

Personalized Medication Response Prediction for Attention-Deficit Hyperactivity Disorder: Learning in the Model Space Versus Learning in the Data Space

Hin Kwan Wong¹, Paul A. Tiffin^{2*}, Michael J. Chappell³, Thomas E. Nichols¹, Patrick R. Welsh⁴, Orla M. Doyle⁵, Boryana C. Lopez-Kolkovska¹, Sarah K. Inglis⁶, David Coghil⁷, Yuan Shen⁸, Peter Tiño⁸

¹Warwick Manufacturing Group, University of Warwick, United Kingdom, ²Department of Health Sciences, University of York, United Kingdom, ³School of Engineering, University of Warwick, United Kingdom, ⁴School of Psychology, Newcastle University, United Kingdom, ⁵Centre for Neuroimaging Sciences, King's College London, United Kingdom, ⁶Ninewells Hospital and Medical School, University of Dundee, United Kingdom, ⁷Royal Children's Hospital, Australia, ⁸School of Computer Science, University of Birmingham, United Kingdom

Submitted to Journal:
Frontiers in Physiology

Specialty Section:
Computational Physiology and Medicine

Article type:
Methods Article

Manuscript ID:
243114

Received on:
16 Nov 2016

Revised on:
17 Mar 2017

Frontiers website link:
www.frontiersin.org

Conflict of interest statement

The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest

Author contribution statement

HW was responsible for performing feature extraction of the clinical data, literature search, causal factor model simplification, the software implementation of the learning in model space framework, and determined benchmarking protocol. PT provided the clinical context of ADHD, constructed the prior knowledge used by the Bayesian framework, aided development on the causal factor model, and designed the literature search protocol. PW aided HW on the literature search and written the background on ADHD. OD implemented the SVM and GP machine learning algorithms. BK constructed the linear mixed effects models. YS helped with the technical implementation of the Bayesian linear regression algorithm. PT provided the mathematical formulation of the learning in model space approach and guidance on software implementation. MC and TN suggested candidate modeling approaches compatible with the framework, suggested a number of conventional approaches to be evaluated, appraised the output of the Bayesian approach, and provided recommendations for the improvement of the implementation and the methodology to ensure fair comparisons. SI coordinated access to the clinical data, facilitated digitization of paper-based clinical records, and maintained the database. DC provided the clinical context of ADHD, suggested causal mechanisms for treatment response, and defined the scope of clinical data collection.

Keywords

attention-deficit hyperactivity disorder, Bayesian inference, machine learning, Methylphenidate, mixed-effects model, personalized medicine, prognosis, treatment response

Abstract

Word count: 327

Attention-Deficit Hyperactive Disorder (ADHD) is one of the most common mental health disorders amongst school-aged children with an estimated prevalence of 5% in the global population (1). Stimulants, particularly methylphenidate (MPH), are the first-line option in the treatment of ADHD (2, 3) and are prescribed to an increasing number of children and adolescents in the US and the UK every year (4, 5) though recent studies suggest that this is tailing off, e.g. Holden et al. (6).

Around 70% of children demonstrate a clinically significant treatment response to stimulant medication (7, 8, 9, 10). However, it is unclear which patient characteristics may moderate treatment effectiveness. As such, most existing research has focused on investigating univariate or multivariate correlations between a set of patient characteristics and the treatment outcome, with respect to dosage of one or several types of medication. The results of such studies are often contradictory and inconclusive due to a combination of small sample sizes, low-quality data or a lack of available information on covariates.

In this paper, feature extraction techniques such as latent trait analysis were applied to reduce the dimension of on a large dataset of patient characteristics, including the responses to symptom-based questionnaires, developmental health factors, demographic variables such as age and gender, and socioeconomic factors such as parental income. We introduce a Bayesian modeling approach in a 'learning in the model space' framework that combines existing knowledge in the literature on factors that may potentially affect treatment response, with constraints imposed by a treatment response model. The model is personalized such that the variability among subjects is accounted for by a set of subject-specific parameters. For remission classification, this approach compares favorably with conventional methods such as support vector machines and mixed effect models on a range of performance measures. For instance, the proposed approach achieved an area under receiver operator characteristic curve of 82-84%, compared to 75-77% obtained from conventional regression or machine learning ('learning in the data space') methods.

Funding statement

We gratefully acknowledge the support from the Engineering and Physical Sciences Research Council (EPSRC), grant number EP/L000296/1. The ADDUCE project from which this piece of research borrowed clinical data is funded by the European Union's Seventh Framework Programme for research, technological development and demonstration under grant agreement number 260576.

Ethics statements

(Authors are required to state the ethical considerations of their study in the manuscript, including for cases where the study was exempt from ethical approval procedures)

Does the study presented in the manuscript involve human or animal subjects: Yes

Please provide the complete ethics statement for your manuscript. Note that the statement will be directly added to the manuscript file for peer-review, and should include the following information:

- Full name of the ethics committee that approved the study
- Consent procedure used for human participants or for animal owners
- Any additional considerations of the study in cases where vulnerable populations were involved, for example minors, persons with disabilities or endangered animal species

As per the *Frontiers* authors guidelines, you are required to use the following format for statements involving human subjects:

This study was carried out in accordance with the recommendations of 'name of guidelines, name of committee' with written informed consent from all subjects. All subjects gave written informed consent in accordance with the Declaration of Helsinki.

The protocol was approved by the 'name of committee'.

For statements involving animal subjects, please use:

This study was carried out in accordance with the recommendations of 'name of guidelines, name of committee'. The protocol was approved by the 'name of committee'.

If the study was exempt from one or more of the above requirements, please provide a statement with the reason for the exemption(s).

Ensure that your statement is phrased in a complete way, with clear and concise sentences.

This is secondary data analysis of the ADDUCE ADHD study. ADDUCE has received a favourable ethical opinion from the relevant UK National Health Service (NHS) Research Ethics Committee (REC reference 11/ES/0016). All subjects gave written informed consent in accordance with the Declaration of Helsinki. The protocol was approved by the REC. Consent to information access and secondary data analysis has been granted. All data were handled as per relevant guidelines and rules for such information, including the Data Protection Act 1998 and as per the initial ethical approval.

Personalized Medication Response Prediction for Attention-Deficit Hyperactivity Disorder: Learning in the Model Space Versus Learning in the Data Space

Hin Kwan Wong¹, Paul A. Tiffin², Michael J. Chappell^{3,*}, Thomas E. Nichols¹, Patrick R. Welsh⁴, Orla M. Doyle⁵, Boryana C. Lopez-Kolkovska¹, Sarah K. Inglis⁶, David Coghill⁷, Yuan Shen⁸ and Peter Tiño⁸

¹*Institute of Digital Healthcare, Warwick Manufacturing Group, University of Warwick, Coventry, UK*

²*Department of Health Sciences, University of York, York, UK*

³*School of Engineering, University of Warwick, Coventry, UK*

⁴*School of Psychology, Newcastle University, Newcastle upon Tyne, UK*

⁵*Centre for Neuroimaging Sciences, King's College London, London, UK*

⁶*Division of Maternal and Child Health Sciences, Ninewells Hospital and Medical School, University of Dundee, Dundee, UK*

⁷*Royal Children's Hospital, Melbourne, Australia*

⁸*School of Computer Science, University of Birmingham, Birmingham, UK*

Correspondence*:

Mental Health and Addiction Research Group, Department of Health Sciences, Faculty of Science Area 4, A/RRC/204, Alcuin Research Resource Centre, University of York, Heslington, York, YO10 5DD, UK
paul.tiffin@york.ac.uk

ABSTRACT

Attention-Deficit Hyperactive Disorder (ADHD) is one of the most common mental health disorders amongst school-aged children with an estimated prevalence of 5% in the global population (1). Stimulants, particularly methylphenidate (MPH), are the first-line option in the treatment of ADHD (2, 3) and are prescribed to an increasing number of children and adolescents in the US and the UK every year (4, 5), though recent studies suggest that this is tailing off, e.g. Holden et al. (6).

Around 70% of children demonstrate a clinically significant treatment response to stimulant medication (7, 8, 9, 10). However, it is unclear which patient characteristics may moderate treatment effectiveness. As such, most existing research has focused on investigating univariate or multivariate correlations between a set of patient characteristics and the treatment outcome, with respect to dosage of one or several types of medication. The results of such studies are often contradictory and inconclusive due to a combination of small sample sizes, low-quality data or a lack of available information on covariates.

In this paper, feature extraction techniques such as latent trait analysis were applied to reduce the dimension of on a large dataset of patient characteristics, including the responses to symptom-based questionnaires, developmental health factors, demographic variables such as age and gender, and socioeconomic factors such as parental income. We introduce a Bayesian modeling approach in a ‘learning in the model space’ framework that combines existing knowledge in the literature on factors that may potentially affect treatment response, with constraints imposed by a treatment response model. The model is personalized such that the variability among subjects is accounted for by a set of subject-specific parameters. For remission classification, this approach compares favorably with conventional methods such as support vector machines and mixed effect models on a range of performance measures. For instance, the proposed approach achieved an area under receiver operator characteristic curve of 82–84%, compared to 75–77% obtained from conventional regression or machine learning (‘learning in the data space’) methods.

Keywords: attention-deficit hyperactivity disorder, Bayesian inference, machine learning, methylphenidate, mixed effects model, personalized medicine, prognosis, treatment response

1 INTRODUCTION

The ability to predict treatment response (or non-response) in patients with mental health issues is potentially beneficial to both clinicians and patients in a number of ways. First, any treatment is accompanied by the risk of adverse effects—where non-response is a probable outcome then the risks of treatment may outweigh the benefits. Second, prediction of treatment response may guide both the dose and choice of medication. For example, where adverse events are dose-dependent then a clinician may choose to abandon a treatment course if a patient was a probable non-responder. Third, response prediction helps to calibrate both clinician and patient expectations of treatment outcomes. Finally, identifying non-responders may prompt a re-appraisal of the diagnosis and formulation of a patient’s problem—misdiagnosis being one potential cause of non-response. These benefits certainly apply to Attention-Deficit Hyperactive Disorder (ADHD), which is one of the most common developmental disorders among school-aged children with an estimated prevalence of 5% in the general population worldwide (1). Stimulants, particularly methylphenidate (MPH), are the first-line option in the treatment of ADHD (2, 3). Stimulants are prescribed to an increasing number of children and adolescents in the US and the UK every year (4, 5), though recent studies suggest that this trend is tailing off e.g. Holden et al. (6). The beneficial effects of stimulant medication on the core symptoms of ADHD have been demonstrated by numerous clinical trials, reviews and meta-analyses (11, 12, 13, 14). Nevertheless, adverse effects of the medications are also common (14). The findings from previous research suggest that around 70% of children demonstrate a clinically significant treatment response to stimulant medication (7, 8, 9, 10). However, it is unclear which patient characteristics may moderate treatment effectiveness and whether non-response can be predicted.

To date, achieving accurate predictions of the clinical outcomes for patients with ADHD has proven elusive—most of the literature has focused on investigating the potential correlations between a set of patient characteristics and the outcome following treatment with one or more types of medication. Information relating to patient characteristics has mostly been in the form of subjective questionnaire ratings, clinical notes and qualitative psychometric data; for example, the ratings from symptom-based questionnaires such as the Swanson, Nolan and Pelham (SNAP) questionnaire (15, 16, 17, 18), along with demographic variables such as age, sex and social economic background. The results from such studies are often contradictory and inconclusive due to small sample sizes and/or limited availability and quality of data, especially in the temporal (longitudinal) domain.

Along with more conventional statistical approaches, machine learning has also shown promise in predicting treatment response or prognosis in healthcare applications. Indeed, recently a random forest regression analysis was used to predict outcome in a group of patients affected by Obsessive Compulsive Disorder (OCD) from a relatively small pool of questionnaire items, with a reported error rate of 24.6% (19). Likewise, there has been a previous attempt to use machine learning techniques to predict treatment response in ADHD (20); support vector machine classification from this study was reported as 84.6% accurate (not to be confused with the balanced accuracy measure used in this paper). However, in addition to demographic and clinical questionnaire-derived data, the study used genetic as well as neuroimaging and neuropsychological information as inputs. Such data are unlikely to be readily available to clinicians in routine practice.

In this paper we investigate whether the inclusion of prior knowledge relating to the potential mechanism behind the presentation of a mental health condition and characteristics of individual patients can add value in predicting treatment. Thus, we hypothesized that a pragmatic machine learning approach based on a mechanistic or parametric model (a ‘learning in the model space’ framework) for treatment response prediction may offer an advantage over more conventional methods (21, 22, 23, 24). This method represents each newly observed patient through a model; the models are personalized such that individual differences are accounted for by a set of subject-specific parameters. In the case of ADHD, developing a plausible mechanistic model is not straightforward—despite decades of research, the underlying mechanism for the disorder is not well understood. In addition, any mechanistic model would have to be based on data that are likely to be available in good, but routine, clinical practice.

This paper documents, within the ‘learning in the model space’ framework, a Bayesian linear regression model for the prediction of treatment response in a cohort of children diagnosed and treated for ADHD in the UK. The performance of this new approach is then compared with conventional regression and machine learning methods (‘learning in the data space’) to assess whether or not the new approach offers benefits, and if so under what circumstances.

2 MATERIAL AND METHODS

2.1 Participants

The children enrolled in the study were drawn from the ADHD Drug Use and Chronic Effects (ADDUCE) cohort study (25), covered by a data sharing agreement with patient consent. The participants were from the UK NHS Tayside region who had attended the ADHD treatment clinics held at Dundee and Perth, UK. 262 families of eligible children were contacted, of which 181 (70%) were recruited and data on 173 of them were obtained for the purpose of this study. In addition, data were available on 94 healthy controls. Out of the 173 patients (whose baseline data are available), 157 of them started dose optimization studies and therefore longitudinal (temporal) data are available (See Section 3.1). To be eligible for the ADHD group, children had to be 6–17 years of age, have a clinical diagnosis of ADHD (see below), have had no previous medical history of methylphenidate use (medication-naïve) and have parental and child consent/assent to commence. The criteria for the healthy control group were similar apart from them having no current or previous psychiatric diagnoses. The recruitment was carried out over a 30-month period from January 2012 to August 2014.

All patients in the ADHD group had already been clinically diagnosed with ADHD; this diagnosis was based on the clinical judgement of the assessing physician, informed by structured interviews with parents/carers, information provided by the child’s school, direct observation of the child at the clinic, and

at times, in their educational setting. Thus, the physician had to be satisfied that the child fulfilled the diagnostic criteria for a hyperkinetic disorder according to the International Classification of Diseases 10th edition (ICD-10) (26), or ADHD as defined by the Diagnostic and Statistical Manual 4th edition (DSM-IV) (27). This means that the child had to demonstrate disabling and pervasive inattentiveness, hyperactivity and impulsivity across a range of settings. The clinic was designed to implement a ‘dose optimization titration’ scheme of medication in children diagnosed with ADHD. This involved giving increasing doses of methylphenidate (as the first line medication) at roughly weekly intervals until remission from symptoms was achieved or problematic adverse effects were encountered. If remission was not achieved with a first line medication within recommended dosage limits, or if problematic side-effects were encountered then a second line drug was initiated, and again, increased in dosage, as before.

