Individualized Optimal Behavioral Interventions by Predicting Relapse

Nada Aldarrab, Victor Ardulov, Biswarup Bhattacharya,
Su Lei, Han Ching Ou, Yilei Zeng

Abstract

There exists a number of prevailing therapies and behavioral interventions that are used to help individuals recovering from substance abuse. The Global Appraisal of Individual Needs (GAIN) [7] data set provides information over a large set of adolescents and emerging adults presenting diagnostically relevant demographic variables and reported substance abuse, and relapse frequency. The data are recorded over 4 reporting periods separated evenly over the course of a year (i.e. every 90 days). Of particular interest, GAIN also provides details regarding individuals' attempted treatments for substance abuse during the course of the year. Often treatment options provided are a result of availability rather than specific diagnostic decisions. Our aim is to (1) predict whether an individual will relapse given the treatment received and other demographic and feature information, (2) analyze the effectiveness of treatment options across particular demographic variables, to see if there exist correlations between participant backgrounds and the effectiveness of certain behavioral interventions. The eventual goal is to prescribe personalized treatments in an effort to maximize their likelihood of recovery and minimize their likelihood of relapsing.

1 Introduction

Implementing effective interventions that treat adolescent and young adult substance use is an important step to improve outcomes. Across the U.S., the quality of treatment service varies, and developmentally appropriate treatment is not always available [10]. Since the early 2000s, the U.S. Department of Health and Human Services, Substance Abuse and Mental Health Services Administration (SAMHSA) has encouraged the use of evidence-based treatments to address adolescent and young adult substance use. Nonetheless, recovery rates one-year post treatment are typically less than 50% [6].

The primary goal of this study is to determine which substance abuse treatment modalities (i.e., cognitive-behavioral therapy, motivational interviewing, contingency management, among others) are associated with better outcomes one-year post-treatment and particularly, whether specific treatment modalities are more effective at addressing certain substance use disorders (i.e., alcohol use disorder, cannabis use disorder, opioid use disorder, among others) and groups of individuals (geographic, age, etc.).

The discussion below begins with an analysis of the population represented in the data, and their relative treatment outcomes. It continues into discussing the process of cleaning and data preparation necessary to construct a number of models that will be used to predict the treatment outcome. Next, we introduce an unsupervised method for finding nonlinear embedding of patient data, which enables for more robust data processing, and missing data reconstruction. Finally, our conclusion outlines methods in which these technologies and models can be used and interpreted to for more individualized prescription of behavioral interventions.

2 Background

2.1 Data

Global Appraisal of Individual Needs (GAIN) is one of the largest national datasets of adolescent substance use treatment [7]. The survey is administered by treatment staff at treatment baseline as a bio-psycho-social

