

GP-HD: Using Genetic Programming to Generate Dynamical Systems Models for Health Care

Mark Hoogendoorn
m.hoogendoorn@vu.nl
Vrije Universiteit Amsterdam,
Department of Computer Science
Amsterdam, The Netherlands

Ward van Breda
w.r.j.van.breda@vu.nl
Vrije Universiteit Amsterdam,
Department of Computer Science
Amsterdam, The Netherlands

Jeroen Ruwaard
j.j.ruwaard@vu.nl
Vrije Universiteit Amsterdam,
Department of Clinical, Neuro and
Developmental Psychology
Amsterdam, The Netherlands

ABSTRACT

The huge wealth of data in the health domain can be exploited to create models that predict development of health states over time. Temporal learning algorithms are well suited to learn relationships between health states and make predictions about their future developments. However, these algorithms: (1) either focus on learning one generic model for all patients, providing general insights but often with limited predictive performance, or (2) learn individualized models from which it is hard to derive generic concepts. In this paper, we present a middle ground, namely parameterized dynamical systems models that are generated from data using a Genetic Programming (GP) framework. A fitness function suitable for the health domain is exploited. An evaluation of the approach in the mental health domain shows that performance of the model generated by the GP is on par with a dynamical systems model developed based on domain knowledge, significantly outperforms a generic Long Term Short Term Memory (LSTM) model and in some cases also outperforms an individualized LSTM model.

ACM Reference Format:

Mark Hoogendoorn, Ward van Breda, and Jeroen Ruwaard. 2019. GP-HD: Using Genetic Programming to Generate Dynamical Systems Models for Health Care. In *IEEE/WIC/ACM International Conference on Web Intelligence (WI '19)*, October 14–17, 2019, Thessaloniki, Greece. ACM, New York, NY, USA, 8 pages. <https://doi.org/10.1145/3350546.3352494>

1 INTRODUCTION

Within the domain of health we are faced with an ever increasing amount of data that can be exploited for the benefit of the patient. There are many examples of insights that can be obtained from such data. One case is gaining understanding into how health states evolve over time and how they influence each other. Take the domain of mental health for instance, we might be interested to know what the mood of a depressed patient will be like in a few days time, and how the sleep quality influences this future mood. To derive such patterns from the data, we can apply temporal learning algorithms. When doing so, we have to make a choice in whether we are aiming for a

one-size-fits-all model or models per individual (cf. [10]). One-size-fits-all models provide generic insights, but often suffer from limited predictive value due to the inherent heterogeneity among patients. Individual models are tailored towards a person, but deriving useful information across all patients is difficult. In addition, limited data is typically available per patient, making it difficult to generate models that generalize well.

A technique that provides a middle ground between generic and individual models is dynamical systems modeling. These models represent the health states of a patient as a numerical value, and specify influence relationships between the states over time by means of difference equations. The equations include parameters that express the strength of the relationships. Hence, equations model the generic relationships and parameter values allow for individual tailoring. Unfortunately until now these kind of models need to be specified by exploiting domain knowledge rather than finding relationships in the data automatically (see e.g. [1, 3, 4, 19]). This makes the development time consuming, open to interpretation (as theories are often not precise enough to specify a difference equation), and it does not allow one to find new relationships in the data.

In this paper, we propose an approach that is able to generate dynamical systems models for health using Genetic Programming (GP). While several approaches have been proposed to generate these types of models using GPs (see e.g. [5, 17]) the health domain poses very different challenges. Rather than fitting these models towards a single dataset, in the health domain models should predict well across sets of patients, and should be able to cope with the variability of the patients by means of their parameter values. This has implications for the fitness function used by the GP. To develop such an approach for the health domain, we take an existing GP approach as a basis (cf. [5]). We extend the approach with a fitness function which is based on an evaluation framework for more knowledge driven dynamical systems models for the health domain (cf. [20]). Overall, this results in an approach to develop accurate and insightful predictive models for the domain of health. We refer to the approach as GP-HD (for GP Health-state Dynamics). We aim to answer the following research question in this paper:

Is the predictive performance of models generated by GP-HD better than state-of-the-art data-driven and knowledge-driven approaches for a real-life case study in the health domain?

To answer this question we evaluate the approach by means of a case study in mental health using a real dataset: forecasting the mood and perceived sleep quality of depressed patients up to three days

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

WI '19, October 14–17, 2019, Thessaloniki, Greece

© 2019 Association for Computing Machinery.
ACM ISBN 978-1-4503-6934-3/19/10...\$15.00
<https://doi.org/10.1145/3350546.3352494>

ahead. We compare the resulting model with an existing knowledge-driven dynamical systems models [1], individualized Long-Short Term Memory (LSTM) neural networks (cf. [9]) and a single generic LSTM model. We also evaluate the influence of several of the key hyperparameters of the system, while fixing others to values reported in prior literature.