2.2 Assessment

A range of baseline social and demographic factors was recorded, including parental marital status, family composition, and socioeconomic status as indicated by the Scottish Index of Multiple Deprivation (SIMD) 2012 (28) derived from the family home postcode. A history of any previous psychiatric or non-psychiatric medication exposure was recorded, as were any physical health issues. Verbal and non-verbal intellectual functioning was estimated from parental reports and any educational issues noted. Problems with anxiety and low mood were rated using the short form of the Mood and Feelings Questionnaire (MFQ) with the parents, and where appropriate, the child as informants (29). Dystonia and abnormal movements were recorded using the Abnormal Involuntary Movement Scale (AIMS) (30, 534–7). Oppositional and ADHD symptoms and behaviours were rated, according to parental report, using the Swanson, Nolan and Pelham (SNAP-IV) questionnaire (15). Any substance used by the participants was recorded using the Substance Use Questionnaire (SUQ). Fine motor issues were recorded using the Developmental Coordination Disorder Questionnaire 2007 (DCDQ’07). Several sections of the Development and Well-Being Assessment (DAWBA) were used (31); these were 1) Rapidly Changing Mood (child and parent versions), 2) Tic disorders, including the Tourette syndrome, 3) Awkward and troublesome behaviour. Tic severity (where present) was also rated using the Yale Global Tic Severity Scale (YGTSS) (32). Possible behaviours associated with an underlying Autism Spectrum Disorder (ASD) were evaluated using the Social Communication Questionnaire (SCQ) (33). The Strengths and Difficulties Questionnaire (SDQ) (34) was used to rate parental perceived levels of pro-social behaviour, hyperactivity/impulsivity, conduct problems, emotional symptoms and peer relationship problems. The overall clinical impression was recorded using the Clinical Global Impression — Severity scale (CGI-S) (30, 218–22) and Children’s Global Assessment Scale (CGAS) (35).

Responses to medication, in terms of levels of ADHD symptoms, were reported by parents and recorded using the SNAP-IV questionnaire at each visit. Likewise, any potential adverse effects and co-morbidity problems were reported using the standard clinic proforma, along with weight, height and blood pressure of the child at each visit.

2.3 Feature extraction/factor analysis

The aforementioned questionnaires included a large number of items with categorical (binary or ordinal) response formats. Thus, in order to facilitate model development by reducing the dimensionality of the data whilst minimising the loss of information, a series of factor (latent trait) analyses were conducted.

The key questionnaires used in the modeling process were the SCQ, the SDQ, and the SNAP-IV (see the previous section). In particular, the SNAP-IV scores served as the outcome variables, which indicated

whether symptomatic remission had been achieved, following the dose-optimized titration of medication. The factor analyses sought to identify the dimensionality underlying the responses to the questionnaires and, consequently, the standardized factor scores represented the level of trait for each patient in that underlying dimension or construct.

In order to estimate the dimensionality, the sample of 173 patients and 94 healthy controls was randomly divided into two roughly equal exploratory and confirmatory datasets. A parallel analysis (36), adapted for categorical data, was then implemented in the freeware FACTOR (37) using unweighted least squares (ULS) estimation method. A weighted 'promax' rotation was deployed to achieve factor simplicity (38). The maximum number of plausible factors (latent variables) was assumed to be indicated at the point where the eigenvalues of the factors in randomly generated data exceeded those observed in the real data. A series of exploratory factor analyses (EFAs - adapted for categorical dependent variables) were then conducted to aid interpretation of the factors. Oblique 'geomin' rotation was used (39), assuming that, as in almost all psychological measures, underlying latent traits would be correlated with each other to some extent (40). A series of confirmatory factor analyses (CFAs) were then conducted using the held-back, confirmatory data (see Section 3.1 on cross-validation), in order to ensure that the factor structures derived fitted the data adequately. All EFAs and CFAs were conducted in the Mplus software version 7.1, using robust weighted least squares with mean and variance adjustment (WLSMV) as the estimation method (41). Remission was defined by a child having a reported factor score in the hyperactive and inattentive dimensions (both elicited from factor analysis) equivalent to a mean item score in the SNAP-IV of one or less, which is conventionally taken to indicate symptomatic remission (42, 43). The resulting symptom score thresholds are only slightly different for inattentiveness and hyperactivity (-0.97 versus -0.92).

3 MODELING APPROACH

The causal factor model, shown in Fig. 1, was derived using a rapid review approach to appraise and synthesize the existing evidence (44). This model also took into account the nature of the data available in the cohort and was modified accordingly. The goal is not for the causal model to be comprehensive or definitive, but to identify from the literature as many potential factors relating to treatment response as there are available from the dataset, as well as helping to elicit the Bayesian prior distributions (Section 3.2.1). Model development was based on a literature review. This involved running searches in the EMBASE, MEDLINE and PsycINFO databases using the synonyms for ADHD (e.g. hyperkinesis) combined with terms relating to treatment outcome or response, and the names of the medications (both scientific and trademarks, full and abbreviated) prescribed in the cohort. The medications include 1) immediate release methylphenidate (IR-MPH, e.g. Ritalin[®]), 2) long-acting methylphenidate (XR-MPH, e.g. Concerta XL[®], Equasym XL[®], Medikinet XL[®]), 3) dextroamphetamine (DEX, e.g. Dexedrine[®]) including its prodrug lisdexamfetamine dimesylate (e.g. Elvanse[®]), and 4) atomoxetine (ATOM, e.g. Strattera[®]). Secondary sources were followed up. The quality of trial-based studies could be appraised using the CONSORT checklist (45) and observational studies via the STROBE guidance (46). Two of the authors (HKW and PAT) then made a judgement, based on the findings reported in the literature and the perceived likelihood of bias or uncertainty as to what extent variables in the model might be related to treatment response in ADHD. The model derived was consequently used to populate prior distributions for the patient-specific model parameters (i.e. the hyperpriors). Where the evidence was uncertain or inconsistent, the variances (i.e. imprecision) of the hyperpriors were increased.

Not every piece of information mentioned in Section 2.2 was used for the purpose of modeling, because of insufficient data or multicollinearity between the variables. The causal factor model was then simplified

based on the breadth of available data from the cohort, leading to a much reduced model as shown in Fig. 2. Some factors were combined through another layer of feature extraction; for example, the *motor* and *control* latent factors, themselves also obtained from applying feature extraction to the DCDQ’07 questionnaire data (see Section 2.2), were combined with the *non-verbal communication* factor from the SCQ questionnaire to obtain a developmental adversity factor. Some factors were not obtained from standard questionnaires; for example, the *perinatal adversity* factor (see Fig. 2) was constructed from birth weight and gestation age; the *family size and socioeconomic status* factor combined the number of siblings, parental house ownership (owned, mortgaged or rented) and the SIMD 2012 index (28).

3.1 Data

There were 267 subjects whose baseline characteristics were measured (173 clinically diagnosed with ADHD and 94 healthy controls) at the first clinical appointment. Of the 173 non-controls, 157 were enrolled in dose optimization titration studies with parental consent, for whom longitudinal (temporal) data are available. The 157 patients with longitudinal data were randomized and 10-fold cross-validation partitions were constructed. Subjects were partitioned into 10 subgroups of roughly equal size in a patient-coherent fashion, i.e. data from a single patient only appeared in a single fold.

For all models investigated in this paper, a single fold was used as the validation dataset and the remaining nine folds were combined to serve as the training dataset. This process was iterated until each fold had served as validation data exactly once.

3.1.1 Baseline characteristics

We labeled the patient subjects by the indexing variable $s = 1, 2, \dots, N$. A set of L patient-specific baseline continuous latent factors, encoded in a row vector $\mathbf{b}_s \in \mathbb{R}_{1 \times L}$ was obtained by performing feature extraction as described in Section 2.3 over the questionnaires detailed in Section 2.2. Referring to Fig. 2, $L = 14$ factors were used for the baseline. Data from the controls in addition to the training dataset were utilized during feature extraction to ensure that the resulting latent factor models can sufficiently encompass the entire range of characteristics from ADHD patients to normal children. The resulting continuous latent factors would, in theory, be sufficiently representative of the information conveyed by the categorical questionnaire response variables.

To ensure that validation data were strictly not used for the model building, feature extraction was first performed using only training data from each of the folds (plus all the controls). This resulted in 10 sets of factor scores corresponding to each fold. The factor model structures (e.g. the number of factors per questionnaire) over the folds did not change across the folds, as statistical fit indices and Chi-square difference tests did not suggest that any changes were necessary. The factor models were then used to estimate the baseline factor scores for the validation sets in each of the folds.

Each of the 10 cross-validation runs resulted in a set of corresponding continuous latent factors, which were used as inputs to subsequent models. The models were trained and validated using the same training-validation partitioning used in the feature extraction process.

3.1.2 Longitudinal data

Each of the 157 subjects with longitudinal data visited the clinic a varying number of times—from titration, stabilization to continuing care; the number of doctor’s appointments, A_s , varies from 1 to 22. At each appointment, the parent or guardian of the patient was asked to fill in an 18-item SNAP-IV questionnaire, which measures the degree of inattentiveness and hyperactivity. The responses were

entered into a factor model (identified through feature extraction) to extract a continuous symptom score for inattentiveness and hyperactivity. We denote the appointment number by the indexing variable a so that $a = 1, 2, \dots, A_s$. Let the independent ‘input’ variables $m_{a,1}, m_{a,2}, m_{a,3}, m_{a,4}$ be the four types of medications, respectively, IR-MPH, XR-MPH, DEX and ATOM for subject s at appointment a . Using datasheets for the medicines used, the dosages of DEX and ATOM were normalized to an equivalent daily dosage (EDD) of IR-MPH. For all a and s , this results in input and output matrices of the form:

$$\text{Input: } \mathbf{M}_s = \begin{bmatrix} m_{1,1} & m_{1,2} & m_{1,3} & m_{1,4} \\ m_{2,1} & m_{2,2} & m_{2,3} & m_{2,4} \\ \vdots & \vdots & \vdots & \vdots \\ m_{a,1} & m_{a,2} & m_{a,3} & m_{a,4} \\ \vdots & \vdots & \vdots & \vdots \\ m_{A_s,1} & m_{A_s,2} & m_{A_s,3} & m_{A_s,4} \end{bmatrix} \quad (1)$$

$$\text{Output: } \mathbf{y}_s = [r_1 \ r_2 \ r_3 \ \dots \ r_{A_s}]^T, \quad (2)$$

where r is the symptom severity measure and can be either the inattentiveness factor score or the hyperactivity factor score.

The combined EDDs of medications (for the 4 types) prescribed over the appointment number for all patients are plotted as a boxplot in Fig. 3. One can observe that as forced titration progressed over the appointments, the overall dosage level increased.

Fig. 4 shows the distribution of inattentiveness and hyperactivity symptom factor scores for the patients for each appointment. The lower the factor scores, the less severe the symptoms are. In terms of a general trend, one can clearly see an effective and quick reduction in symptom levels over the first 5 appointments, as stimulant medication prescription ramps up during forced titration. The symptom scores cease to improve for appointments 6–8, after which a slight increase can be observed. This hints at adherence or persistence issues, but the available data do not allow further investigation—as such issues are not consistently reported by the parent/guardian or recorded in the clinical notes. While the model has no mechanism for modeling such effects, the adaptive learning nature of the Bayesian algorithm is able to self-correct and compensate for small deviations, for example, by ‘learning’ to weight down the dose-response parameter for a given medication when the patient has a low adherence.

3.2 Treatment response model formulation

The treatment outcome is modeled as a linear combination of the baseline variables and the medication dosage,

$$\mathbf{y}_s = \mathbf{X}_s \boldsymbol{\omega}_s + \boldsymbol{\epsilon} \quad (3)$$

where

$$\mathbf{X}_s = \begin{bmatrix} \mathbf{b}_s & \vdots & \vdots & 1 \\ \vdots & \mathbf{M}_s & \vdots & \vdots \\ \mathbf{b}_s & \vdots & \vdots & 1 \end{bmatrix}, \quad (4)$$

$\epsilon \sim \mathcal{N}(0, \sigma_\epsilon, s^2) \in \mathbb{R}_{A_s \times 1}$ is an error term and $\omega_s \in \mathbb{R}_{P \times 1}$ is the *subject-specific* parameter vector moderating the effect of the baseline variables \mathbf{b}_s on the treatment response, i.e. it accounts for how large an effect each of the various baseline variables or medication types has on treatment outcome. ‘Subject-specific’ means that the parameter vector was allowed to be different for each subject so that patients with similar baseline characteristics can still have a different prediction outcome. The number of free parameters required is $P = L + 4 + 1 = 19$ for $L = 14$ (see Fig. 2).

The baseline variables remain unchanged over different appointments while the medication dosage may vary according to the titration regime specified by the clinician. Hence, every row of the matrix \mathbf{X}_s contains the same baseline characteristic vector for an individual patient, combined with the medication dosage vector. The row number in \mathbf{X}_s corresponds to the appointment number.

Because the number of visits A_s of the subjects was usually fewer than the dimension of the parameter space P , the problem is mathematically underdetermined. Hence, classical least squares regression methods would fail without an appropriate regularization (47). To this end, we employ a Bayesian formulation. In particular, the Bayesian linear regression was used to model the temporal evolution of the dose-response relationship for each patient.

In essence, a Bayesian approach allows prior or expert knowledge to be encoded into the problem formulation and this enables a probabilistic solution to be found despite the limited data available.

3.2.1 Prior distributions and knowledge

The joint prior probability density function $\Pr(\omega, \sigma^2)$ is given in the Appendix by Eqn. (15). In this exercise, the causal treatment response model based on the literature was used to constrain the prior of ω and its covariance matrix Λ_0^{-1} .

When eliciting the prior, quantitative information from the literature was not used, e.g. setting the mean of the prior distribution of the parameter vector ω to a specific numerical value. This is because the demography, sample sizes and effect sizes across literature vary and there is no correct way to normalize them. Instead, only the sign (direction) of the effect was encoded. For example, there is evidence that methylphenidate improves treatment response in the literature (a positive dose results in lower symptom score), therefore a negative value of -1 was specified for the columns of τ_0 corresponding to $m_{a,1} \dots m_{a,4}$ in Eqn. (1). For positive associations with symptom scores, $+1$ was used instead. The same magnitude is used in other factors (i.e. either $+1$ or -1).

Each cohort study or clinical trial from the rapid review was appraised, respectively, using the STROBE and the CONSORT checklists by counting the number of pass and fail items out of the total. Evidence from the literature was marked as good quality when both the checklist score was similar to other studies on the same topic (within 20% from the best) and effect sizes were statistically significant as reported by the authors for the sample size used. The diagonals of the covariance matrix Λ_0^{-1} were assigned an initial value of one; for every contradicting evidence (the effect sizes are opposite in direction) satisfying these criteria, 0.5 was added to the corresponding variance in the covariance matrix. A higher value may be specified if necessary, to ensure that the prior distribution of parameter ω spans both positive and negative sides sufficiently—within one standard deviation of τ_0 . On the other hand, if the effect sizes are positive the variance was reduced by 0.1 for each supporting studies, at the same time ensuring the variance does not go below 0.5. The (lack of) proposed existence of causal links between the variables in the causal model (see Fig. 1) ensures the sparsity of the covariance matrix Λ_0^{-1} .

While these numbers may not be completely objective, the amount of data available means that the sensitivity of the results to the prior is low—sensitivity analysis shows that the effect of scaling the prior covariance between 50% to 150% of its original values changes the errors by about 5% of the training root mean squared (rms) error, and 3% for the validation rms error.