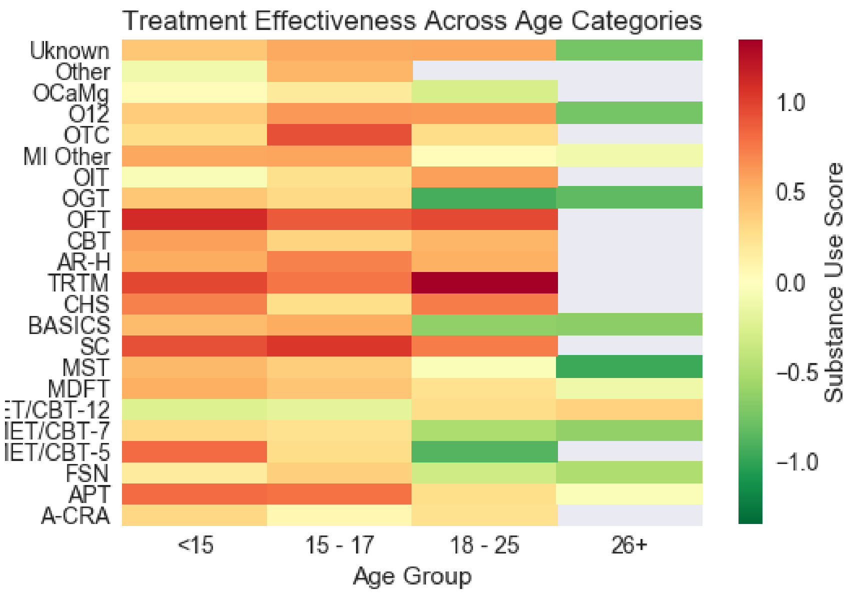

Figure 1: Heatmap representing how varying the age has different expected outcomes based on original treatment types

clinical assessment tool (GAIN-I). The GAIN M90 is administered at 3 months post-treatment and again at 6 and 12 months. The follow-up surveys utilize a time line follow back procedure where participants reflect on the past 90 days, essentially providing a full calendar year of data on each individual. The GAIN has a total sample size of $N = 32452$ individuals ($n = 23436$ adolescents; $n = 4728$ young adults; $n = 9061$ adults).

2.2 Preparation

For the purposes of our modeling and experiments, we concentrated on the adolescent population since they form the majority in the dataset and are the highest priority group due to their young age.

A large portion of the data had missing values for many variables due to the methodology of data acquisition. Missing values were expected as a result of the unbalanced nature study design in which the data collection period began upon entering treatment, but individuals varied in the amount of time passed from baseline times and the end of data collection. Thus, the majority of missing data are explained by censoring - that is, individuals who did not have an opportunity to provide data.

Due to the nature of our study, for a portion of our experiments we chose not to consider the individuals who have missing data for the experiments. This can be handled in the future by using smoothing techniques to replace the missing data [9, 8]. The data was processed to remove features with overwhelming amounts of missing data leaving only 123 features across $n = 11668$ participants. Furthermore the relapse value was condensed to binary classes $\{0, 1\}$ representing a relapse in the year following the initial baseline questions.

In Figure 1 we can see the how age can effect the likelihood of positive outcome for a particular therapy. This is suggestive that not all individuals will respond to the treatments in the same manner. However many of the treatments initially presented fall under larger classes. We elect to use these classes as our label for treatment, since it generalizes to an approach. Furthermore some specific treatments are highly specific to a very small group of participants and yield very low confidence in an interpolated results. As a result the the study uses a simplified index of 7 treatment classes:

- Adolescent Community Reinforcement Approach/Assertive Continuing Care (ACRA-ACC)
- Cognitive Behavioral Therapy (CBT)
- Seven Challenges

- Multidimensional Family Therapy for Adolescent Drug Abuse (MDFT)
- Other Evidence-Based Therapy (EBT)
- Specific Manualized Program
- Other

2.3 Related Work

Previous work on fair and efficient prescription and resource allocation for homeless youth has been studied in the context of prioritizing housing options for intermediate housing [1]. A primary contrast between the approaches is that in housing allocation there is a limited resource constraint that must be optimized, which is not present in substance abuse treatment options.

Our study draws on the use of treatment decision for medical diagnosis and prescription. By taking advantage of statistical guarantees in machine learning algorithms, it is possible to train neural network models to optimize for individual treatment effectiveness rather than population averages [12]. This enables doctors to prescribe medications and treatments to patients which are optimal for the specific background of the patient, taking into consideration the diverse set of features and history that a patient may present.

Due to the nature of the studies, and the large amounts of missing data, we also look to leverage previous results in medical record imputation. By utilizing an unsupervised embedding learning, information otherwise missing from health records can be interpolated by first project the present data into a subspace and the re-projecting back onto the original space [2].

3 Methods

3.1 Mixed Effects Model

Mixed-effects models [13] (or mixed models) are statistical models that incorporate both fixed-effects parameters and random effects. They include additional random-effect terms, and are often appropriate for representing clustered, and usually correlated, data. Examples include, hierarchical data collections, observations taken for various individuals, or data gathered over time for the same individuals. Model equations can be expressed as: Outcome(dependent variable) $\sim 1 +$ fixed effects $+$ random effects.