This paper is organized as follows. First, we introduce our approach to generate the dynamical systems models in Section 2. We then present the dataset we use to evaluate our approach in Section 3, while Section 4 provides the experimental setup. The results are presented in Section 5, and finally Section 6 concludes the paper with a discussion.

2 APPROACH

Before going into detail on GP-HD itself, let us first consider the dynamical systems models we aim to learn. As said, we focus on the development of health states over time, i.e. we are faced with time series per patient and want to predict future values based on historical values. We assume that the equations of the dynamical systems model to make these predictions express the value of a state at time point $t + \Delta t$ based on values of states at t , for example: $s_1 t + \Delta t = s_1 t + \gamma_1 \cdot s_2 t$. The right hand side can contain a combination of states and parameters. As we are interested in not just modeling a single state, but a whole range of states, we have a system of such difference equations. Next, we will explain the GP part of the approach, followed by an explanation of the fitness function used.

2.1 Genetic Programming approach

Let us first focus on the generic GP loop, shown in Figure 1. This approach is taken from [5] with a few minor modifications.

Individual representation: We encode the individuals that represent the systems of equations by means of a vector of trees. Each tree T_i in the vector represents the difference equation to compute the new values for state i . Assuming m states we have a vector of m trees: $\langle T_1, \dots, T_m \rangle$. A tree can have mathematical operators as nodes (we use $*$, $+$, and $-$) while the leaves of the tree can either contain a parameter or a state. We put a limit to the depth of the tree (d_{max}) and the number of parameters used in the total vector of trees (λ_{max}).

Population initialization: We create an initial population of n individuals. For each tree, we start with an empty tree and randomly select one of the mathematical operators (with probability p_{op}) or randomly select one of the terminals with probability $1 - p_{op}$ (terminals are states or parameters) in case we have not reached d_{max} yet. In the event that we have reached d_{max} we always select a terminal.

Parent selection: We select parent(s) based on a standard tournament selection approach of size t_s .

Variation operators: According to the standard GP loop (cf. [8]) we perform either mutation, crossover, or we copy the original solution. We start by selecting one parent. We copy the parent with a fixed probability p_r . The mutation rate depends on the fitness of the parent that has been selected. Assuming the fitness of the parent being f

and the fittest one in the population being f_{max} we take the mutation probability as (based on [5]): $p_m = 0.1 + 0.2 \cdot \left(1 - \left(\frac{f}{f_{max}}\right)\right)$. Crossover is performed with the remaining probability (i.e. $p_c = 1 - p_m - p_r$).

In case of mutation, a random tree of the parent is selected from which we pick a random node or leaf, and replace it by a randomly generated subtree with a depth such that d_{max} is not exceeded. In case the crossover operator is selected, we select an additional parent using the same parent selection approach. After that, we apply either tree level crossover or vector level crossover with equal probability. In tree level crossover we merge the two parent trees for each position in the vector while in the vector level crossover for each position we randomly select a tree of the two parents with equal probability. Merging of trees is done by selecting a crossover point in each one of the two trees, and substituting the subtree from the crossover point of the first parent by the subtree at the crossover point in the second parent. Since it can result in trees that exceed the maximum depth, we try this ϕ times and in case we do not succeed the crossover fails and we select different parents.

Next Generation: We create a new population, and add the best individual of the previous generation (elitist approach). We generate children using the variation operators and fill up the new population until we reach the desired population size (n).

2.2 Determining the Fitness Value

Determining a fitness value is certainly not trivial for the type of health data we assume. We use the evaluation framework presented in [20] to determine the fitness value of a dynamical systems model. To make the paper self-containing the most important aspects of the approach are presented below.

In the evaluation framework, it is assumed that we have data which expresses the discretized measured values of all relevant states m of p patients over time. Here, the step size equals the step size used in the models (Δt), i.e. $\{x_{i,t_{start}}, x_{i,t_{start} + \Delta t}, \dots, x_{i,t_{end}}\}$. We split this data up into a training, validation, and test set, taking the first fraction as training data, the middle part as validation set, and the last part as test data. The goal of our models is to predict n time steps ahead and minimize the difference (in terms of the mean squared error) between the values of the states in the model and the real data. Depending on the goal of the model either all measured health states should be predicted well, or a subset thereof (this is a choice for the user of the model in the end). Hence, it is very likely that we want to optimize the predictive performance of multiple health states. This is known as a multi-objective optimization problem. Given the individual candidate model we have (as part of the population in the GP) and the criteria we want to optimize, we need to find values for the set of parameter that are present in the model $\{\gamma_1, \dots, \gamma_k\}$ (where $k \leq \lambda_{max}$). We do this *per patient*. A model with instantiated parameter values is called a *model instance*. Since this is a multi-objective optimization problem, multiple model instances can result that reside on a Pareto front. Each of these model instances obtains a certain score on each objective while not being dominated by other model instances (not scoring worse on one criterion while not performing better on any other). To derive such model instances, we use NSGA-II (cf. [7]). We run the algorithm r times per patient.