3.2.2 Posterior distributions

The posterior distribution is given in the Appendix by Eqn. (16), where the parameters of the distributions are obtained through Eqn. (17). In the training set, Bayesian learning uses data from all of the appointments that a subject had, in which case $n = A_s$ where, as before, A_s is the total number of visits or appointments a subject has and data are available for. We introduce the simplified notations after Bayesian update has been applied to the training set using Eqn. (17), so that

$$\Lambda_{A_s}^{-1} \rightarrow \hat{\Lambda}_s^{-1}, \quad \tau_{A_s} \rightarrow \hat{\tau}_s, \quad \alpha_{A_s} \rightarrow \hat{\alpha}_s \text{ and } \beta_{A_s} \rightarrow \hat{\beta}_s. \quad (5)$$

This notation will be used in later sections. The reader should be reminded that the parameters are derived from each subject and hence are different across subjects. Notice that in a prediction exercise (instead of retrospective regression formulated here), the learning can be applied incrementally for each future observation with each update using just the new observation.

3.3 Virtual patient profile

When a new patient (denoted by $s = *$) is received, one can measure their baseline variables \mathbf{b}_* , but not their model parameter space ω_* . The goal is to estimate a *virtual patient profile* that is believed to best describe the new patient using only the available baseline measurements. To do this one derives the mathematical mapping functions from the baseline characteristics of a patient to their posterior parameters $\mathbf{b}_* \mapsto \Pr(\omega_*)$ and $\Pr(\sigma_*^2)$, such that a prediction can be made from the baseline variables. These functions are forged using machine learning on the existing pool of training data. Since $\Pr(\omega_s)$ is parameterized by (τ_s, Λ_s^{-1}) and likewise $\Pr(\sigma_s^2)$ by (α_s, β_s) , one has to learn the mappings from the baseline variables to the parameters. The learnt mathematical mapping functions can then be used to obtain estimates of $(\hat{\tau}_*, \hat{\Lambda}_*^{-1})$ and $(\hat{\alpha}_*, \hat{\beta}_*)$, which represent the virtual patient profile for the new patient in the model space.

Due to the conjugate nature of the priors, one does not need to derive the hyperparameters $\hat{\alpha}_*$ and $\hat{\beta}_*$ from Eqn. (5) for the purpose of having point estimates for the treatment response prediction. However, these hyperparameters are necessary in order to derive the posterior distribution of the predicted value—commonly referred to as the posterior predictive distribution. Knowing the distribution allows us to approximate the uncertainties of the estimates, e.g. 95% confidence intervals. Two methods were proposed for learning the mappings from the baseline variables to the virtual patient profile and they are introduced in the following subsections.

3.3.1 Method 1: Generalized linear regression

To determine the mappings, one finds the functions: a) $f_\tau(\mathbf{b}_*) \approx \hat{\tau}_*$ and b) $f_u(\mathbf{b}_*) \approx \hat{\Lambda}_*^{-1}$ where \mathbf{u}_s is a row vector containing non-zero elements of the upper (or lower) triangular part of $\hat{\Lambda}_s^{-1}$. Since $\hat{\Lambda}_s^{-1}$ is a covariance matrix (hence symmetric), knowledge of the lower/upper half of the off-diagonal elements plus the diagonal elements is sufficient to fully recreate the matrix.

The mappings are learnt from the training data, in which the posterior distributions of $\hat{\tau}_s$ and $\hat{\Lambda}_s^{-1}$ are already available through Eqn. (17). Linear regression models of the form $\mathbf{Y} = \hat{\mathbf{P}}\mathbf{B}$ were used to model the two mappings, where the matrix \mathbf{B} has rows of \mathbf{b}_s vectors—one for each subject in the training set—and

similarly \mathbf{Y} is composed of rows of a) τ_s for determining f_τ or b) \mathbf{u}_s for determining $f_{\mathbf{u}}$. The least squares solutions for the models are given by the Moore-Penrose pseudo-inverse,

$$\hat{\mathbf{Q}} = (\mathbf{B}^\top \mathbf{B})^{-1} \mathbf{B}^\top \mathbf{Y}. \quad (6)$$

For prediction, f_τ and $f_{\mathbf{u}}$ can both be formulated as $f(\mathbf{b}_*) = \mathbf{b}_* \hat{\mathbf{Q}}$.

The posterior estimates for the hyperparameters for a new patient are taken as the averaged values of $\hat{\alpha}_s$ and $\hat{\beta}_s$ across all subjects in the training set, resulting in $\hat{\alpha}_* = 5.5$ and $\hat{\beta}_* = 1.7$.

3.3.2 Method 2: Gaussian kernel weighted averaging

An alternative method is to find $\hat{\tau}_*$ and $\hat{\Lambda}_*^{-1}$ using a weighted average of $\hat{\tau}_{s'}$ and $\hat{\Lambda}_{s'}^{-1}$, with $s' \in \mathbb{S}_*$ being a subset of subjects in the existing training pool whose baseline variables ($\mathbf{b}_{s'}$) were ‘‘similar’’ to those of the new patient (\mathbf{b}_*). Highly similar subjects will have a higher influence on the value of $\hat{\tau}_*$ and $\hat{\Lambda}_*^{-1}$. The ‘(dis)similarity’ d_s is measured using the pairwise euclidean distance between \mathbf{b}_s and \mathbf{b}_* , such that

$$d_s = (\mathbf{b}_* - \mathbf{b}_s) (\mathbf{b}_* - \mathbf{b}_s)^\top. \quad (7)$$

This is then sorted and the 17.5% of subjects in \mathbb{S}_* with the smallest ‘dissimilarity’ values are kept; this percentage value was chosen as it resulted in the lowest validation error. The weighting w_s was taken as the normalized Gaussian kernel

$$w_s = \frac{\exp(-\lambda \cdot d_s)}{\sum_{\forall s' \in \mathbb{S}_*} \exp(-\lambda \cdot d_{s'})} \quad (8)$$

where the parameter value $\lambda = 1.15$ was chosen as it again resulted in the lowest validation error. Using equations (7) and (8), one can estimate $\hat{\tau}_*$ and $\hat{\Lambda}_*^{-1}$ as

$$\hat{\tau}_* = \sum_{\forall s \in \mathbb{S}_*} w_s \tau_s, \quad (9a)$$

$$\hat{\Lambda}_*^{-1} = \sum_{\forall s \in \mathbb{S}_*} w_s \Lambda_s^{-1}, \quad (9b)$$

and similarly the estimates of the hyperparameters are calculated using

$$\hat{\alpha}_* = \sum_{\forall s \in \mathbb{S}_*} w_s \alpha_s, \quad (9c)$$

$$\hat{\beta}_* = \sum_{\forall s \in \mathbb{S}_*} w_s \beta_s. \quad (9d)$$

3.4 Prediction using the posterior predictive distribution

When a new subject visits the clinician, their \mathbf{b}_* vector may be measured and used to approximate $\hat{\tau}_*$, $\hat{\Lambda}_*^{-1}$, $\hat{\alpha}_*$ and $\hat{\beta}_*$ using either of the methods in the previous subsections. Given a hypothetical medication input \mathbf{x}_* , the treatment response for the new subject can then be predicted through the posterior predictive

distribution

$$\Pr(y_*) = t_\nu \left(\mathbf{x}_* \hat{\boldsymbol{\tau}}_*, \frac{\hat{\beta}_*}{\hat{\alpha}_*} \left(\mathbf{I} + \mathbf{x}_* \hat{\boldsymbol{\Lambda}}_* \mathbf{x}_*^\top \right) \right) \quad (10)$$

where the number of degrees of freedom for the Student's t -distribution is given by $\nu = 2\hat{\alpha}_*$.

Equation (10) may be used to predict the treatment response for this new patient over their course of the treatment directly without learning; that is to treat each appointment as independent and the parameters are not updated. On the other hand, it is possible to perform incremental Bayesian learning over the course of treatment, by using Eqn. (17) to update the parameters given the treatment outcome measured for each new visit and the associated inputs. Through incremental learning, the model corrects for discrepancies between the true profile and the virtual patient profile of the new patient. As such, one would expect the prediction to improve as data from more visits to the clinic become available. The implementation of these methods is discussed in more detail in Section 4.4.

At some point, the profile of the new patient in terms of the treatment outcome, inputs, and baseline characteristics can be added to the pool of existing patient profiles (training set) to improve the model's generalizability for future patients.

3.5 Training and validation

As discussed in Section 3.1, 157 patients with longitudinal data were randomized and 10-fold cross-validation partitions were constructed resulting in 10 folds of training-validation data partitions.

First, for each fold, the framework detailed in Section 3.2 was followed and Bayesian linear regression was performed to fit patient-specific parameters to each patient in the training dataset. Second, either of the methods specified in Section 3.3 was used in order to construct virtual patient profiles for each patient in the validation set, using the patient-specific parameters. Finally, the procedure outlined in Section 3.4 was followed in order to obtain a prediction for patients in the validation set; effectively treating each patient as new.

3.6 Dichotomous remission prediction

Although the model was initially formulated to predict a continuous scale of symptom scores, one can explore dichotomizing the outcome into patients who have shown reduced symptoms and those who have not. The justification is that clinicians and doctors are less likely to be interested in a predicted SNAP-IV score or symptom severity scale as opposed to a simple “yes/no” answer as to whether the patient will be in remission for a given medication. A simple way to adapt the current model to do this is to apply a threshold to the continuous symptom score prediction, below which the patient is predicted to be in remission.

Some of the literature loosely defines remission in ADHD as having a large majority of SNAP-IV responses rated in category 0 (not at all) or 1 (a little) (42, 43). Therefore in this paper, the thresholds were chosen such that the approximate continuous symptom score corresponds to the raw responses from the 18-item SNAP-IV questionnaire all lying in category 1. The resulting symptom score thresholds are only slightly different for inattentiveness and hyperactivity (−0.97 versus −0.92, see Section 2.3).

Using these thresholds, it was found that the proportion of visits when measurements were taken indicates that remission was relatively rare, 160 out of a total of 1147 (13.95%) for the inattentiveness score and 139 (12.1%) for the hyperactivity score. This is expected, as forced titration initially starts with a low medication dosage and one would not expect an effective reduction in symptom ratings to remission levels

before the dosage was ramped up in later appointments; in addition, because of medical persistence issues patients can drop out before the clinicians are able to find an effective dose.

Note that all the methods in this paper were first used to predict the continuous symptom scores by regressing the baseline variables, medication prescribed to the treatment response at following appointments. Dichotomized remission prediction only occurs at a later stage. Right censoring (where patients prematurely drop out of dose optimisation stage without achieving remission) is therefore not an issue; regression methods can utilize the remaining appointment information to model treatment response regardless of whether remission was achieved or not.

4 PERFORMANCE METRICS

To facilitate a comparison between the performance of the different approaches, several performance metrics were used. For the regression tasks, one is interested in the deviation in the predicted symptom scores against the true symptom scores; whereas for the remission classification tasks, one is interested in the performance of the classifiers with regard to the probabilities or ratios of true positive, false positive, true negative and false negative cases.

4.1 Regression task

The root mean squared (rms) error measure is defined as the square of the averaged squared error across the 10 folds, across subjects and across all appointments for each individual, i.e.

$$\sqrt{\text{rms}} = \frac{1}{10|\mathbb{S}|} \sum_{k=1}^{10} \sum_{\forall s \in \mathbb{S}} \sum_{a=1}^{A_s} A_s^{-1} |y_s - \hat{y}_s|^2 \quad (11)$$

where \mathbb{S} is the set of all subjects considered (e.g. those in the validation set), $|\mathbb{S}|$ denotes the number of subjects in \mathbb{S} ; y_s , \hat{y}_s and A_s are, respectively, the true outcome symptom score, the fitted or predicted outcome symptom score, and the total number of appointments for the individual subject s .

4.2 Classification task

4.2.1 Sensitivity and specificity

The sensitivity (SEN, also known as the true positive rate or recall) is defined as $N_{\text{TP}}/N_{\text{P}}$ where N_{TP} is the number of *true positives*—appointments where measurements indicated remission and were correctly predicted as such; and N_{P} is the actual number of positive cases, i.e. the number of appointments where the corresponding subjects were indeed in remission. This is reported in addition to the specificity (SPC, also known as the true negative rate or fall-out), defined as $N_{\text{TN}}/N_{\text{N}}$, where N_{TN} is the number of *true negatives*—those *not* in remission and correctly predicted as such; and N_{N} is the actual number of negative cases (48). Note that if one lets N_{FP} and N_{FN} be the number of false positives and false negatives respectively, then $N_{\text{P}} = N_{\text{TP}} + N_{\text{FN}}$ and $N_{\text{N}} = N_{\text{TN}} + N_{\text{FP}}$ (48).

Sensitivity characterizes the ability of a classifier to rule out false negative predictions (type-II errors) given that a condition is true. On the other hand, specificity measures the ability of a classifier to rule out false positive predictions (type-I errors) given that a condition is false. In this exercise, the sensitivity measure is more important; due to the rarity of remission, and the goal is to try to predict what level of medication is required to achieve remission, the ability of a classifier to recall remission cases (ruling out type-II errors) is more important than ruling out type-I errors.

4.2.2 PPV and NPV

The positive predictive value (PPV, also known as the precision) is the proportion of true positives in the *predicted* positive cases and is the probability of remission given a positive prediction by the algorithm. As such, the PPV is a measure of the “quality” of a given positive prediction. PPV is given by $N_{TP}/(N_{TP} + N_{FP})$. Conversely, the negative predictive value (NPV) is the proportion of true negatives in the *predicted* negative cases, and is the probability of non-remission given a negative prediction. NPV is given by $N_{TN}/(N_{TN} + N_{FN})$ (48).

By the argument outlined above, the PPV is more important for this exercise than the NPV.

4.2.3 Balanced accuracy

The overall accuracy of a dichotomous predictor is defined by

$$\text{Accuracy} = (N_{TP} + N_{TN})/N,$$

where $N = 1147$ is the total number of appointments across all subjects.

However, the overall accuracy measure is known to be problematic when the prevalence of success/failure is low (49), i.e. the data are imbalanced (see also the end of Section 3.6). Due to this, some of the literature uses the balance accuracy (BAC) measure, defined as the average of sensitivity and specificity (50). This is the accuracy measure used throughout this paper. Note that, numerically, the BAC is closely related to the Youden’s J -statistic (51), also known as ‘informedness’ or ‘DeltaP’ (52), since it is equal to sensitivity plus specificity minus one.

4.2.4 ROC and AUC

The receiver operating characteristic (ROC) curve is commonly used in the medical and the machine learning community to evaluate the performance of binary classifiers (53). It plots the true positive rate (sensitivity) against the false positive rate (one minus specificity) for a given classifier. A curve is obtained when its classification performance can be tuned through setting a threshold or changing a parameter, trading off the true positive rate against the false positive rate. Binary classifiers that can achieve good compromise between sensitivity and specificity have a large area-under-the-curve (AUC), and this single metric may be used to compare the performance between the different classifiers (54).

4.3 Trading off sensitivity and specificity

From Section 3.6, the proportion of appointments without remission is $(100 - 13.95)\% = 86.05\%$. Therefore, given this statistic, one would expect that a null model guessing the result randomly would have a sensitivity of 13.95% and a specificity of 86.05%. Simply using point estimates of the continuous symptom score from the learning in the model space approach and thresholding them to give dichotomous predictions of remission results in classifiers with low sensitivity values between 22% to 28% and high specificity values of 94% to 97%. Due to the low number of remission cases compared to the non-remission cases, the classification is biased against predicting the remission cases, leading to low sensitivity (but high specificity). A classifier can be tuned to improve its sensitivity performance by trading off specificity to a certain degree. A good compromise would be maximising both sensitivity and specificity equally, which is in essence maximising the BAC or the Youden’s J -statistic in Section 4.2.3.