In this model, MixedLM is used to perform a mixed effects analysis on the relationship between the 123 features and the relapse of each individual. The features, excluding the outcome and the treatment type, were modeled with fixed effects. For the Random effects, the intercepts were for treatment type and age groups, as well as by-treatment and by age group random slopes for the effect of relapse. P-values were obtained by likelihood ratio tests of the full model with the effect in question against the model without the effect in question.

3.1.1 Preliminary Study

First, a visual inspection is performed on the residual plots of features. This analysis did not initially reveal any obvious deviations from homoscedasticity or normality.

PCA (Principle component analysis) is implemented to give an visual clue of the 7 treatment types' innate pattern. Since neither, 2 component nor, 3 component PCA analysis were able to reveal valuable intuition, mixed models are then carried out for smaller granularity. The result of PCA analysis is presented below, 2 component PCA can be seen in Figure 2a, while 3 component PCA is in Figure 2b.

3.1.2 Individual Domain Results

The preprocessed 123 features are categorized into 12 domains including the relapse domain, with each domain name listed below. Now we represent the demographic domain's result as an example.

1. Demographics

Demographics domain result is shown in Figure 3. Results shows that four features including 'female','nonwhite','B2a_0' and 'POPIgrp' have shown significant fixed effects. Feature 'AGEGM4' means

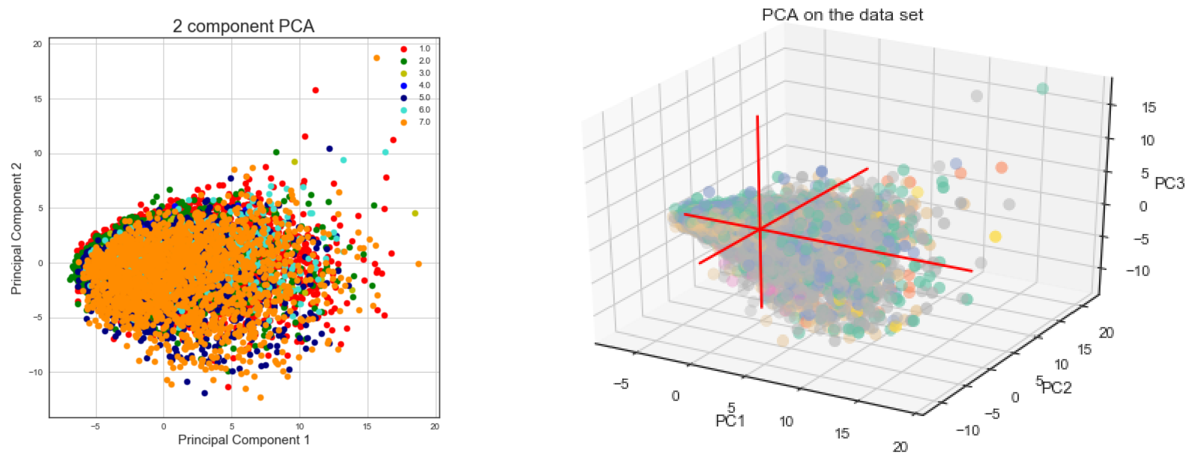

(a) 2 Component PCA

(b) 3 Component PCA

Figure 2: **Comparison of *2 Component* vs. *3 Component*PCA.** (a) 2 Component PCA and (b) 3 Component PCA