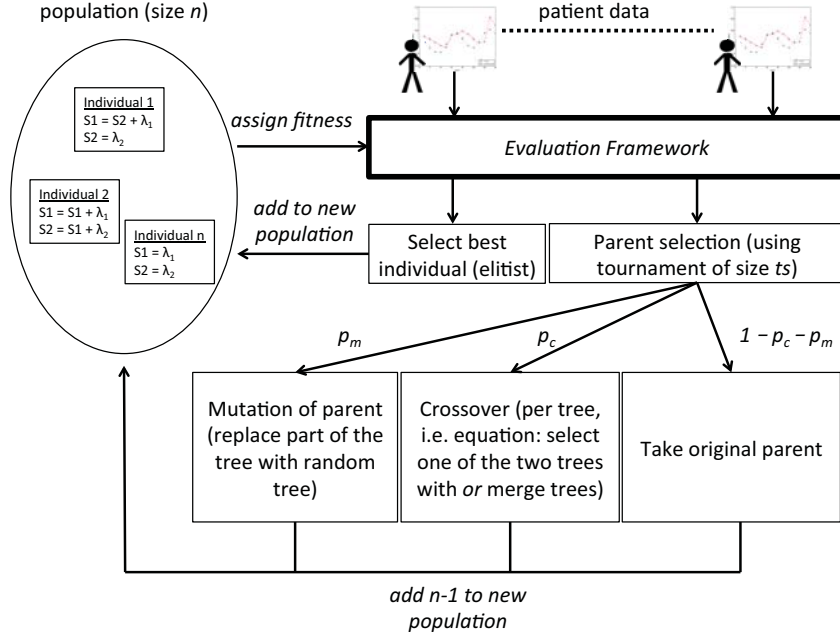


Figure 1: Genetic Programming loop for our framework.

Using these model instances, we compute a score for the candidate model M based on four criteria (see [20] for a more precise formulation and the rationale for each aspect):

2.2.1 Descriptive Capability. As said, multiple model instances result from the optimization process, each having their own niche (and error score) in terms of the objectives, none being dominated by another. To compute the score for the descriptive capability (i.e. the error on the training set) we compute the hypervolume from the worst point in the error space (an error of 1 on all evaluation criteria, as values are assumed to be normalized) to the Pareto front. This is called the dominated hypervolume. A value of 1 is the best value (i.e. the Pareto front including model instances with all zero error values). We compute the average dominated hypervolume per patient (remember, we run the algorithm r times) and compute the average over all patients p (μ_d) as well as the standard deviation σ_d . The descriptive score of an individual M (i.e. model) is:

$$descriptive_score_M = \mu_d 1 - \sigma_d \quad (1)$$

Hence, the higher and more consistent the score is over all patients, the higher the score in terms of descriptive capabilities.

2.2.2 Predictive Capability. In order to compute the predictive capability, we apply the model instances we have found for the training data to unseen validation data (and again, predict n steps ahead for as long as we have validation data, starting from the real initial values). This results in an error associated with each criterion for each model instance for each patient. We compute the mean μ_{ap} and standard deviation σ_{ap} over all of these errors and compute the predictive score:

$$predictive_score_M = 1 - \mu_{ap} 1 - \sigma_{ap} \quad (2)$$

2.2.3 Parameter Sensitivity. Parameters should be useful and have an impact on the performance of the model. A correlation analysis is performed between the parameter values and the error on the training set for each objective. A parameter is deemed useful in case the highest found correlation (of all correlations over all model instances, patients, and objectives) is above 0.35. The total number of useful parameters is then divided by the total number of parameters:

$$sensitivity_score_M = \frac{\sum_{l=1}^k useful_{M,\lambda_l}}{k} \quad (3)$$

2.2.4 Model Complexity. Finally, model complexity is weighed, the more complex the model, the lower the score for this aspect will be. It is defined by the number of parameters in the model, divided by the maximum (since we always have all states in the model):

$$complexity_score_M = \frac{k}{\lambda_{max}} \quad (4)$$

These scores are combined using a weighed sum:

$$fitness_M = w_1 \cdot descriptive_score_M + w_2 \cdot predictive_score_M + w_3 \cdot sensitivity_score_M + w_4 \cdot complexity_score_M \quad (5)$$

3 DATASET

We want to investigate how well our approach performs compared to alternative models. For this comparison, we require a dataset that includes a substantial number of patients for whom measurements of multiple health states have been performed over time. In our case, we have obtained such a dataset from the domain of mental health. Nowadays, interventions in mental health are becoming more and more digitized. For example, apps are being developed that can aid depressed patients to battle their depression (see e.g. [16]). Next to cost effectiveness, such apps also bring benefits when it comes to tracking the health state of patients as people carry their phone with them all the time. The tracking involves asking patients to score various aspects of their mental health on a regular basis (e.g. their mood) using pop ups on their mobile device. This is commonly referred to as Ecological Momentary Assessment (EMA). Our dataset originates from the E-COMPARED project. This project is focused on studying the effectiveness of interventions for depression. Within the project, a comparison is made between treatment as usual and blended care. The blended care setup features an app the depressed patients can use. Table 1 shown an overview of the questions posed to the patients. Note that not all questions are posed on a daily basis.