The thresholds for remission are defined by the SNAP-IV symptom factor scores as in Section 2.3, and this defines the ground truth of whether a patient is in remission or not. However, one can take advantage

of the fact that the predicted continuous symptom scores from the learning in the model space approach form full posterior distributions with uncertainties associated, and the levels of uncertainty are known (e.g. see the error bars in Figures 8 and 9). One may define a critical value as the lower bound of the prediction, above which the probability of the prediction being correct is $x\%$. Instead of the remission thresholds comparing against the point estimates, they may be compared against the point estimates minus a critical value. The larger the critical value, the higher the prediction score has to be in order to be classified as not in remission. This in effect is equivalent to raising the threshold, classifying more and more cases into remission, which increases the sensitivity and lowers the specificity. The range of ‘thresholds’ or classifier parameter settings that makes this trade-off can be used to generate a ROC plot (Section 4.2.4). A similar trade-off can be made with classical machine learning algorithms and will be discussed in Section 5.

The training data are used to find an optimal classifier setting in order to achieve the highest BAC, and the same classifier setting is then used to classify the validation data. This ensures that the validation data are not used to minimize the validation error.

4.4 Benchmarking and implementation

The Bayesian learning in the model space approach relies on prior knowledge (Section 3.2.1) and virtual patient profiles in the model space (Section 3.3), as well as iterative learning (Bayesian update) in order to function. To assess whether these components contribute to the prediction capability of the model, several implementation strategies are investigated, namely:

1. **Appointment-independent prediction (AI):** treating each appointment as independent (as the first appointment) and giving a prediction only using the virtual patient profile;
2. **Incremental Bayesian linear regression (BR):** The first prediction is performed in exactly the same manner as the appointment-independent case. Then, Bayesian linear regression using elicited priors (see Section 3.2.1) is applied progressively. That is, the effect/outcome of medication prescribed in appointment 1, then observed at appointment 2, is used in the regression model. Then, at appointment 3, outcomes from appointments 1 and 2 are used. Similarly, at appointment 4, outcomes from appointments 1, 2 and 3 are used, and so on. This means that except for the first prediction, the incremental Bayesian linear regression learns from scratch the patient-specific parameters (at each appointment) using the elicited priors. This essentially disregards any information already learnt from the current training set (the virtual patient profiles), treats the validation set as a new ‘training’ set, and performs basic Bayesian linear regression fitting. However, instead of all appointment outcomes being available for each new patient, as is the case during the training phase, one simulates the fact that information is progressively collected during the course of treatment for new patients. Since the virtual patient profiles are not utilized, this serves as a benchmark reference to evaluate the effectiveness of the constructed virtual patient profiles when compared with the next case; this represents a method that can be implemented even when no training data exist.
3. **Incremental Bayesian learning/update (BU):** The first appointment is predicted as for the previous two cases, but then when the true value is observed (in appointment 2), it is fed-back into the Bayesian learning model, i.e. Eqn. (17). This updated model is then used to generate a prediction. This progressive updating continues up to the most recent appointment. The crucial difference between this method and the BR method is that, here, the priors used were derived from the virtual patient profiles, as opposed to the elicited priors used in the BR method. When compared to the BR case, this highlights whether the model space offers any utility in aiding the prediction of treatment response.

The Bayesian approach was implemented ad hoc in MATLAB software with custom routines. The performance of the prediction was compared across the three different implementations above, using the validation data.

5 COMPARISON WITH CONVENTIONAL METHODS

To provide context to the results achieved using the learning in the model space approach, the performance of conventional linear regression methods and machine learning methods was also investigated.

5.1 Mixed effects models

Linear mixed effects models (MEM) are widely used in many fields; for instance, biology (55), ecology (56), linguistics (57) and social sciences (58). They extend upon classical linear regression techniques to support data that have some form of grouping. For example, in this paper, each patient had one or more clinical appointments, and the data from each subject form a group. For each patient s with a number of appointments (from 1 to A_s), the severity symptom score vector \mathbf{y}_s (as in Eqn. (2)) is, for simplicity, assumed to have a linear relationship with the baseline and treatment effect via the following formulation:

$$\mathbf{y}_s = q_s + b_0 + \hat{\mathbf{X}}_s \boldsymbol{\omega}_s + \boldsymbol{\epsilon}, \quad (12)$$

where $\hat{\mathbf{X}}_s$ is similar in structure to \mathbf{X} defined in Eqn. (4) but without the last column of ones; $\boldsymbol{\omega}_s \in \mathbb{R}_{P \times 1}$ is the subject-specific parameter vector for the fixed effects, $q_s \sim N(0, \sigma_q^2)$ is the random effect affecting only the intercept b_0 , and $\boldsymbol{\epsilon} \sim N(0, \sigma_{\epsilon,s}^2) \in \mathbb{R}_{A_s \times 1}$ is an error term assumed to have a normal distribution. Observe that in this model, the scalar intercept b_0 is a fixed component for all of the patients in the population and the random effect a_s is a subject-specific scalar and is not grouped under any other parameters.

A linear mixed effects model was constructed within the R software (59), using the package ‘*lme4*’ (60). Severity score predictions were produced by performing out-of-sample forecasts, i.e. on the validation data for each of the folds, using the ‘*predict*’ function in the R software. To generate a prediction, the random effects are assumed to be zero and the population intercept was used. The continuous symptom score regression results for the MEM are prefixed MEM.

For a dichotomized clinical remission classification, the symptom score thresholds 0.92 and 0.97 from Section 2.3 for hyperactivity and inattentiveness were used to generate the ground truths. Following the rationale in Section 4.3, the symptom scores predicted by the MEM are given thresholds at different levels to produce a set of classifiers trading off sensitivity against specificity. These threshold-adjusted classifiers are labeled τ_{aMEMC} . The best (in terms of Youden’s statistic) threshold settings found using the training data were used for the validation data; the thresholds were 0.05 and -0.20, respectively, for the inattentiveness and hyperactivity symptom scores. In addition, the ‘*melogit*’ function in the Stata software (61) was used to directly estimate a mixed effects logistic regression model—a MEM with a logistic link function that predicts the probability of the binary remission outcome. In this case, the threshold procedure was applied to the probabilities rather than the raw symptom scores. The resulting classifier is labeled τ_{rMEMC} .

5.2 Support vector machines and Gaussian processes

In addition to MEM, machine learning classification approaches using support vector machines (SVM) and Gaussian processes (GP) were benchmarked. Both the SVM and GP learning methods are kernel machines and were implemented using linear (dot product kernel: $k(\mathbf{x}_i, \mathbf{x}_j) = (\mathbf{x}_i \cdot \mathbf{x}_j)$) and nonlinear kernels (the Gaussian kernel: $k(\mathbf{x}_i, \mathbf{x}_j) = \exp[-\gamma(\|\mathbf{x}_i - \mathbf{x}_j\|^2)]$). Readers are invited to refer to Burges

(62) for a detailed description of support vector machines and to Rasmussen and Nickisch (63) for a detailed description of GP learning. Compared to MEM and the learning in the model space approach, the SVM and GP are non-parametric methods—there are no subject specific parameters to identify; the models map the subject-specific inputs, such as baseline characteristics and the medication dosage, to the output symptom scores.

For the SVM, nested cross-validation was employed to optimize the parameters in the model. A broad log range spanning $[10^{-3} : 10^2]$ was arbitrarily chosen as the search range for the regularization parameter C . Similarly, the gamma parameter of the Gaussian kernel was optimized in the log range spanning $[10^{-4} : 10^1]$. For the GP, the model parameters were optimized using conjugate gradient descent, avoiding the need for nested cross-validation. SVM and GP learning approaches were employed as regression models (support vector regression SVR and Gaussian process regression GPR) for the linear and nonlinear kernels to predict the clinical scores. Dichotomous remission predictions were obtained by thresholding the distance from the hyperplane for the SVM, and for thresholding the probabilistic predictions of class membership for GP. These binary classifiers are respectively labeled as SVC and GPC.

From Section 3.6, the number of remission cases outweighed non-remission cases by a ratio of roughly 1:7. For a classification task, this imbalance of data is problematic for many classification algorithms (64). To help alleviate this, a downsampling approach was implemented for both linear and nonlinear kernels of the SVM and GP classifiers $dsSVC$ and $dsGPC$. During the training phase, the non-responder class was downsampled randomly to match the number of training instances in the remission class. By repeating this downsampling procedure, an ensemble of 1000 classifiers was trained. A classification prediction was generated by majority voting of the ensemble. Additionally, for the SVMs, an alternative is to learn the regularisation parameters C on a per-class basis. The rationale is that a higher penalty for errors can be placed on the more abundant class (65); this method is referred to as the weighted SVM ($rwSVC$). The per-class C parameters C^+ and C^- were optimized using the ranges $[10^{-3} : 10^1]$ and $[10^{-2} : 10^2]$, respectively. Finally, for the Gaussian process classifier, it is possible to calibrate the probabilistic predictions in order to help account for imbalanced data (66); this approach is referred to as a re-calibrated GP ($rcGPC$).

Using MATLAB software, the SVM was implemented using the `libsvm` toolbox (67), the Gaussian process learning was implemented using the `GPML` toolbox (63).

6 RESULTS AND DISCUSSIONS

6.1 Continuous symptom score prediction

The rms errors across all models are reported in Table 1.

6.1.1 Learning in the model space approach

Looking at the learning in the model space approach, one can observe that the virtual patient profile construction method, labeled Method 2, resulted in lower errors overall compared to Method 1. In Method 1, the mappings were learnt using simple linear regression from baseline variables to the parameter space. In addition to this simple linear regression, low degree polynomial (quadratic to quartic) basis functions were tried; whilst degrees up to a cubic resulted in a slightly lower training error, there was worse generalizability (i.e. higher validation error). For Method 2, the incremental Bayesian learning (BUR) approach performed the best overall; its performance advantage over the appointment-independent prediction (AIR) approach is expected given that it allows the model to adapt to a new patient as the treatment continues. The

performance advantage over the retrospective Bayesian linear regression (BRR) approach can be attributed to the fact that the virtual patient profile (Section 3.3) had utilized the prior whilst the Bayesian linear regression only uses the elicited prior (see Section 3.2.1). This supports the fact that the training population was able to add valuable information to the prediction task.

We recall that the BUR constructs virtual patient profiles while the BRR only uses the prior knowledge. It is interesting to note that the BRR outperforms the BUR using Method 1, suggesting that Method 1 was not an effective method for incorporating information from existing patient models.

Figure 5 shows the rms values averaged across all subjects during the validation phase and sorted by the clinical appointment (visit) number. Data above 15 visits are not shown as only a single patient had more than 15 visits. There is a slight downward trend visible with the ν BUR; suggesting that incremental Bayesian learning approach is able to reduce the prediction error as more data are known about a new patient through repeated appointments. The ν BRR also shows a downward trend, but the error is slightly higher than the ν BUR. This is because the ν BRR starts with only the elicited prior and performs learning (fitting) when more data are available, unlike the ν BUR which starts off with information from the training set in the form of a constructed/estimated virtual patient profile.

Figure 8 illustrates some examples of Bayesian linear regression performed during the training phase and their associated fitting rms errors. It can be seen that Bayesian linear regression fits similarly well for both the inattentiveness and hyperactivity symptom scores. Switching to a new type of medicine is usually associated with larger uncertainty (error bars). Looking at subject #74 (Figure 8b) in particular, it can be seen that, despite having the same input dosage from appointments 6–8, there were variations in the severity of the ADHD symptoms. It is not possible to know the exact reason for the variation for this subject during this particular period, without further information—perhaps this was due to adherence issues (the patient not taking their medication as prescribed), physiological factors, measurement ‘noise’, or perhaps something else entirely. By design, Bayesian linear regression can only fit the same outcome given the same input. This does highlight the fact that the current model may not have enough information in the form of covariates to account for some of these factors.

A subset of results for the validation phase is plotted in Figure 9. For brevity, only prediction outcomes for hyperactivity symptom scores are shown and retrospective Bayesian linear regression results (BRR) are omitted. The lack of solid lines connecting the predictions in the topmost subplot serves as a reminder that the model does not incorporate temporal aspects for the case of appointment-independent (AIR) prediction, which treats each appointment as the first (new) appointment for a new patient. This is also why the 95% confidence intervals for AIR are larger (more uncertain) than those for the BUR. Also note that, by design, the prediction results for the first appointment are identical for both approaches.

The figures illustrate that, during prognosis, incremental learning does not always improve the prediction error compared to simply predicting at every appointment without updating the model using new information. However, based on the rms errors in Table 1a, one expects incremental learning to perform better overall across subjects, especially for subjects with a prediction offset, such as over- or under-estimates. This is illustrated by the results for subject #148 given in Figure 9b, where the virtual patient profile for this patient consistently underestimates the actual hyperactivity score. Here, the incremental Bayesian learning was able to adapt the parameter ω and shifted the prediction upwards, resulting in lower prediction errors over the subsequent appointments.

6.1.2 Conventional machine learning approaches

Looking at the results in Table 1b, the conventional approaches yield similar performance, with linear SVR and GPR methods performing better than their nonlinear counterparts. The mixed effects model has slightly worse results. Errors of linear SCR and linear GPR are similar to each other, and to those for the learning in the model space approach BUR with Method 2. We conclude that for the task of predicting continuous symptom scores with the dataset investigated, the learning in the model space approach performs comparably with conventional approaches.

6.2 Dichotomous remission prediction

6.2.1 Learning in the model space approach

For the learning in the model space approach, the ROC curves for the dichotomous predictor are plotted in Figure 6 for both of the virtual patient profile (Section 3.3) construction methods. As the results for inattentiveness and hyperactivity scores were similar, only the ROC curves for inattentiveness are shown. The AUC values are given in the legend. The crosses on the lines mark the resulting classifier performance if one uses point estimates for the continuous symptom score from the model and simply applies the clinical remission thresholds. The squares mark the classifiers that have critical values based on maximizing the Youden's J -statistic (or the BAC, see Section 4.2.3) for the training set—this is equivalent to the sensitivity and specificity measures being maximized equally as a function of the critical values. Lastly, the circles mark the best classifier for the validation set in terms of the Youden's J -statistic. The closer the squares are to the circles, the better optimized the classifier is assuming no knowledge of the validation dataset. Those optimized classifiers marked by squares in the graph are used to generate various binary classifier performance metrics (see Section 4.2) in Table 2.

Looking at Table 2. The confusion matrices (CFM) show the number of true positives and false negatives in the first column, and false positives and true negatives in the second column. These may be used to calculate any classification performance metrics not included in this paper, such as the F -measure.

Similar to the continuous symptom prediction task, the BRC outperforms the AIC showing that posterior information is utilized effectively. As in the continuous task, the virtual patient profile construction method labeled Method 2 is better overall than Method 1, but the difference is much smaller in the classification task and the advantage is not universal across all metrics, especially for the AIC. Note that the virtual patient profile construction method has little effect on the BRC as it does not use it. The BRC achieves higher sensitivity values but a lower PPV compared to the BUC, meaning that the BRC is better at recalling remission cases, but the remission predictions by the BUC are more reliable. The SPC achieved by the BUC is notably higher, being better at ruling out false positives.

6.2.2 Conventional machine learning approaches

Table 3 shows the binary classifier performance metrics for the conventional machine learning approaches. Apart from the AUC, all of the other metrics in the table were derived from classifier settings (set-points) that had optimized the balanced accuracy (BAC) during the training stage. Apart from rCGPC, the BAC values across the different approaches are similar. The MEC classifiers perform well compared with GPC and SVC, with consistently high AUC values for both inattentiveness and hyperactivity. However, the set-points of the MEC classifiers achieve lower sensitivity (but higher specificity) than the SVC. As mentioned in Section 4.2, a higher sensitivity is more important for this exercise. PPV is the other measure of interest; the lrMEC, in particular, achieved the highest PPV amongst all the conventional approaches—partially helped by its low sensitivity.