that all individuals in the data set are categorized into 4 age groups. 'B2a_0' feature is the exact age feature, while 'POPIgrp' represents SES information. The coefficient of the first three features have indicated negative effect on relapse, while for "POPIgrp" coefficient revealed positive effect on relapse. "Intercept RE","AGEGM4 RE" and "Intercept RE*AGEGM4" are random slope for treatment type,random slope for age groups, as well as by-treatment and by age group random slopes for the random effect of relapse.

```
              Mixed Linear Model Regression Results
==================================================================
Model:                 MixedLM     Dependent Variable:     Relapse
No. Observations:      11668       Method:                 REML
No. Groups:            7           Scale:                  0.2003
Min. group size:       131         Likelihood:             -7226.9023
Max. group size:       5347        Converged:              Yes
Mean group size:       1666.9
------------------------------------------------------------------
                         Coef.  Std.Err.   z     P>|z| [0.025 0.975]
------------------------------------------------------------------
Intercept                 0.817    0.084  9.695 0.000  0.652  0.982
female                   -0.063    0.010 -6.502 0.000 -0.082 -0.044
nonwhite                 -0.056    0.009 -6.188 0.000 -0.074 -0.038
B2a_0                    -0.009    0.003 -2.630 0.009 -0.016 -0.002
AGEgm4                   -0.002    0.025 -0.072 0.943 -0.050  0.046
POPIgrp                   0.012    0.001  8.576 0.000  0.009  0.014
mstat                     0.009    0.013  0.719 0.472 -0.016  0.035
V1_0                      0.004    0.003  1.251 0.211 -0.002  0.011
V7_0                     -0.008    0.005 -1.454 0.146 -0.018  0.003
Intercept RE              0.015    0.022
Intercept RE x AGEgm4 RE  0.001    0.007
AGEgm4 RE                 0.002    0.004
==================================================================
```

Figure 3: Demographics Domain Model: Relapse $\sim$ Demographics features + (1|treatment_new) + (1|age_group)

Other domains are: 2. Environment 3. Substance Use 4. Family Health 5. Mental Health 6. Victimization 7. Physical Health 8. Risk Behaviors 9. Social Support 10. Criminal Activity 11. Treatment 12. Relapse. To succinctly represent our results we only show one domain's result as an example.

	Coef.	Std.Err.	z	P>|z|	[0.025	0.975]
Intercept	0.354	0.061	5.772	0.000	0.234	0.475
female	-0.057	0.010	-5.805	0.000	-0.076	-0.038
nonwhite	-0.047	0.009	-5.338	0.000	-0.065	-0.030
B2a_0	-0.003	0.003	-1.256	0.209	-0.009	0.002
POPIgrp	-0.005	0.002	-2.158	0.031	-0.010	-0.000
V1_0	0.004	0.003	1.217	0.224	-0.002	0.010
algrps	0.111	0.011	10.296	0.000	0.090	0.132
mjgrps	0.104	0.013	8.041	0.000	0.079	0.130
ccgrps	0.021	0.011	1.939	0.053	-0.000	0.041
DSS9_0	0.007	0.002	3.281	0.001	0.003	0.011
dldiag	0.047	0.009	5.097	0.000	0.029	0.065
TSS_0	-0.004	0.002	-2.499	0.012	-0.008	-0.001
EPS7p_0	0.001	0.000	2.607	0.009	0.000	0.001
E9e18	0.027	0.014	1.980	0.048	0.000	0.054
HDS_0	0.004	0.004	1.101	0.271	-0.003	0.011
HPS3p_0	0.001	0.000	2.351	0.019	0.000	0.001
SxRS_0	0.011	0.003	4.430	0.000	0.006	0.016
txtypen	0.000	0.000	0.452	0.651	-0.001	0.001
S6	0.007	0.010	0.733	0.464	-0.012	0.027
S8h_0	0.056	0.012	4.797	0.000	0.033	0.078
TxPI_0	0.010	0.002	5.160	0.000	0.006	0.014
S2a1_0	0.002	0.000	5.695	0.000	0.001	0.003
S2c1_0	0.002	0.000	13.634	0.000	0.002	0.002
Intercept RE	0.012	0.018				
Intercept RE x AGEgm4 RE	0.002	0.004				
AGEgm4 RE	0.001	0.002				

Figure 4: Aggregated Domains Results: Relapse $\sim$ selected significant features + (1|treatment_new) + (1|age_group)