Table 1: The EMA questions that are present in the dataset.

Abbreviation	EMA question
Mood	How is your mood right now?
Worry	How much do you worry about things at the moment?
Self-Esteem	How good do you feel about yourself right now?
Sleep	How did you sleep tonight?
Activities done	To what extent have you carried out enjoyable activities today?
Enjoyed activities	How much have you enjoyed the days activities?
Social contact	How much have you been involved in social interactions today?

We obtain a dataset of 60 patients from the project. These are patients that have a long enough history to make up an interesting time series (at least 40 days of measurements). Some of the questions are posed multiple times a day (the mood), while others are only posed once a day or even less. To create a suitable dataset we: (1) normalize them on a scale in the range 0, 1; (2) aggregate the values per question on a daily basis by averaging in the event of multiple answers per day, and (3) in case a question does not have any answer on a day we linearly interpolate it based on the last known value and the first value in the future.

The resulting dataset contains on average 119.66 days of data per patient, with a large standard deviation of 70.55 days. Figure 2 shows a boxplot covering the different questions and the distribution of the answers the patients gave. Mood shows the narrowest distribution while the sleep and worrying questions seem to have the largest spread in answers.

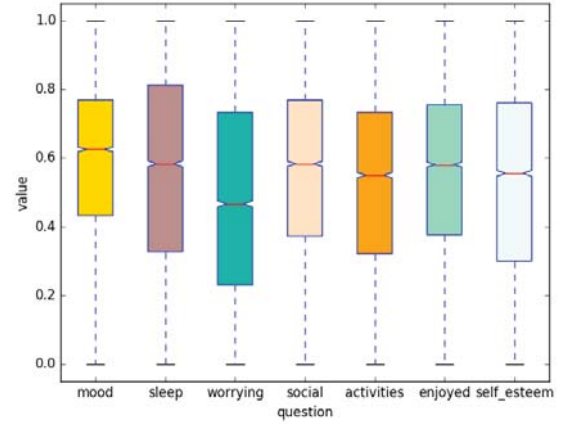


Figure 2: Distributions of the responses to the EMA questions.

4 EXPERIMENTAL SETUP

In our experimental setup, two phases are distinguished: (1) exploring hyperparameters of the GP framework and generating a dynamical systems model with the best settings, and (2) comparing the performance of the resulting model with alternative approaches. Before diving into those details, we first explain how we further prepare the dataset.

4.1 Dataset preparation

We select all data from a random sample of 30 patients as an *in sample* dataset and the 30 remaining patients as *out of sample* dataset. The data of each patient is split in a training, validation, and test set, containing 60, 20, and 20% of the data respectively, in a time ordered fashion. We evaluate the performance for two states in the dataset: *mood* and *sleep*, both deemed highly relevant by domain experts. The others are less important to predict in the future, but could contain important predictive information. The root mean squared is used as an evaluation metric. We predict 1, 2, and 3 time points (i.e. days) ahead.

4.2 Parameter Settings and GP runs

GP-HD contains a number of hyperparameters. An overview is shown in Table 2. Some hyperparameters we fix based on the literature (those that have been shown to work well for multiple types of problems), while we study the influence of others.

To optimize the varied hyperparameters of the NSGA-II algorithm we perform a number of runs with a small (random) sample of 10 patients from the *in sample* dataset. We create a fixed population of individuals and consider the mean and standard deviation of the fitness values we obtain over multiple runs of the NSGA-II algorithm. For the hyperparameters of the GP algorithm, we have based most of the settings on literature and performed some initial runs to set the more problem dependent hyperparameters appropriately. Initial experiments using the *in sample* data with λ_{max} and d_{max} showed that one parameter per state (i.e. 7 in total) works best, as well as a maximum depth of 6 (simple, yet sufficiently powerful models). We

Table 2: Parameter settings for GP-HD.

Parameter	Meaning	Values
<i>Fixed</i>		
p_{op}	Probability to select operator in random generation of (sub)trees	0.5 (cf. [5])
t_s	Tournament size	3 (cf. [5])
r_{max}	number of NSGA-II runs per patient	3 (cf. [20])
p_r	probability of retaining a copy of a parent	0.1 (cf. [5])
ϕ	Number of tries for crossover	3 (cf. [5])
w_1, \dots, w_4	Weight of evaluation criteria	All set to 0.25 (cf. [20])
<i>Varied</i>		
λ_{max}	Maximum number of parameters	various (7 selected)
d_{max}	Maximum depth of tree that represents model	various (6 selected)
pop_{nsga_II}	population size of the NSGA-II algorithm	{5, 10, 20, 50}
gen_{nsga_II}	number of generations for the NSGA-II algorithm	{5, 10, 20, 50}
pop_{gp}	population size of the GP	{50, 100}
gen_{gp}	number of generations of the GP	{50, 100}

study the influence of two hyperparameters in more detail, namely the population size and numbers of generations. We report results of experiments for those hyperparameters and focus on the convergence and the overall quality of the solutions in terms of fitness. We select the best hyperparameter setting, run it 10 times using the *in sample* data and select the best individual (i.e. model) we encounter. We continue with this model in the next phase.

