arXiv:1810.03548*, 2018.
- M.W. Weiner, D.P. Veitch, P.S. Aisen, , L.A. Beckett, N.J. Cairns, R.C. Green, D. Harvey, C.R. Jack Jr, W. Jagust, J.C. Morris, et al. Recent publications from the alzheimer’s disease neuroimaging initiative: Reviewing progress toward improved ad clinical trials. *Alzheimer’s & Dementia*, 13(4):e1–e85, 2017.