3.1.3 Aggregated Domains Results

We selected features which presented significant fixed effects in each individual domain's result described in the previous section. Our aggregated model is then built based on these selected features. Result is shown in Figure 4. All the significant feature in sequence indicates that gender, ethnicity, SES, Alcohol use problem, Marijuana use problem, Cocaine/Crack use problem, Depressive Symptom Scale, Dual diagnosis (both substance abuse and mental health problem),Traumatic Stress Scale, Emotional Problems Scale, GVS: Age first time abused, Health Problem Scale,Sex Risk Scale,TMI: Will need to come back to Tx 1/more times,Treatment Pressure Index,How many days used any alcohol (baseline),How many days used cannabis (baseline) have significant fixed effects on relapse.

3.2 Random Forest

Decision trees are considered to be simple and easily interpretable, but they have a poor predictive performance and poor generalization on the test set. Random forests are a popular ensemble technique which improves the predictive performance of decision trees by averaging across a number of different decision tree models, effectively reducing the variance of the prediction [3]. We developed a random forest classifier to predict the relapse variable.

3.3 Support Vector Machine

In machine learning, support vector machines (SVMs) are supervised learning models that analyze data and provide output for classification and regression analysis [5]. SVM is a non-probabilistic binary linear classifier, such that it classifies examples to be falling into one of two classes. The SVM margin, which is the separation between the two classes, is trained to be as optimal as possible. New examples are then mapped into that same training space and predicted to belong to one of the classes depending on which side of the hyperplane they fall in. We trained our SVM classifier with state of the art libSVM [4] toolbox, which support parameter optimization and feature selection, to predict whether an individual will relapse or not given the input features.

3.4 Multi-Layer Perceptron (MLP)

A multilayer perceptron is a feed-forward artificial neural network model that has one layer or more of hidden units and nonlinear activations [11]. We attempted two architectures (50×10, 5×2) and compared their performance in predicting relapse.

3.5 Autoencoders

Extracting decisions from data can often be done when the data is projected down into a nonlinear subspace. This can combine many of the variable and extract a more noise invariant vectorization of the data. One technique to accomplish this is through the use of an autoencoder.

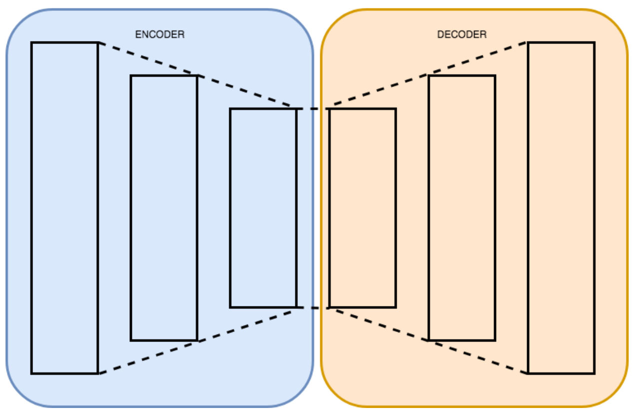

Figure 5: Visualization of Autoencoder architecture

The autoencoder will is a neural network that is used to approximate the identity function. The network first projects the data down to a lower dimension through its "encoding' and then takes this lower dimensional input and will try "decode" it back into the original input. Figure 5 depicts the architecture, showing the 2 major components of the neural network. After training the neural network to convergence, the decoder can be removed and the encoder now represents a nonlinear subspace that the data can still be accurately reconstructed from.

This can be used to more compactly represent the initial data set. In particular in this study he original 128 features were projected down to a 32 dimensional feature space before using a simplified MLP trained on the projections to identify the relapse outcome.

The autoencoder also provide a level of interoperability often lost with other deep networks. The encoder robustness to noise also manifests in that the projection can still be constructed with missing values [2], the missing data can be reconstructed or interpolated through projections.