4.3 Performance evaluation

We compare the best individual we have generated with three alternative approaches:

- (1) a dynamical systems model we obtained from the literature (cf. [1], referred to as the *literature model*). This does not predict *sleep* but does include nearly all other states we have in our dataset. It includes a total of 25 parameters.
- (2) an LSTM model per individual patient (*individual LSTM*).
- (3) a single LSTM model across all patients (*generic LSTM*).

To train the dynamical systems models (*GP model* and the *literature model*), we optimize the parameters for each patient individually, by applying NSGA-II on the training portion of the data of that patient, and select the model instance for the patient that minimizes the sum of the errors on the validation set. We do this for both the *in sample* and *out of sample* data. Note that the GP model itself has of course been generated using only the *in sample* data.

We have chosen an LSTM model as benchmark machine learning model as this has shown to work best on this type of data (see [14]). For the LSTM models we use the combination of the training and validation part of the patient data as training set. Six output neurons are used (we have three future time points we want to predict for two states). We have experimented with various parameter settings of the LSTM in some initial runs, which showed that taking a batch size of 7 combined with a single layers of 128 hidden neurons work best. We train for 30 epochs. Here, the *individual LSTM* is trained and applied per patient (both for the *in sample* and *out of sample* patients), while the *generic LSTM* is trained on all patients in the *in sample* data.

We measure the performance of all algorithms on the test set part of the data of each patient, both for the *in sample* and *out of sample* data. We perform only single runs of the algorithms per patient as the number of patients will allow us to tackle the stochasticity.

5 RESULTS

In the results, we first focus on the parameter settings of our GP-HD approach, followed by the resulting behavior of the algorithm and the best model we find. We then compare the performance of that model to the benchmark algorithms.

5.1 Parameter Settings and GP runs




Figure 3: Development of the fitness values over the generations over 10 runs, note that the y-axis does not start at 0.

When we consider the influence of the hyperparameter settings of NSGA-II on the fitness values found, we hardly see any difference between the different parameter settings we identified, neither

in terms of the absolute fitness value, nor in terms of the standard deviation. Hence, we select the cheapest option in terms of computation ($pop_{NSGA-II} = 5$ and $gen_{NSGA-II} = 5$). When we look at the hyperparameters of the GP, we see that the process converges before 50 generations independent of the population size. During the initial runs, the larger population size seemed slightly better, which we therefore use (so, $pop_{gp} = 100$ and $gen_{gp} = 50$). Figure 3 shows the median fitness values over 10 runs and the 25-and 75% inter quartile ranges. We see that the fitness nicely converges and obtain best fitness values of around 0.92.

The best model we obtain has a fitness value of 0.9372 (not remarkably higher compared to other models we observe in the same run or other runs of our framework). The model has the following specification:

$$s_{mood}^{t+1} = s_{mood}^t + \gamma_1 \cdot s_{sleep}^t \cdot \gamma_1 - s_{mood}^t \quad (6)$$

$$s_{sleep}^{t+1} = s_{sleep}^t \quad (7)$$

$$s_{worrying}^{t+1} = s_{enjoyed}^t - s_{social}^t \cdot s_{enjoyed}^t - s_{social}^t \cdot s_{social}^t \quad (8)$$

$$s_{social}^{t+1} = s_{worrying}^t - \gamma_1 - s_{social}^t \cdot s_{worrying}^t \cdot s_{sleep}^t + s_{worrying}^t - \gamma_1 - s_{social}^t \cdot s_{sleep}^t * s_{sleep}^t + \gamma_1 \quad (9)$$

$$s_{activities}^{t+1} = s_{social}^t \quad (10)$$

$$s_{enjoyed}^{t+1} = s_{self_esteem}^t \quad (11)$$

$$s_{self_esteem}^{t+1} = s_{mood}^t \quad (12)$$

The model only has one parameter (γ_1). Furthermore, we see that for *sleep* an extremely simple relationship is found (just take the previous value). While one could argue this does not show the benefit of our approach for this aspect, it does show that it does not generate unnecessarily complex models. For the *mood*, a more complex relationship is observed, drawing advantage of the previous values measured for both *mood* and *sleep*. A relationship between *mood* and *sleep* has been reported in the psychological literature (see e.g. [18]). When comparing the outcome to the literature model ([1] only focusing on mood), we see the literature model being much more complex and less insightful (with 25 parameters) while the prediction of mood depends on three factors that are not included in the resulting GP model: the *social* interactions, number of *activities*, and how much the patient *enjoyed* the activities, whereas the literature model does not use *sleep*. A remarkable difference that will be analyzed in more detail with clinical psychologists.