The ROC plot for the MEC is shown in Fig. 7. The $1r_{MEC}$ variant fitted the training set better but both the $1r_{MEC}$ and the t_{aMEC} achieve similar validation performance. Tracing the ROC values, the $1r_{MEC}$ seems more suitable for high specificity settings while t_{aMEC} appears to be more suitable for high sensitivity classification.

For the machine learning approaches GPC and SVC, linear models work better. This was similarly observed in the continuous symptom score prediction task. Comparing methods in tackling data imbalance, the weighted SVM classifier rw_{SVC} method performed better than the downsampled ds_{SVC} method, while the downsampled Gaussian process classifier ds_{GPC} method performed better than the re-calibrated rc_{GPC} . Looking at both the BAC and AUC metrics, rw_{SVC} and ds_{GPC} perform similarly, with the former slightly better at classifying remission of hyperactivity, whereas the latter is slightly better for inattentiveness.

Overall for the conventional methods, the rw_{SVC} achieves the best compromise between SEN and PPV, meaning that it can identify remission cases more readily and at the same time the remission predictions are more reliable. Comparing Tables 2 and 3, it can be seen that the learning in the model space approach is superior overall. With respect to the BAC and AUC measures, the best performing BUC approach has an advantage of about 6–7%. This is interesting given the similar performance in the continuous symptom score prediction task amongst all approaches. The rms error measure in the symptom score prediction task was based on point estimate calculations, and thus used no information on the shape of the posterior distribution. The posterior predictive distribution (Section 3.4) for the BUC has a Student's t -distribution specific to each patient. The distributions were used to construct a probabilistic threshold in trading off specificity and sensitivity. This subject-specific nonlinear thresholding procedure may have contributed to its performance advantage over other approaches.

6.2.3 Comparison with literature

As far as the authors are aware, Kim et al. (20) is the only published literature on treatment response prediction of ADHD patients using machine learning techniques. Their best attempt achieved an AUC value of 0.84 and 86.4% classification accuracy (that is, the percentage of correct predictions, different from the BAC measure used in this paper) using a wide range of information types including demographical, clinical, genetic, environmental, neuropsychological and neuroimaging measures. In comparison, this paper includes only the more readily obtainable demographical and clinical information and is able to achieve best-case AUCs of 0.82–0.84. Restricting to demographical and clinical information, the highest performing method using SVMs in Kim et al. (20) had an AUC of 0.69. Granted, the comparison is imprecise because the quality, quantity and sources of demographical and clinical information are different between this paper and Kim et al. (20). Judging from the AUC values achieved by SVMs in this paper of about 0.71 (see Table 3), the results appears to be very close to those in Kim et al. (20). Due to this similarity, the previous comparisons should be valid.

7 CLINICAL UTILITY AND FURTHER WORK

The proposed learning in the model space approach is capable of predicting, for an individual, the minimum dosage of a particular medication required to have a user-defined chance of achieving symptomatic remission. It is highly flexible and potentially can be extended to any disease or disorder where medication is used in the course of treatment, speeding up and reducing the cost of the dose optimization/forced titration process, and potentially improving the quality of life for patients by ending the treatment sooner.

The current model, however, does not take into account adverse drug reactions (ADRs), minimization of which is another goal of a dose optimization titration process. To improve clinical utility, it is essential that ADRs are modeled. While data on this are available from the clinical notes accompanying the ADDUCE trial, a different modeling approach is required to incorporate the many different types of ADRs, with prevalence ranging from infrequent to very rare.

While the proposed approach achieves excellent performance in terms of treatment response classification, there is room for improvement. One obvious way to achieve this is to incorporate more data, especially covariates that are functions of time. In this exercise for example, the body mass index and age variables measured at baseline (first appointment) of the patients contribute to the latent factors, which in turn form the baseline variables. As such, they do not vary over time. It may be worth investigating whether the addition of temporal covariates, such as blood pressure, would improve the model.

Another venue for potential improvement is to extend the linear model to a nonlinear model—there is no guarantee that all the covariates have a linear relationship with treatment response. Identifying the level and nature of nonlinear relationships is the first challenge. In the current Bayesian framework, the introduction of nonlinearities increases computational complexity for Bayesian inference, requiring the use of techniques such as Gibbs sampling.

There are other areas of interest. For example, what is the optimum strategy, in terms of timing and requirements, for incorporating semi-new patient data to the model space to improve the generalizability of the model for other new patients? How can medical adherence/concordance be modeled? Does gender of the patient play a role in their treatment response?

8 CONCLUSION

A learning in the model space framework has been utilized to develop a personalized medicine approach to treatment response prediction. First of all, factor analysis was performed to extract latent factors from a large clinical dataset, collected from a UK sample of 157 patients suffering from attention-deficit hyperactivity disorder. The resulting reduced-order patient information was then encoded in a model parameter space resulting in a cloud of personalized models. Then, the patient-specific model space parameters were used to train a Bayesian linear regression model. New patients are then matched to existing patients most similar to themselves to obtain a virtual patient profile, which in turn forms a prior parameter set for the Bayesian linear regression model. Through a Bayesian update algorithm, new data are continuously integrated to improve the prediction performance for a given patient. In addition, the parameters of the “new” patients can be added to the model parameter space (once sufficient data are available) to improve the generalizability of the model for future patients.

Comparisons were made between the learning in the model space approach with conventional data-driven machine learning and regression approaches. In terms of the prediction of the continuous symptom factor scores, the performance of the learning in model space framework was on a par with conventional approaches. However, the new approach is shown to outperform support vector machines, Gaussian processes and linear mixed effects classifiers in the prediction of symptomatic remission. The effective gain in classification performance of the new model can potentially speed up and reduce the cost of a forced titration or dose optimization titration process, which is normally manually performed by the clinician to assess the effective dosage of medication. Further work includes incorporating the prediction of adverse drug reactions, which is also an important element in the dose optimization titration process.

9 APPENDIX

9.1 Prior distributions and knowledge

With respect to equations (3) and (4), dropping the subscript for subject s for convenience, the outcome likelihood can then be modeled as:

$$\begin{aligned} \Pr(\mathbf{y} \mid \mathbf{X}; \boldsymbol{\omega}, \sigma^2) &= \mathcal{N}(\boldsymbol{\mu} = \mathbf{X}\boldsymbol{\omega}, \sigma^2\mathbf{I}) \\ &= \left(\frac{1}{2\pi\sigma^2}\right)^{0.5A_s} \exp\left\{-\frac{1}{2\sigma^2}(\mathbf{y} - \mathbf{X}\boldsymbol{\omega})^\top(\mathbf{y} - \mathbf{X}\boldsymbol{\omega})\right\} \end{aligned} \quad (13)$$

where \mathbf{I} is the $A_s \times A_s$ identity matrix.

The conjugate prior of the unknowns—the parameter vector $\boldsymbol{\omega}$ and the error variance σ^2 are given by, respectively, $\Pr(\boldsymbol{\omega} \mid \sigma^2) \sim \mathcal{N}(\boldsymbol{\tau}_0, \sigma^2\boldsymbol{\Lambda}_0^{-1})$ and $\Pr(\sigma^2) \sim \mathcal{G}^{-1}(\alpha_0, \beta_0)$, where the \mathcal{G}^{-1} is the Inverse-Gamma distribution (68), such that

$$\Pr(\sigma^2) = \frac{\beta_0^{\alpha_0}}{\sigma^{2(\alpha_0+1)}\Gamma(\alpha_0)} \exp\left(-\frac{\beta_0}{\sigma^2}\right); \quad \sigma^2, \alpha_0, \beta_0 > 0.$$

Assuming $\boldsymbol{\omega}$ and σ^2 are independent, the joint prior probability density function $\Pr(\boldsymbol{\omega}, \sigma^2)$ is then given by the Normal-Inverse-Gamma ($\mathcal{N}\mathcal{G}^{-1}$) prior, such that

$$\Pr(\boldsymbol{\omega}, \sigma^2) = \Pr(\boldsymbol{\omega} \mid \sigma^2) \Pr(\sigma^2) = \mathcal{N}\mathcal{G}^{-1}(\boldsymbol{\tau}_0, \boldsymbol{\Lambda}_0^{-1}, \alpha_0, \beta_0) \quad (14)$$

$$\begin{aligned} &= \frac{\beta_0^{\alpha_0}}{(2\pi)^{P/2} |\boldsymbol{\Lambda}_0^{-1}|^{1/2} \Gamma(\alpha_0)} \left(\frac{1}{\sigma^2}\right)^{\alpha_0+1+P/2} \times \\ &\quad \exp\left\{-\frac{1}{\sigma^2} \left[\beta_0 + \frac{1}{2}(\boldsymbol{\omega} - \boldsymbol{\tau}_0)^\top \boldsymbol{\Lambda}_0^{-1}(\boldsymbol{\omega} - \boldsymbol{\tau}_0)\right]\right\}. \end{aligned} \quad (15)$$

9.2 Posterior distributions

The posterior distribution can be parameterized as

$$\begin{aligned} \Pr(\boldsymbol{\omega}, \sigma^2 \mid \mathbf{y}) &\propto \Pr(\boldsymbol{\omega} \mid \sigma^2; \mathbf{y}) \Pr(\sigma^2 \mid \mathbf{y}) \\ &= \left(\frac{1}{\sigma^2}\right)^{\alpha_n+1+P/2} \exp\left\{-\frac{1}{\sigma^2} \left[\beta_n + \frac{1}{2}(\boldsymbol{\omega} - \boldsymbol{\tau}_n)^\top \boldsymbol{\Lambda}_n(\boldsymbol{\omega} - \boldsymbol{\tau}_n)\right]\right\}, \end{aligned} \quad (16)$$

where the two probabilities on the right-hand side correspond to, respectively, the multivariate normal distribution $\mathcal{N}(\boldsymbol{\tau}_n, \sigma^2\boldsymbol{\Lambda}_n^{-1})$ and the inverse-gamma distribution $\mathcal{G}^{-1}(\alpha_n, \beta_n)$. The parameters of the distributions are obtained through Bayesian learning:

$$\begin{aligned} \boldsymbol{\Lambda}_n &= \mathbf{X}^\top\mathbf{X} + \boldsymbol{\Lambda}_0 & \boldsymbol{\tau}_n &= \boldsymbol{\Lambda}_n(\boldsymbol{\Lambda}_0\boldsymbol{\tau}_0 + \mathbf{X}^\top\mathbf{y}) \\ \alpha_n &= \alpha_0 + \frac{n}{2} & \beta_n &= \beta_0 + \frac{1}{2}(\mathbf{y}^\top\mathbf{y} + \boldsymbol{\tau}_0^\top\boldsymbol{\Lambda}_0\boldsymbol{\tau}_0 - \boldsymbol{\tau}_n^\top\boldsymbol{\Lambda}_n\boldsymbol{\tau}_n) \end{aligned} \quad (17)$$

where n is the number of observations (68).

9.3 Marginal distributions

The marginal distribution of treatment outcome $\Pr(\mathbf{y})$ is found by integrating out (marginalising) ω_* and σ^2 over $\Pr(\mathbf{y} | \omega, \sigma^2)$ to obtain $\Pr(\mathbf{y})$. First of all, one computes the distribution $\Pr(\mathbf{y} | \sigma^2)$ by integrating out ω , which is

$$\begin{aligned}\Pr(\mathbf{y} | \sigma^2) &= \int \Pr(\mathbf{y} | \omega, \sigma^2) \Pr(\omega | \sigma^2) d\omega \\ &= \int \mathcal{N}(\mathbf{X}\omega, \sigma^2\mathbf{I}) \mathcal{N}(\tau_0, \sigma^2\Lambda_0^{-1}) d\omega.\end{aligned}$$

This can be shown to result in

$$\Pr(\mathbf{y} | \sigma^2) = \mathcal{N}(\mathbf{X}\tau_0, \sigma^2(\mathbf{I} + \mathbf{X}\Lambda_0^{-1}\mathbf{X}^\top)). \quad (18)$$

The marginal density $\Pr(\mathbf{y})$ can then be obtained by integrating out a \mathcal{NG}^{-1} density function, as

$$\begin{aligned}\Pr(\mathbf{y}) &= \int \Pr(\mathbf{y} | \sigma^2) \Pr(\sigma^2) d\sigma^2 \\ &= \int \mathcal{N}(\mathbf{X}\tau_0, \sigma^2(\mathbf{I} + \mathbf{X}\Lambda_0^{-1}\mathbf{X}^\top)) \mathcal{G}^{-1}(\alpha_0, \beta_0) d\sigma^2 \\ &= \int \mathcal{NG}^{-1}(\mathbf{X}\tau_0, \sigma^2(\mathbf{I} + \mathbf{X}\Lambda_0^{-1}\mathbf{X}^\top), \alpha_0, \beta_0) d\sigma^2 \\ &= t_\nu\left(\mathbf{X}\hat{\tau}_s, \frac{\hat{\beta}_s}{\hat{\alpha}_s}(\mathbf{I} + \mathbf{X}\hat{\Lambda}_s\mathbf{X}^\top)\right)\end{aligned} \quad (19)$$

where t_ν is the multivariate Student's t -distribution with the number of degrees of freedom $\nu = 2\hat{\alpha}_*$.

CONFLICT OF INTEREST STATEMENT

The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

AUTHOR CONTRIBUTIONS

HW was responsible for performing feature extraction of the clinical data, literature search, causal factor model simplification, the software implementation of the learning in model space framework, and determined benchmarking protocol. PT provided the clinical context of ADHD, constructed the prior knowledge used by the Bayesian framework, aided development on the causal factor model, and designed the literature search protocol. PW aided HW on the literature search and written the background on ADHD. OD implemented the SVM and GP machine learning algorithms. BK constructed the linear mixed effects models. YS helped with the technical implementation of the Bayesian linear regression algorithm. PT provided the mathematical formulation of the learning in model space approach and guidance on software implementation. MC and TN suggested candidate modeling approaches compatible with the framework, suggested a number of conventional approaches to be evaluated, appraised the output of the Bayesian approach, and provided recommendations for the improvement of the implementation and the methodology

to ensure fair comparisons. SI coordinated access to the clinical data, facilitated digitization of paper-based clinical records, and maintained the database. DC provided the clinical context of ADHD, suggested causal mechanisms for treatment response, and defined the scope of clinical data collection.

FUNDING

We gratefully acknowledge the support from the UK Engineering and Physical Sciences Research Council (EPSRC), grant number EP/L000296/1. The ADDUCE project from which this piece of research borrowed clinical data is funded by the European Union's Seventh Framework Programme for research, technological development and demonstration under grant agreement number 260576.

ACKNOWLEDGMENTS

The authors would like to thank the staff of Ninewells Hospital and Medical School, the University of Dundee and NHS Tayside; in particular trial database manager Emma McKenzie for coordinating the data exchange and clinical research nurse Jacqueline Paton for her meticulous data digitization efforts. The authors also would like to acknowledge the invaluable feedback by fellow POEMS (Predictive modeling for hHealthcare through MathS) network members and members in the ADDUCE consortium.