4 Experiments and Results

4.1 Nonlinear Models to predict relapse

All categorical variables were one-hot encoded. After encoding the categorical variables, there were 1462 input features. We consider relapse to have occurred when an individual relapses during any of the months after the baseline. All results are averaged over 20 simulations and are shown in Table 1.

	Accuracy	**Precision**	**Recall**	**F1-score**
SVM	**0.766**	0.767	0.935	**0.843**
RF	0.741	0.750	**0.958**	0.841
MLP (50×10)	0.665	**0.783**	0.773	0.778
MLP (5×2)	0.711	0.734	0.833	0.780

Table 1: Performance of non-linear models in predicting relapse

4.1.1 Random Forest

We developed a random forest classifier to predict the relapse variable. The top five features which were the most important in making the predictions are shown in Table 3.

Rank	Variable	Description	Domain
1	hmlsrisk_0	Homelessness Risk	Environmental
2	s2s1d_0	1+ days any substance use (baseline)	Post-Treatment
3	s2a1_0	How many days used any alcohol (baseline)	Post-Treatment
4	amgrps	Amphetamine use problem	Substance Use
5	ESI_0	Environmental Strengths Index	Social support

Table 2: Most important features in predicting relapse (RF)

We observe that the most important features in predicting relapse is in fact *homelessness*, and it is closely followed by substance use factors.

4.1.2 Support Vector Machine (`libsvm`)

We trained our SVM classifier to predict whether an individual will relapse or not given the input features. We utilized the `libsvm` toolbox that automatically scales the data, optimizes parameters in the kernel function using cross validation, and supports feature selection, visualizations and various other tools.

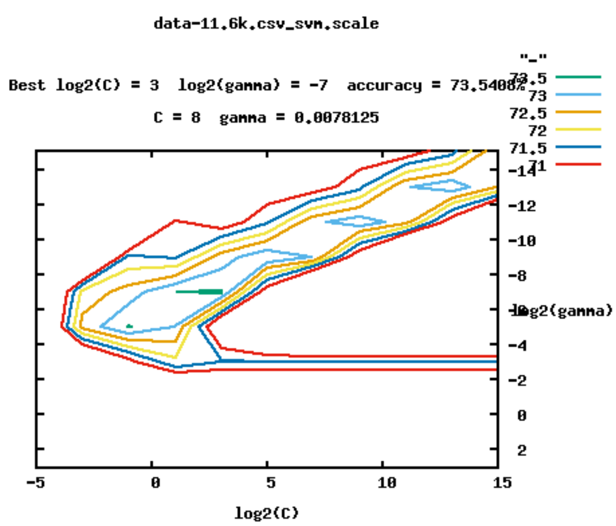

Figure 6: LibSVM optimizing parameters

4.1.3 Multi-Layer Perceptron (MLP)

We also attempted two architectures (50×10, 5×2) and compared their performance in predicting relapse. The results are shown in Table 1.

Rank	Variable	Description	Domain
1	hmlsrisk_0	Homelessness Risk	Environmental
2	s2s1d_0	1+ days any substance use (baseline)	Post-Treatment
3	s2a1_0	How many days used any alcohol (baseline)	Post-Treatment
4	IDS_0	Inattentive Disorder Scale	Mental Health
5	TSS_0	Traumatic Stress Scale	Mental health

Table 3: Most important features in predicting relapse (SVM)

The most important features in predicting relapse identified by these models are found to be:

1. Homelessness risk
2. Previous substance use
3. Previous alcohol use
4. Features related to environmental support and/for mental health

4.2 Autoencoder

In order to construct an autoencoder for this data set, the data is encoded using 2 hidden layers with 64 and 32 ReLU activated units respectively, then decoded with 2 more hidden layers with 32, and 64 ReLU activated units respectively as well. Using mean-squared error and K-fold cross validation($K = 5$) training regiment, the reconstruction accuracy was driven to approximately 91%.