5.2 Performance Evaluation

Let us explore how well the model performs compared to other approaches. Table 3 shows the errors we obtain over the different algorithms, evaluation criteria and number of time steps ahead. We have statistically compared the difference in performance of the other algorithms compared to our GP-HD model using a Wilcoxon ranked sum test ($p = 0.05$), the results of the significance test are also shown in the table.

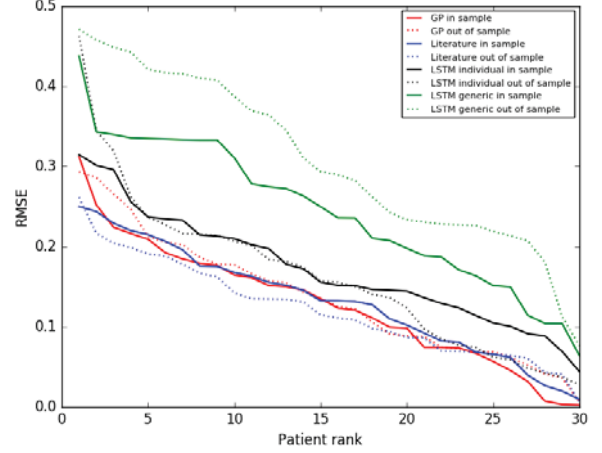


Figure 4: Ranked RMSE's for prediction of mood at time $t + 1$ for all algorithms (in and out of sample).

We observe that our approach is not significantly outperformed by any other approach, while it significantly outperforms the generic LSTM in all cases and the individual LSTM for the *sleep* state. The literature model and our GP-HD model perform equally well. Our approach also seems to generalize well (considering the performance on *out of sample* patients), especially for the *mood* state. Figure 4 shows the difference in performance over different patients for *mood* at $t + 1$ (other prediction intervals show similar patterns) for both the *in sample* and *out of sample* patients. We can observe similar patterns across all algorithms, though performance of the generic LSTM model is a lot poorer for unseen patients.

Finally, Figure 5 shows an example *out of sample* patient (the patient with the highest variation in performance scores) and the accompanying predictions for *mood* at time $t + 1$ using the different approaches. It can clearly be observed that the GP and literature model predict quite reasonable. The LSTM individual model however does follow the trends, but provides a prediction closer to the average value. This holds even more extreme for the generic LSTM model.

6 DISCUSSION

In this paper, we have presented GP-HD to generate dynamical systems models for predicting developments of health states over time. This approach tailored an existing GP approach (cf. [5]) using a fitness function based on an evaluation framework for dynamical systems models (cf. [20]). The research question we posed was: Is the predictive performance of models generated by GP-HD better than state-of-the-art data-driven and knowledge-driven approaches for a real-life case study in the health domain? Based on the results we have obtained for the mental health case studied in this paper we can answer this question with a partial yes. The approach is on par with a literature based dynamical systems model (while being a simpler and more insightful model), outperforms a generic LSTM model, and scores at least as well as individual LSTM models. Of course, it is hard to generalize these results over other datasets. The

Table 3: Median RMSE performances of algorithms. Scores in *italics* indicates that the performance is significantly worse ($p < 0.05$, Wilcoxon ranked sum test) than the performance of the GP, significantly better compared to the GP is not observed.

Algorithm	$t+$	RMSE (in sample)		RMSE (out of sample)	
		mood	sleep	mood	sleep
GP	1	0.129	0.030	0.130	0.078
	2	0.162	0.068	0.148	0.122
	3	0.173	0.093	0.169	0.158
Literature	1	0.132	-	0.113	-
	2	0.149	-	0.148	-
	3	0.174	-	0.172	-
LSTM ind.	1	<i>0.153</i>	<i>0.189</i>	0.156	<i>0.198</i>
	2	0.155	<i>0.190</i>	0.167	<i>0.207</i>
	3	0.154	<i>0.185</i>	0.184	<i>0.226</i>
LSTM gen.	1	<i>0.242</i>	<i>0.283</i>	<i>0.291</i>	<i>0.302</i>
	2	<i>0.239</i>	<i>0.295</i>	<i>0.265</i>	<i>0.283</i>
	3	<i>0.292</i>	<i>0.254</i>	<i>0.309</i>	<i>0.307</i>