REFERENCES

- 1 .American Psychiatric Association. *Diagnostic and Statistical Manual of Mental Disorders* (Washington, DC: APA), 5th edn. (2013).
- 2 .Reeves G, Schweitzer J. Pharmacological management of attention-deficit hyperactivity disorder. *Expert Opinion on Pharmacotherapy* **5** (2004) 1313–20.
- 3 .Dopheide JA, Pliszka SR. Attention-deficit-hyperactivity disorder: an update. *Pharmacotherapy: The Journal of Human Pharmacology and Drug Therapy* **29** (2009) 656–79.
- 4 .Safer DJ, Zito JM, Fine EM. Increased methylphenidate usage for attention deficit disorder in the 1990s. *Pediatrics* **98** (1996) 1084–8.
- 5 .McCarthy S, Asherson P, Coghill D, Hollis C, Murray M, Potts L. Attention-deficit hyperactivity disorder: treatment discontinuation in adolescents and young adults. *The British Journal of Psychiatry* (2009) 273–7.
- 6 .Holden SE, Jenkins-Jones S, Poole CD, Morgan CL, Coghill D, Currie CJ. The prevalence and incidence, resource use and financial costs of treating people with attention deficit/hyperactivity disorder (ADHD) in the United Kingdom (1998 to 2010). *Child and Adolescent Psychiatry and Mental Health* **7** (2013) 34.
- 7 .Barbarese WJ, Katusic SK, Colligan RC, Weaver AL, Leibson CL, Jacobsen SJ. Long-term stimulant medication treatment of attention-deficit/hyperactivity disorder: Results from a population-based study. *Journal of Developmental & Behavioral Pediatrics* **27** (2006) 1–10.
- 8 .Schachter HM, Pham B, King J, Langford S, Moher D. How efficacious and safe is short-acting methylphenidate for the treatment of attention-deficit disorder in children and adolescents? A meta-analysis. *Canadian Medical Association Journal* **165** (2001) 1475–88.
- 9 .Swanson JM, Kraemer HC, Hinshaw SP, Arnold LE, Conners CK, Abikoff HB. Clinical relevance of the primary findings of the MTA: Success rates based on severity of ADHD and ODD symptoms at the end of treatment. *Journal of the American Academy of Child & Adolescent* **40** (2001) 168–79.

- 10 .Spencer T, Biederman J, Wilens T, Harding M, O'Donnell D, Griffin S. Pharmacotherapy of attention-deficit hyperactivity disorder across the life cycle. *Journal of the American Academy of Child & Adolescent Psychiatry* **35** (1996) 409–32.
- 11 .Banaschewski T, Coghill D, Santosh P, Zuddas A, Asherson P, Buitelaar J, et al. Long-acting medications for the hyperkinetic disorders. a systematic review and European treatment guideline. *European Child & Adolescent Psychiatry* **15** (2006) 476–495.
- 12 .Greenhill L, Kollins S, Abikoff H, McCracken J, Riddle M, Swanson JM. Efficacy and safety of immediate-release methylphenidate treatment for preschoolers with ADHD. *Journal of the American Academy of Child & Adolescent Psychiatry* **45** (2006) 1284–93.
- 13 .van der Oord S, Prins PJM, Oosterlaan J, Emmelkamp PMG. Efficacy of methylphenidate, psychosocial treatments and their combination in school-aged children with ADHD: A meta-analysis. *Clinical Psychology Review* **28** (2008) 783–800.
- 14 .Storebø OJ, Ramstad E, Krogh HB, Nilausen TD, Skoog M, Holmskov M, et al. Methylphenidate for children and adolescents with attention deficit hyperactivity disorder (ADHD). *Cochrane Database of Systematic Reviews* (2015). doi:10.1002/14651858.CD009885.pub2. Issue 11. Art. No.: CD009885.
- 15 .Swanson JM, Sandman CA, Deutsch C, Baren M. Methylphenidate hydrochloride given with or before breakfast: I. Behavioral, cognitive, and electrophysiologic effects. *Pediatrics* **72** (1983) 49–55.
- 16 .Atkins MS, Pelham WE, Licht MH. A comparison of objective classroom measures and teacher ratings of attention deficit disorder. *Journal of Abnormal Child Psychology* **13** (1985) 155–167.
- 17 .Swanson JM. *School-based assessments and interventions for ADD students* (Irvine, CA: KC publishing) (1992).
- 18 .Bussing R, Fernandez M, Harwood M, Hou W, Garvan CW, Swanson JM, et al. Parent and teacher SNAP-IV ratings of attention deficit/hyperactivity disorder symptoms: Psychometric properties and normative ratings from a school district sample. *Assessment* **15** (2008) 317–328.
- 19 .Askland KD, Garnaat S, Sibrava NJ, Boisseau CL, Strong D, Mancebo M, et al. Prediction of remission in obsessive compulsive disorder using a novel machine learning strategy. *International Journal of Methods in Psychiatric Research* **24** (2015) 156–169.
- 20 .Kim J, Sharma V, Ryan ND. Predicting methylphenidate response in ADHD using machine learning approaches. *International Journal of Neuropsychopharmacology* **18** (2015) 1–7.
- 21 .Doyle OM, Tsaneva-Atansaova K, Harte J, Tiffin PA, Tiño P, Diaz-Zuccarini V. Bridging paradigms: hybrid mechanistic-discriminative predictive models. *IEEE Transactions on Biomedical Engineering* **60** (2013) 735–742.
- 22 .Brodersen KH, Schofield TM, Leff AP, Ong CS, Lomakina EI, Buhmann JM, et al. Generative embedding for model-based classification of fMRI data. *PLoS Computational Biology* **7** (2011) 1–19. doi:10.1371/journal.pcbi.1002079.
- 23 .Chen H, Tiño P, Rodan A, Yao X. Learning in the model space for cognitive fault diagnosis. *IEEE transactions on neural networks and learning systems* **25** (2014) 124–136.
- 24 .Shen Y, Tiño P, Tsaneva-Atanasova K. A classification framework for partially observed dynamical systems. *ArXiv e-prints* (2016). Unpublished.
- 25 .The ADDUCE Consortium. The ADDUCE (attention deficit/hyperactivity disorder drugs use chronic effects) project (2016). .
- 26 .World Health Organization. *ICD-10*. World Health Organization, 10 edn. (2010).
- 27 .American Psychiatric Association. *Diagnostic and Statistical Manual of Mental Disorders* (Washington, DC: APA), 4th, text rev. edn. (2000).

- 28 .APS Group Scotland. Scottish index of multiple deprivation. Tech. rep., The Scottish Government, Edinburgh (2012).
- 29 .Angold A, Costello EJ, van Kämnen W, Stouthamer-Loeber M. Development of a short questionnaire for use in epidemiological studies of depression in children and adolescents: factor composition and structure across development. *International Journal of Methods in Psychiatric Research* **5** (1996) 251–262.
- 30 .Guy W. *ECDEU assessment manual for psychopharmacology*. DHEW No. ADM 76-338 (Rockville, MD: US Department of Health, Education and Welfare, Public Health Service, Alcohol, Drug Abuse and Mental Health Administration, NIMH Psychopharmacology Research Branch, Division of Extramural Research Programs), revised edn. (1974).
- 31 .Goodman R, Ford T, Richards H, Gatward R, Meltzer H. The development and well-being assessment: description and initial validation of an integrated assessment of child and adolescent psychopathology. *Journal of child psychology and psychiatry* **41** (2000) 645–655.
- 32 .Leckman JF, Riddle MA, Hardin MT, Ort SI, Swartz KL, Stevenson J, et al. The Yale Global Tic Severity Scale: initial testing of a clinician-rated scale of tic severity. *Journal of the American Academy of Child & Adolescent Psychiatry* **28** (1989) 566–573.
- 33 .Rutter M, Bailey A, Lord C. *SCQ: The Social Communication Questionnaire*. Western Psychological Services, Torrance, CA (2003).
- 34 .Goodman R. The strengths and difficulties questionnaire: A research note. *Journal of Child Psychology and Psychiatry* **38** (1997-07) 581–586.
- 35 .Shaffer D, Gould MS, Brasic J, Ambrosini P, Fisher P, Bird H, et al. A children's global assessment scale (CGAS). *Archives of General psychiatry* **40** (1983) 1228.
- 36 .Horn JL. A rationale and test for the number of factors in factor analysis. *Psychometrika* **30** (1965) 179–185.
- 37 .Lorenzo-Seva U, Ferrando PJ. FACTOR: A computer program to fit the exploratory factor analysis model. *Behavior research methods* **38** (2006) 88–91.
- 38 .Abdi H. Factor rotations. Lewis-Beck M, Bryman A, Futing T, editors, *Encyclopedia for Research Methods for the Social Sciences* (Thousand Oaks, CA: Sage Publications) (2003), 978–982.
- 39 .Asparouhov T, Muthén BO. Exploratory structural equation modeling. *Structural Equation Modeling: A Multidisciplinary Journal* **16** (2009) 397–438.
- 40 .Thurstone LL. Multiple factor analysis. *Psychological Review* **38** (1931) 406–427.
- 41 .Muthén BO, du Toit SHC, Spisic D. Robust inference using weighted least squares and quadratic estimating equations in latent variable modeling with categorical and continuous outcomes. (1997).
- 42 .Hechtman L. Effects of treatment on the overall functioning of children with ADHD. *The Canadian Child and Adolescent Psychiatry Review* **14** (2005) 10–15.
- 43 .Chou WJ, Chen SJ, Chen YS, Liang HY, Lin CC, Tang CS, et al. Remission in children and adolescents diagnosed with attention-deficit/hyperactivity disorder via an effective and tolerable titration scheme for osmotic release oral system methylphenidate. *Journal of Child and Adolescent Psychopharmacology* **22** (2012) 215–225.
- 44 .Khangura S, Konnyu K, Cushman R, Grimshaw J, Moher D. Evidence summaries: the evolution of a rapid review approach. *Systematic Reviews* **1** (2012) 10. doi:10.1186/2046-4053-1-10.
- 45 .Schulz KF, Altman DG, Moher D. CONSORT 2010 statement: updated guidelines for reporting parallel group randomised trials. *The BMJ* **340** (2010). doi:10.1136/bmj.c332.

- 46 .von Elm E, Altman DG, Egger M, Pocock SJ, Gøtzsche PC, Vandenbroucke JP. Strengthening the reporting of observational studies in epidemiology (STROBE) statement: guidelines for reporting observational studies. *The BMJ* **335** (2007) 806–808.
- 47 .Griffiths TL, Tenenbaum JB, Kemp C. *Bayesian Inference* (Oxford University Press), chap. 3. Oxford Library of Psychology (2012).
- 48 .Fletcher R, Fletcher S. *Clinical Epidemiology: The Essentials*. Epidemiology/Biostatistics (Baltimore, MD: Lippincott Williams & Wilkins) (2005).
- 49 .Alberg AJ, Park JW, Hager BW, Brock MV, Diener-West M. The use of “overall accuracy” to evaluate the validity of screening or diagnostic tests. *Journal of General Internal Medicine* **19** (2004) 460–465.
- 50 .Brodersen KH, Ong CS, Stephan KE, Buhmann JM. The balanced accuracy and its posterior distribution. *Pattern Recognition (ICPR), 20th International Conference on* (2010), 3121–3124. doi:10.1109/ICPR.2010.764.
- 51 .Youden WJ. Index for rating diagnostic tests. *Cancer* **3** (1950) 32–35.
- 52 .Powers DMW. Evaluation: from precision, recall and F -measure to ROC, informedness, markedness & correlation. *Journal of Machine Learning Technologies* **2** (2011) 37–63.
- 53 .Fawcett T. An introduction to roc analysis. *Pattern Recognition Letters* **27** (2006) 861–874. doi:10.1016/j.patrec.2005.10.010.
- 54 .Bradley AP. The use of the area under the roc curve in the evaluation of machine learning algorithms. *Pattern Recognition* **30** (1997) 1145 – 1159. doi:10.1016/S0031-3203(96)00142-2.
- 55 .Rico D, Martin-Diana AB, Bary-Ryan C, Henahan GTM, Frias JM. Simultaneous modelling of the thermal degradation kinetics of pectin methylesterase in lettuce (*lactuca sativa* l.) and carrot (*daucus carota* l.) extracts: Analysis of seasonal variation and tissue type. *Bioscience, Biotechnology, and Biochemistry* **71** (2007) 2383–2392. doi:10.1271/bbb.60484.
- 56 .Stevens MHH, Sanchez M, Lee J, Finkel SE. Diversification rates increase with population size and resource concentration in an unstructured habitat. *Genetics* **177** (2007) 2243–2250.
- 57 .Nooteboom SG, Quené H. Self-monitoring and feedback: a new attempt to find the main cause of lexical bias in phonological speech errors. *Journal of Memory and Language* **58** (2008) 837–861.
- 58 .Kliegl R, Rolfs M, Laubrock J, Engbert R. Microsaccadic modulation of response times in spatial attention tasks. *Psychological Research* **73** (2009) 136–146. doi:10.1007/s00426-008-0202-2.
- 59 .R Core Team. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria (2013).
- 60 .Bates D, Mächler M, Bolker B, Walker S. Fitting linear mixed-effects models using lme4. *Journal of Statistical Software* **67** (2015) 1–48. doi:10.18637/jss.v067.i01.
- 61 .StataCorp. *Stata Statistical Software: Release 14* (College Station, TX: StataCorp LP) (2015).
- 62 .Burges CJ. A tutorial on support vector machines for pattern recognition. *Data mining and knowledge discovery* **2** (1998) 121–167.
- 63 .Rasmussen CE, Nickisch H. Gaussian processes for machine learning (GPML) toolbox. *The Journal of Machine Learning Research* **11** (2010) 3011–3015.
- 64 .He H, Garcia E. Learning from imbalanced data. *Knowledge and Data Engineering, IEEE Transactions on* **21** (2009) 1263–1284.
- 65 .Osuna E, Freund R, Girosi F. Support vector machines: Training and applications. Tech. Rep. AIM-1602, Massachusetts Institute of Technology, Cambridge, MA (1997).
- 66 .Bishop CM. *Pattern recognition and machine learning* (Springer-Verlag) (2006).
- 67 .Chang CC, Lin CJ. Libsvm: A library for support vector machines. *ACM Transactions on Intelligent Systems and Technology (TIST)* **2** (2011) 27.

68 .O'Hagan A, Forester JJ. *Bayesian Inference* (John Wiley & Sons), *Kendall's Advanced Theory of Statistics*, vol. 2b, chap. 11. 2nd edn. (2004).