With the autoencoder constructed and trained the encoder was removed from the model, the data is passed through the encoder and represented in its 32-dimensional subspace. The projected data point was then concatenated with the treatment type and used to predict the relapse outcome with a single 5 unit hidden-layer MLP. Specifically, a 90%, 10%, 10% split was take of data for training, validation, and testing respectively. Table 4 shows the results for using the projected test data in order to predict the relapse outcome.

Metric	Score
Accuracy	0.74
Precision	0.89
Recall	0.77
F1- Score	0.83

Table 4: Relapse Prediction Results using encoded feature space

5 Conclusion

This work has presented a number of methods demonstrating promise in using neural architectures to accurately predict patient relapse outcome. The models have successfully utilized relationships in the patient background information and treatment options indicative of prolonged substance abuse.

Contributed is a methodology for representing patient information in nonlinear subspaces using a unsupervised feature selection and embedding. This enables future reconstruction of missing data and greater robustness to noise in the data

Finally, it can be seen from the presented results that these models are effective at predicting client outcomes and substance abuse relapse. These results have the potential to enable social workers and behavioral therapists the potential to evaluate effective strategies on a individual basis. This opens up a realm of opportunities to more effectively prescribe individualized therapy type.

References

[1] Mohammad Javad Azizi, Phebe Vayanos, Bryan Wilder, Eric Rice, and Milind Tambe. Designing Fair , Efficient , and Interpretable Policies for Prioritizing Homeless Youth for Housing Resources. 2018.

[2] Brett K. Beaulieu-Jones and Jason H. Moore. Missing Data Imputation in the Electronic Health Record Using Deeply Learned Autoencoders. *Biocomputing 2017*, pages 207–218, 2016.

[3] Leo Breiman. Random forests. *Machine learning*, 45(1):5–32, 2001.

[4] Chih-Chung Chang and Chih-Jen Lin. LIBSVM: A library for support vector machines. *ACM Transactions on Intelligent Systems and Technology*, 2:27:1–27:27, 2011. Software available at http://www.csie.ntu.edu.tw/~cjlin/libsvm.

[5] Corinna Cortes and Vladimir Vapnik. Support-vector networks. *Machine learning*, 20(3):273–297, 1995.

[6] Michael L Dennis, Ya-Fen Chan, and Rodney R Funk. Development and validation of the gain short screener (gss) for internalizing, externalizing and substance use disorders and crime/violence problems among adolescents and adults. *The American Journal on Addictions*, 15(s1), 2006.

[7] Michael L Dennis, Janet C Titus, Michelle K White, Joan I Unsicker, and D Hodgkins. Global appraisal of individual needs: Administration guide for the gain and related measures. *Bloomington, IL: Chestnut Health Systems*, 2003.

[8] Craig K Enders. Analyzing longitudinal data with missing values. *Rehabilitation psychology*, 56(4):267, 2011.

[9] Craig K Enders and Deborah L Bandalos. The relative performance of full information maximum likelihood estimation for missing data in structural equation models. *Structural equation modeling*, 8(3):430–457, 2001.

[10] Jeremy Gans, Mathea Falco, Bruce R Schackman, and Ken C Winters. An in-depth survey of the screening and assessment practices of highly regarded adolescent substance abuse treatment programs. *Journal of child & adolescent substance abuse*, 19(1):33–47, 2009.

[11] David E Rumelhart, Geoffrey E Hinton, and Ronald J Williams. Learning internal representations by error propagation. Technical report, California Univ San Diego La Jolla Inst for Cognitive Science, 1985.

[12] Jeremy Weiss, Finn Kuusisto, Kendrick Boyd, Jie Liu, and David Page. Machine Learning for Treatment Assignment: Improving Individualized Risk Attribution. *Proceedings of the American Medical Informatics Association*, pages 1306–1315, 2015.

[13] Bodo Winter. Linear models and linear mixed effects models in r with linguistic applications. *arXiv preprint arXiv:1308.5499*, 2013.