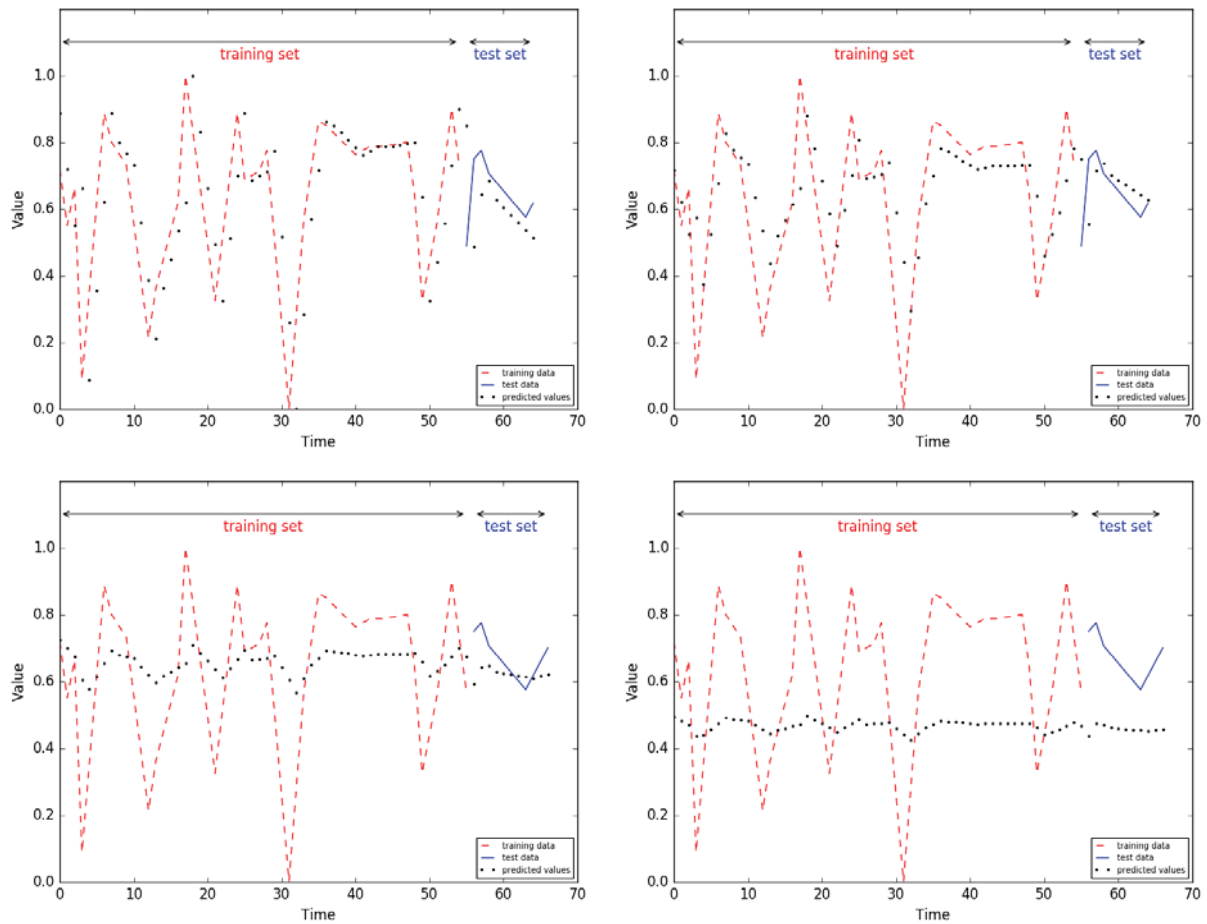


Figure 5: Illustration of predictions for mood at time $t + 1$ for the out of sample patient with the highest variation in performance (black dots are the predictions of the models). Note that the training set is the combination of the training and validation part of the data. Upper left: GP; upper right: literature, lower left: individual LSTM, lower right: generic LSTM.

relationships that are used in the resulting GP model are in line with literature in the psychological domain (e.g. [18]).

A lot of research has been devoted to data-driven predictive models for health. Specifically, a variety of LSTM based approaches have been proposed for modeling temporal data in the health domain (see e.g. [12]). In addition, ways to engineer temporal features are seen (e.g. [2, 10]). However, none make the combination we present in this paper: a generic model with parameters that can be tailored towards individuals. Of course, more GP-based approaches have been developed (e.g. [17]), but none are focused on the specific setting with multiple datasets (one per patient) we have. For the domain of mental health there are only few models that have been developed to provide more fine grained (e.g. daily) predictions of the mental health state, see e.g. [6, 11, 15]. Due to the differences in the characteristics of these groups as well as the measurements performed performances are difficult to compare.

For future work, we want to apply the proposed approach to other health datasets that are larger and contain more variability. We also want to explore the influence of the hyperparameter settings of the evaluation approach more rigorously. On top, we will focus on making the framework more efficient using racing (cf. [13]) and study the scalability of the approach in more depth. Finally, we want to study how we can improve the LSTM performance further. The *generic LSTM* only has access to the *in sample* data while the *individual LSTM* can only use the data of the specific patient. Studying a hybrid solution of the LSTM (e.g. using transfer learning) where we tailor the generic LSTM model based on a small portion of the data of a specific patient is therefore a next step we want to take.

ACKNOWLEDGEMENTS

Just one day after having received the notification of acceptance for this paper, our fellow author Jeroen Ruwaard (1970) unexpectedly passed away. This was a huge shock to us. Jeroen was a wonderful colleague who acted as a bridge between clinical psychologists and more tech driven scientists. We will remember Jeroen as a warm and social person who was always open for collaborations, willing to help, and was unique in that he could truly understand the line of thinking from both the clinical and technological side, and make excellent contributions to both domains. Besides that, we always valued his sense of humour that could really make meetings and trips abroad unforgettable. Our condolences go out to his wife and two children. It has been a pleasure for us to have had the opportunity to work with him over the past years and we will miss him dearly.