FIGURES

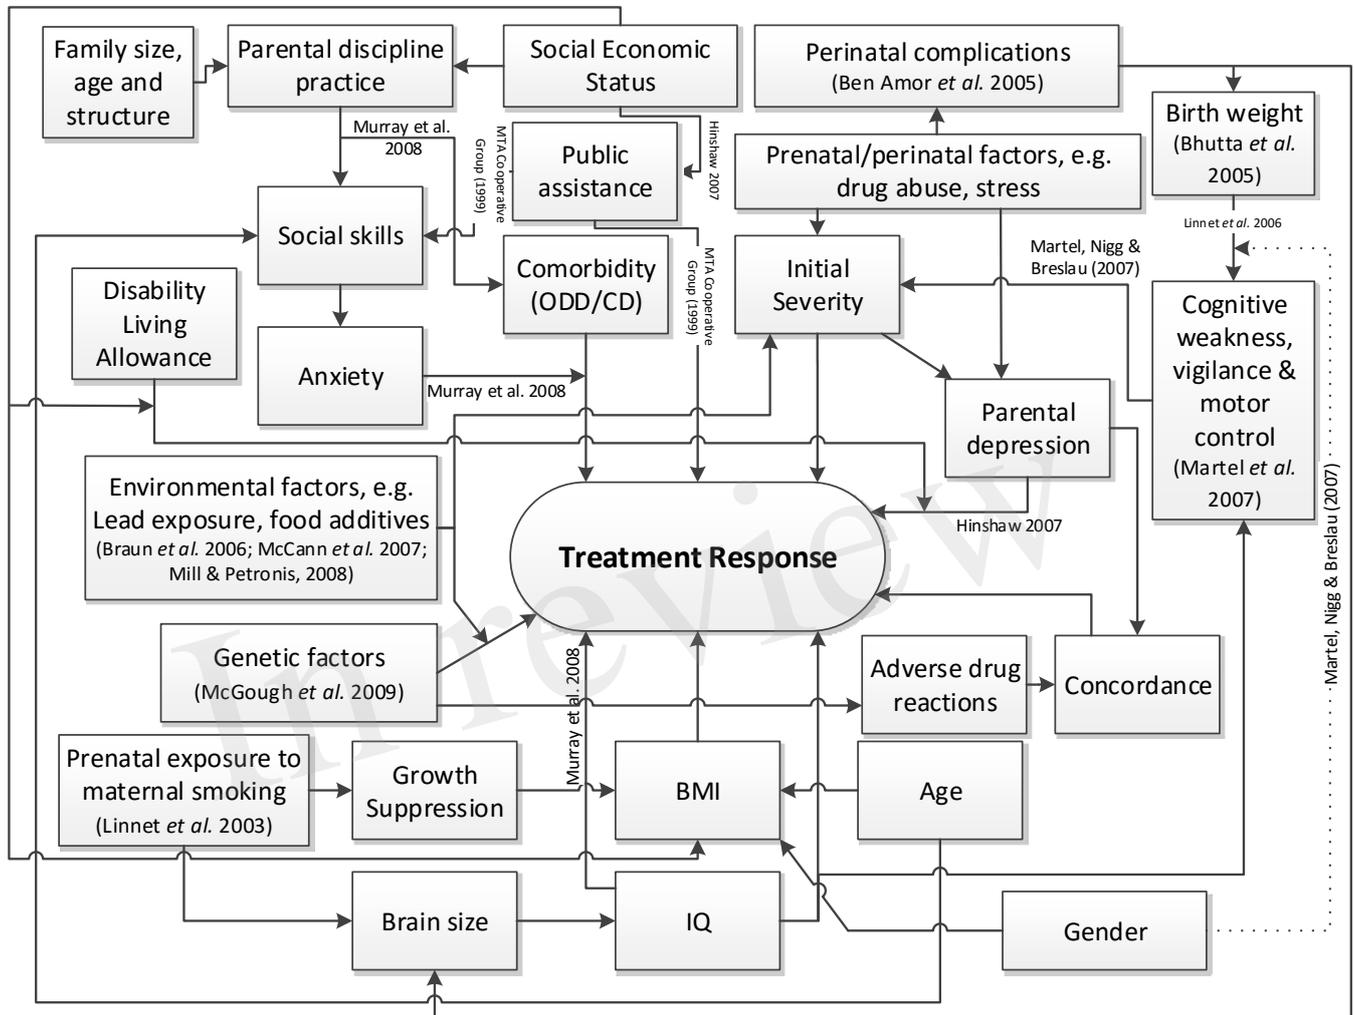


Figure 1. High level causal factor model of treatment response prediction in ADHD

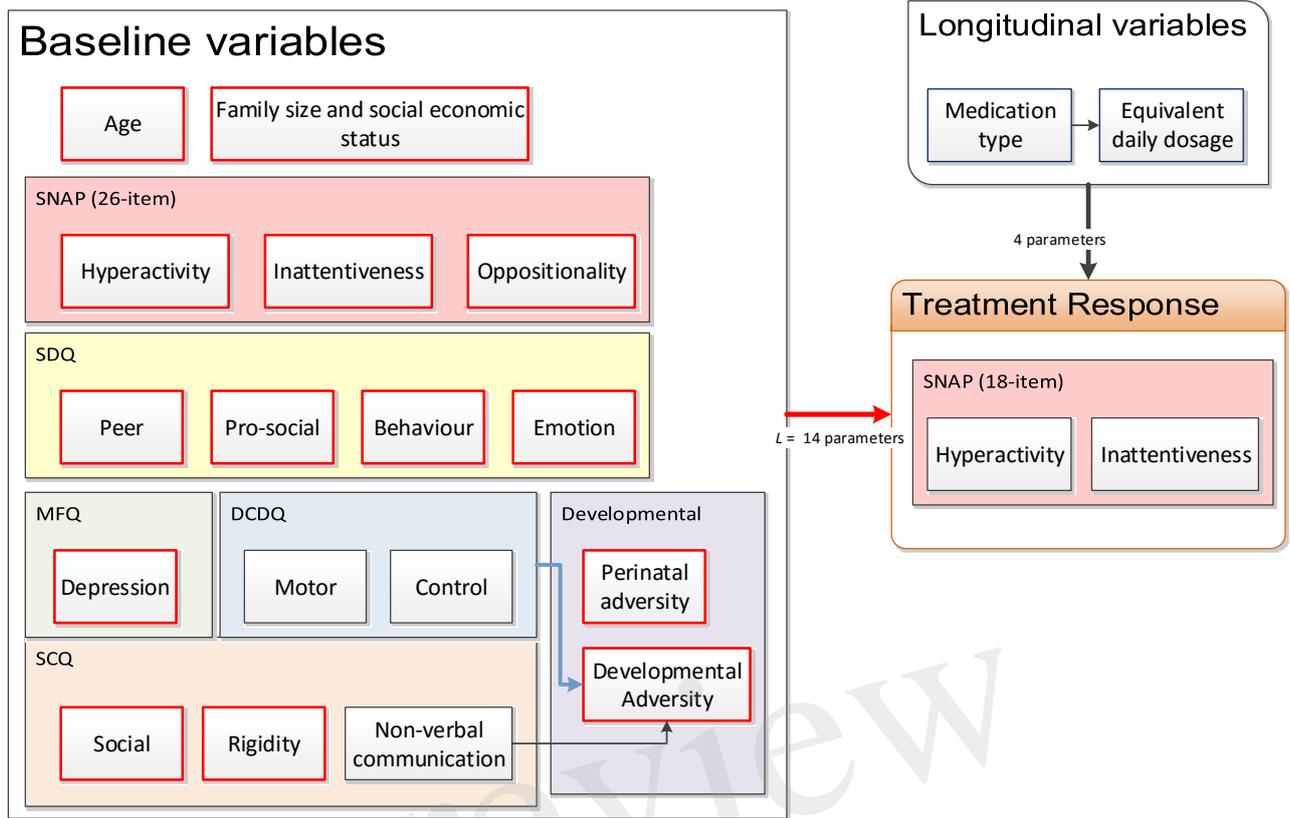


Figure 2. Reduced causal factor model of treatment response prediction in ADHD

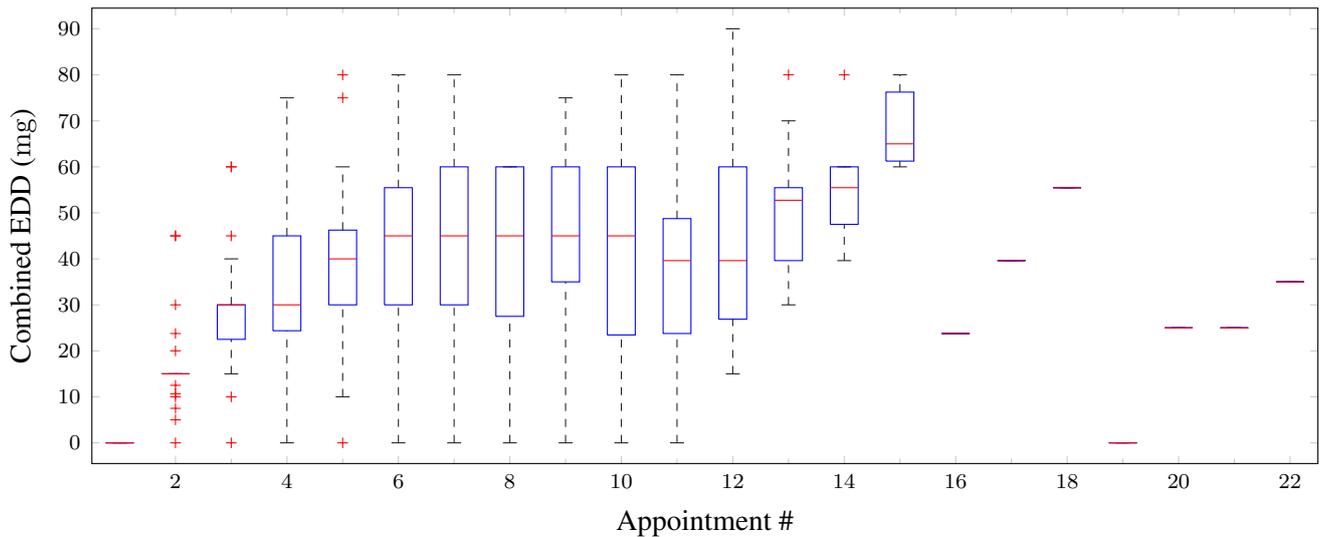


Figure 3. Boxplot of combined equivalent (in IR-MPH units) daily dosage of medications taken for all patients versus appointment number. Red horizontal lines: median; boxes: interquartile range; whiskers: 95% confidence intervals; red crosses: outliers.

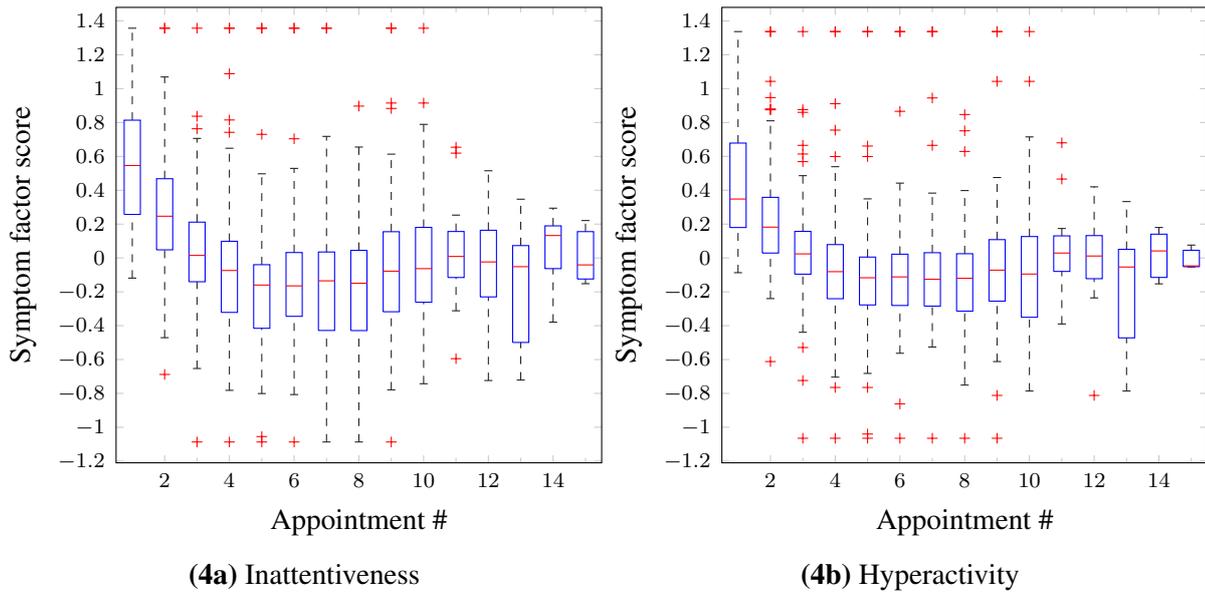


Figure 4. Boxplots of symptom scores across all patients versus appointment number. Red horizontal lines: median; boxes: interquartile range; whiskers: 95% confidence intervals; red crosses: outliers.

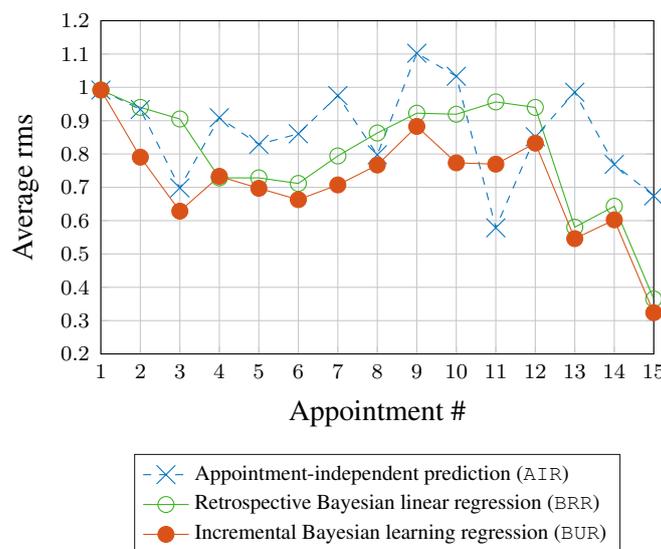


Figure 5. Rms prediction (validation) error averaged across all subjects versus appointment number

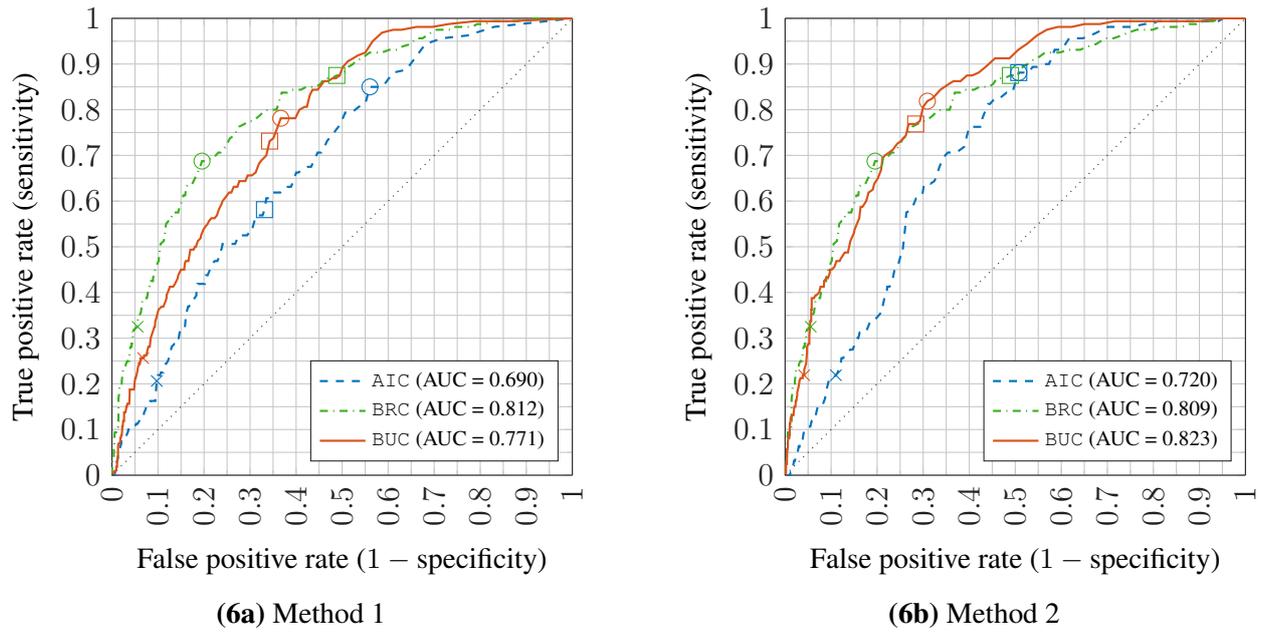


Figure 6. Receiver operator characteristic (ROC) plots of inattentiveness prediction using virtual patient profile constructed by Methods 1 and 2 in the learning in model space approach; crosses: no critical value adjustment (based on point estimates); squares: best performing critical values on training set; circles: best performing critical values on the validation set; AUC: Area under the ROC curve AIC: appointment-independent classifier; BRC: retrospective Bayesian linear regression classifier; BUC: incremental Bayesian learning/update classifier.