This research has been performed in the context of the EU FP7 project E-COMPARED (project number 603098). We want to thank all trial coordinators in the countries where the MoodBuster platform has been used.

REFERENCES

- [1] Altaf Hussain Abro and Michel CA Klein. 2016. Validation of a Computational Model for Mood and Social Integration. In *International Conference on Social Informatics*. Springer, 385–399.
- [2] Iyad Batal, Hamed Valizadegan, Gregory F Cooper, and Milos Hauskrecht. 2013. A temporal pattern mining approach for classifying electronic health record data. *ACM Transactions on Intelligent Systems and Technology (TIST)* 4, 4 (2013), 63.
- [3] Tibor Bosse, Mark Hoogendoorn, Michel CA Klein, and Jan Treur. 2007. An agent-based generic model for human-like ambience. In *European Conference on Ambient Intelligence*. Springer, 93–103.
- [4] Tibor Bosse, Mark Hoogendoorn, Michel CA Klein, and Jan Treur. 2009. A generic architecture for human-aware ambient computing. In *Agent-based ubiquitous computing*. Springer, 35–61.
- [5] Hongqing Cao, Lishan Kang, Yuping Chen, and Jingxian Yu. 2000. Evolutionary modeling of systems of ordinary differential equations with genetic programming. *Genetic Programming and Evolvable Machines* 1, 4 (2000), 309–337.
- [6] Darryl Daugherty, Tairi Roque-Urrea, John Urrea-Roque, Jessica Troyer, Stephen Wirkus, and Mason A Porter. 2009. Mathematical models of bipolar disorder. *Communications in Nonlinear Science and Numerical Simulation* 14, 7 (2009), 2897–2908.
- [7] Kalyanmoy Deb, Amrit Pratap, Sameer Agarwal, and TAMT Meyarivan. 2002. A fast and elitist multiobjective genetic algorithm: NSGA-II. *IEEE transactions on evolutionary computation* 6, 2 (2002), 182–197.
- [8] Agoston E Eiben and James E Smith. 2015. *Introduction to evolutionary computing*. Springer.
- [9] Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation* 9, 8 (1997), 1735–1780.
- [10] Mark Hoogendoorn and Burkhardt Funk. 2017. *Machine Learning for the Quantified Self: On the Art of Learning from Sensory Data*. Springer.
- [11] Robert LiKamWa, Yunxin Liu, Nicholas D Lane, and Lin Zhong. 2013. Moodscope: Building a mood sensor from smartphone usage patterns. In *Proceeding of the 11th annual international conference on Mobile systems, applications, and services*. ACM, 389–402.
- [12] Zachary C Lipton, David C Kale, Charles Elkan, and Randall Wetzell. 2015. Learning to Diagnose with LSTM Recurrent Neural Networks. *arXiv preprint arXiv:1511.03677* (2015).
- [13] Oded Maron and Andrew W Moore. 1997. The racing algorithm: Model selection for lazy learners. In *Lazy learning*. Springer, 193–225.
- [14] Adam Mikus, Mark Hoogendoorn, Artur Rocha, Joao Gama, Jeroen Ruwaard, and Heleen Riper. 2018. Predicting short term mood developments among depressed patients using adherence and ecological momentary assessment data. *Internet interventions* 12 (2018), 105–110.
- [15] Venet Osmani, Alban Maxhuni, Agnes Grünerbl, Paul Lukowicz, Christian Haring, and Oscar Mayora. 2013. Monitoring activity of patients with bipolar disorder using smart phones. In *Proceedings of International Conference on Advances in Mobile Computing & Multimedia*. ACM, 85.
- [16] Heleen Riper, Gerhard Andersson, Helen Christensen, Pim Cuijpers, Alfred Lange, and Gunther Eysenbach. 2010. Theme issue on e-mental health: a growing field in internet research. *Journal of medical Internet research* 12, 5 (2010).
- [17] Michael Schmidt and Hod Lipson. 2009. Distilling free-form natural laws from experimental data. *Science* 324, 5923 (2009), 81–85.
- [18] Dorthe Kirkegaard Thomsen, Mimi Yung Mehlsen, Søren Christensen, and Robert Zachariae. 2003. Rumination-relationship with negative mood and sleep quality. *Personality and Individual Differences* 34, 7 (2003), 1293–1301.
- [19] Jan Treur. 2016. *Network-Oriented Modeling and Its Conceptual Foundations*. In *Network-Oriented Modeling*. Springer, 3–33.
- [20] Ward van Breda, Mark Hoogendoorn, AE Eiben, and Matthias Berking. 2017. Assessment of temporal predictive models for health care using a formal method. *Computers in biology and medicine* 87 (2017), 347–357.