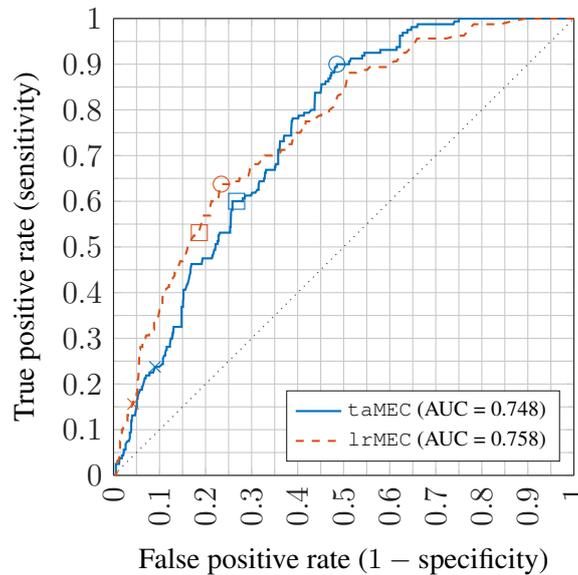


Figure 7. Receiver operator characteristic (ROC) plot of inattentiveness prediction using mixed effects models; crosses: no threshold adjustment (based on point estimates); squares: best performing threshold setting on training set; circles: best performing threshold setting on the validation set; AUC: Area under the ROC curve; taMEC: threshold-adjusted mixed effects classifier; lrMEC: logistic regression mixed effects classifier.

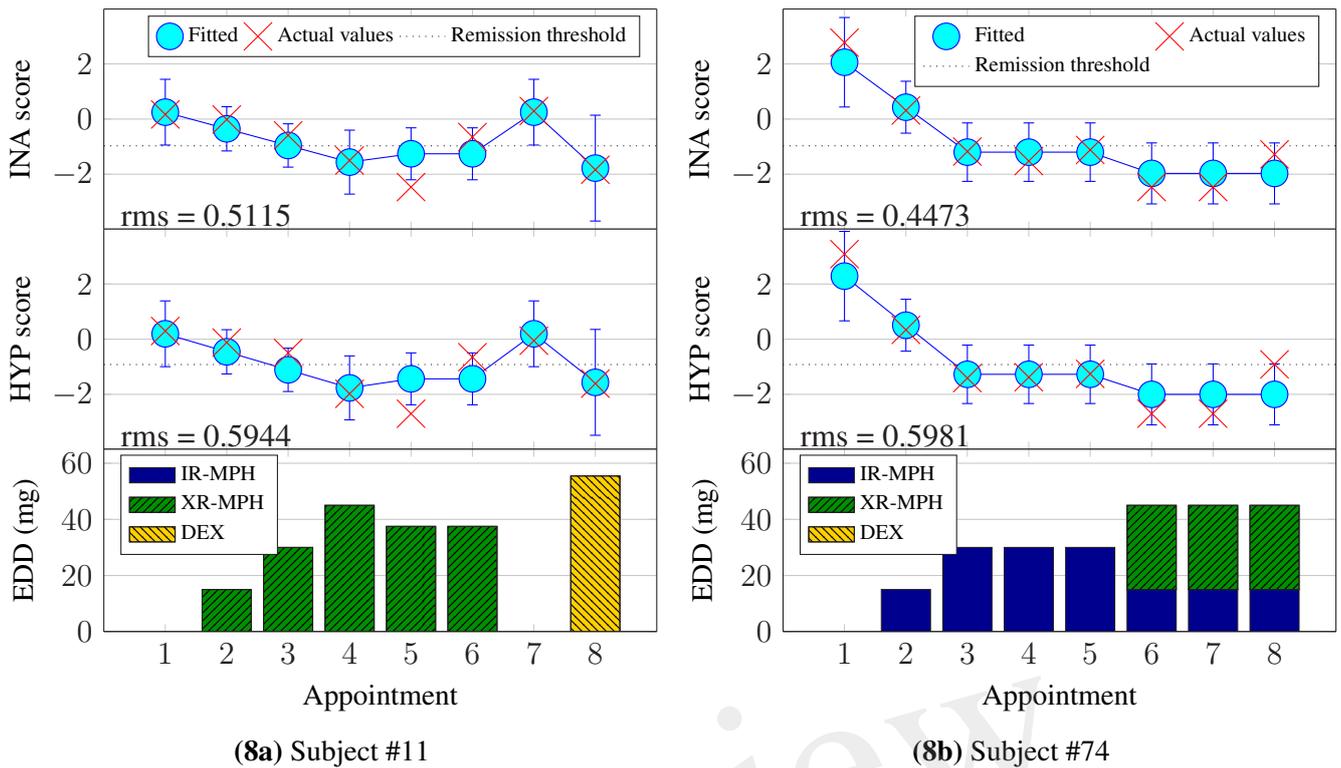


Figure 8. Examples of Bayesian linear regression on continuous inattentiveness (INA) and hyperactivity (HYP) symptom scores with the Bayesian linear regression training dataset (BRR).

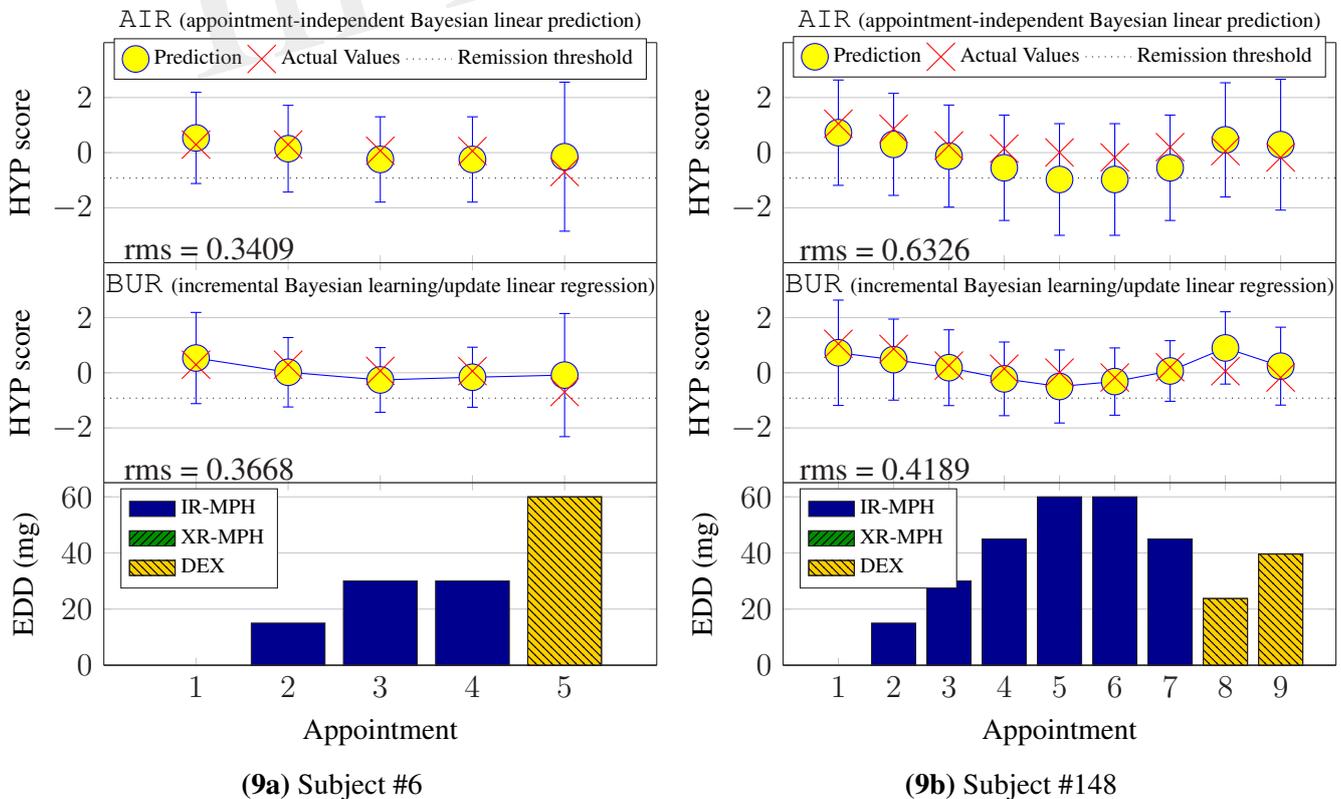


Figure 9. Examples of continuous hyperactivity symptom score prediction with the validation set.

Table 1. Rms errors for predicting symptom scores for inattentiveness (INA) and hyperactivity (HYP) using the (a) learning in model space and (b) conventional approaches.

(1a) Learning in model space				
	Inattentiveness		Hyperactivity	
	Method 1	Method 2	Method 1	Method 2
AIR [*]	0.98	0.84	0.97	0.85
BRR [†]	0.82	0.82	0.84	0.84
BUR [‡]	0.99	0.73	1.01	0.75

^{*} AIR: appointment-independent Bayesian linear prediction
[†] BRR: retrospective Bayesian linear regression [‡] BUR: incremental Bayesian learning/update linear regression

(1b) Conventional approaches				
	Inattentiveness		Hyperactivity	
Kernel:	Linear	Nonlinear	Linear	Nonlinear
SVR [*]	0.73	0.74	0.76	0.81
GPR [†]	0.72	0.77	0.76	0.84
MER [‡]	0.82		0.83	

^{*} SVR: support vector machine regression
[†] GPR: Gaussian processes regression
[‡] MER: mixed effects regression

Table 2. Sensitivity, specificity, accuracy and AUC of the remission classifier with critical values adjusted with respect to uncertainties in the predicted symptom scores.

		Inattentiveness				Hyperactivity			
		Method 1		Method 2		Method 1		Method 2	
CFM ¹	AIC*	93	327	141	501	82	314	121	448
		67	660	19	486	57	694	18	560
	BRC†	140	482	140	483	125	480	125	483
		20	505	20	504	14	528	14	525
	BUC‡	117	338	123	280	97	307	109	253
		43	649	37	707	42	701	30	755
SEN ²	AIC	58.1 %	88.1 %	59.0 %	87.1 %				
	BRC	87.5 %	87.5 %	89.9 %	89.9 %				
	BUC	73.1 %	76.9 %	69.8 %	78.4 %				
SPC ³	AIC	66.9 %	49.2 %	68.9 %	55.6 %				
	BRC	51.2 %	51.1 %	52.4 %	52.1 %				
	BUC	65.8 %	71.6 %	69.5 %	74.9 %				
BAC ⁴	AIC	62.5 %	68.7 %	63.9 %	71.3 %				
	BRC	69.3 %	69.3 %	71.2 %	71.0 %				
	BUC	69.4 %	74.3 %	69.7 %	76.7 %				
PPV ⁵	AIC	22.1 %	22.0 %	20.7 %	21.3 %				
	BRC	22.5 %	22.5 %	20.7 %	20.6 %				
	BUC	25.7 %	30.5 %	24.0 %	30.1 %				
NPV ⁶	AIC	90.8 %	96.2 %	92.4 %	96.9 %				
	BRC	96.2 %	96.2 %	97.4 %	97.4 %				
	BUC	93.8 %	95.0 %	94.4 %	96.2 %				
AUC ⁷	AIC	69.0 %	72.0 %	68.0 %	73.8 %				
	BRC	81.2 %	80.9 %	83.6 %	83.3 %				
	BUC	77.1 %	82.3 %	76.7 %	84.4 %				

¹ CFM: confusion matrix² SEN: sensitivity³ SPC: specificity⁴ BAC: balanced accuracy⁵ PPV: positive predictive value⁶ NPV: negative predictive value⁷ AUC: area under ROC curve

* AIC: appointment-independent classifier

† BRC: retrospective Bayesian linear regression classifier

‡ BUC: incremental Bayesian learning/update classifier

Table 3. Sensitivity, specificity, accuracy and AUC of the remission classifier with critical values adjusted with respect to uncertainties in the predicted symptom scores.

		Inattentiveness		Hyperactivity	
		Linear	Non-linear	Linear	Non-linear
Sensitivity	dsSVC ¹	70.0 %	70.6 %	67.6 %	66.9 %
	dsGPC ²	68.1 %	67.5 %	67.6 %	66.9 %
	rwSVC ³	76.9 %	33.8 %	76.9 %	69.1 %
	rcGPC ⁴	43.1 %	43.2 %	71.3 %	48.2 %
	taMEC ⁵	60.0 %		58.9 %	
	lrMEC ⁶	53.1 %		56.0 %	
Specificity	dsSVC	61.9 %	62.3 %	67.8 %	66.6 %
	dsGPC	67.9 %	69.1 %	67.8 %	71.3 %
	rwSVC	55.9 %	77.6 %	62.1 %	62.1 %
	rcGPC	59.0 %	18.4 %	50.7 %	50.6 %
	taMEC	73.3 %		73.8 %	
	lrMEC	81.4 %		81.4 %	
Balanced accuracy	dsSVC	66.0 %	66.5 %	67.7 %	66.7 %
	dsGPC	68.0 %	68.3 %	67.7 %	69.1 %
	rwSVC	66.4 %	55.7 %	69.5 %	65.6 %
	rcGPC	51.1 %	44.8 %	46.9 %	49.4 %
	taMEC	66.6 %		66.4 %	
	lrMEC	67.2 %		70.6 %	
Positive predictive value	dsSVC	23.0 %	23.3 %	22.4 %	21.6 %
	dsGPC	25.6 %	26.2 %	22.4 %	24.3 %
	rwSVC	22.0 %	19.6 %	21.9 %	20.1 %
	rcGPC	15.6 %	12.4 %	10.8 %	11.9 %
	taMEC	26.7 %		23.7 %	
	lrMEC	31.6 %		29.0 %	
Negative predictive value	dsSVC	92.7 %	92.9 %	93.8 %	93.9 %
	dsGPC	92.9 %	92.9 %	93.8 %	93.4 %
	rwSVC	93.7 %	87.8 %	95.1 %	93.6 %
	rcGPC	86.5 %	79.8 %	86.6 %	87.6 %
	taMEC	91.8 %		92.9 %	
	lrMEC	91.5 %		94.1 %	
Area under ROC curve	dsSVC	71 %	69 %	73 %	71 %
	dsGPC	75 %	71 %	73 %	70 %
	rwSVC	71 %	60 %	76 %	71 %
	rcGPC	49 %	41 %	46 %	48 %
	taMEC	74.8 %		77.5 %	
	lrMEC	75.8 %		77.2 %	

¹dsSVC: down-sampled support vector machine classifier; ²dsGPC: down-sampled Gaussian processes classifier; ³rwSVC: regularization-weighted support vector machine classifier; ⁴rcGPC: regularization-weighted support Gaussian processes classifier; ⁵taMEC: threshold-adjusted mixed effects classifier; ⁶lrMEC: logistic regression mixed effects classifier.