Generating Accurate Personalized Predictions of Future Behavior: A Smoking Exemplar

Aaron J. Fisher and Peter D. Soyster

University of California, Berkeley

Preprint

Manuscript Under Review

Abstract

Objective: The present study sought to apply statistical classification methods to idiographic time series data in order to make accurate future predictions of behavior.

Methods: We recruited 70 individuals who presented as regular smokers; 52 completed experience sampling method (ESM) data collection and provided sufficient time series data. Time stamps from ESM surveys were used to calculate the time of day, day of the week, and continuous time—where the last datum was, in turn, used to calculate 12-hr and 24-hr cycles. Each individual's time series was split into sequential training and testing sections, so that trained models could be tested on future observations. Prediction models were trained on the first 75% of the individual's data and tested on the last 25%. Predictions of future behavior were made on a person by person basis. Two prediction algorithms were employed, elastic net regularization and naïve Bayes classification.

Results: Sample-wide area under the curve was nearly 80%, with some models demonstrating perfect prediction accuracies. Sensitivity and specificity were between 0.78 and 0.81 across the two approaches.

Conclusions: In addition to supporting the prediction accuracy of our models with out-of-sample tests in future data, each model was time-lagged, such that each prediction was for the subsequent measurement. Such a system could be the basis for mobile, just-in-time interventions for substance use, as models that accurately predict future behavior could ostensibly be used for delivering personalized interventions at empirically-indicated moments of need.

Keywords: machine learning; statistical classification; prediction; just-in-time; substance use; tobacco; idiographic; time series.

Personalized and precision medicine are concepts that have received a great deal of interest in recent years (Hamburg & Collins, 2010; Iyengar, Altman, Troyanskya, & FitzGerald, 2015; Schork, 2015), including attention from American (Ashley, 2015) and European (Dzau, Ginsburg, Van Nuys, Agus, & Goldman, 2015) funding initiatives. Although the initial impetus for personalized medicine was to match interventions to patient genotypes (Hamburg & Collins, 2010), after more than 20 years of genomic research in psychiatry, “a sobering picture...reveals daunting genetic and phenotypic complexity” (Hyman, 2018; p. 1), with no reliable candidate genes for common conditions such as depression (Border et al., 2019). Conversely, psychosocial variables such as affect, cognition, behavior, and context may provide rich sources of person-level data for generating personalized interventions (Fisher et al., 2019). Some have argued that passive data approaches—such as digital phenotyping (Insel, 2017)—will be sufficient for designing and implementing these interventions.

However, for behavioral processes, optimal personalization will likely require knowledge of cause, context, *and* timing. That is, detection algorithms and early warning signals may predict when a problem behavior is likely to occur, but do not necessarily provide information about the precipitants and catalysts of the behavior, or the most viable targets for intervention and remediation. For instance, one of the more promising possible avenues for providing personalized care is the just-in-time intervention (Bae, Chung, Ferreira, Dey, & Suffoletto, 2018). The just-in-time model takes advantage of the data from—and access to—mobile phones, by providing intervention materials to individuals in real time, when they are most needed; however, such interventions require theoretical and empirical knowledge of the predictors and correlates of problem behaviors in order to be maximally effective (Nahum-Shani, Hekler, & Spruijt-Metz, 2015; Spruijt-Metz & Nilsen, 2014).

If we take the position that human behavior is predictable, then the problem of making accurate predictions of future behavior becomes a function of identifying and employing the correct set of predictors for a given behavior. Additionally, given compelling evidence that the statistical relationships among variables are almost certain to vary from person to person (Fisher, Medaglia, & Jeronimus, 2018), it is fair to assume that the number and nature of predictors will also vary across persons. This poses a unique challenge: For any target behavior, each individual will have a unique set of predictors, requiring interventionists to construct prediction models on a case-by-case basis. Thus, when designing just-in-time interventions for personalized deployment, we believe that it is important to consider the complete set of possible prediction variables—what the fields of machine learning and pattern-recognition term the feature space—as well as ways to efficiently and effectively reduce the dimensionality of the full space to the person-relevant predictors for each individual. In the present study, we explore a set of procedures for (a) defining a large set of potential predictors, (b) utilizing variable selection methods for dimension reduction, and (c) generating predictions, case-by-case, via personalized prediction models. Specifically, we were interested in applying statistical classification methods to idiographic time series in order to make accurate future predictions of individual behavior.

One area for which behavioral predictions may be useful for creating effective just-in-time interventions is in substance use (Bae et al., 2018). Addressing cravings, social triggers, and other psychosocial influences *in situ* could amplify intervention mechanisms by delivering them when they are needed most. Presently, there are several hundred mobile apps for combating substance use (Hoeppner et al., 2017). Unfortunately, few of these have leveraged empirical evidence, theory, or accepted treatment guidelines in their design (Nahum-Shani et al., 2017). We recruited a sample of cigarette smokers and used experience sampling methods (ESM) to

measure temporal and psychosocial influences on their smoking behavior. Our hope is to provide an empirical basis for creating personalized just-in-time interventions via prediction models that accurately reflect the predictors and temporal coordinates of smoking behavior for each individual.

Tobacco use is the single leading cause of preventable death and disease worldwide, resulting in an estimated 2.5 million deaths globally and 438,000 deaths in the US per year (Siegel, Miller, & Jemal, 2018). Every year, more Americans die from cigarette use than from alcohol use, car accidents, suicide, AIDS, and illegal drug use, combined (Siegel et al., 2018). Cigarette smoking is related to at least 12 types of cancer, and is responsible for approximately one third of all cancer deaths, and almost 90 percent of lung cancer deaths (Warren, Alberg, Kraft, & Cummings, 2014). Among individuals who have already received a cancer diagnosis, cigarette smoking increases all-cause mortality, cancer-specific mortality, risk for second primary cancers, and risk of disease recurrence (Warren et al., 2014). The economic cost of cigarette use is substantial. It has been estimated that cigarette use accounts for approximately \$167 billion in healthcare expenditures and productivity losses annually (Xu, Bishop, Kennedy, Simpson, & Pechacek, 2015). In light of the mortality and impairment associated with tobacco use, reducing the number of people who smoke is a major goal in behavioral and public health. Developing methods to reduce cigarette use would help to reduce substantial personal, social, and financial harms.

Research examining the predictors of tobacco use and other health risk behaviors has predominantly utilized between-subject designs. Yet, between-subjects research designs are poorly-matched for personalized predictions of individual behavior (Fisher et al., 2018). Tobacco use is a behavior that (a) is enacted by an individual, (b) occurs within the context and timing of

that individual's life and (c) is dictated by idiosyncratic predictor variables that vary person by person. Thus, in order to generate models that can accurately predict an individual's behavior, those models must (a) measure the individual over time, (b) accurately represent the *shape* of the behavior's temporal variation (e.g. cyclic, based on time-of-day or day-of-the-week), and (c) incorporate the specific contextual variables that correspond to the behavior, such as thoughts, emotions, and physical/physiological experiences, as well as external cues, triggers, and circumstances. Because these will vary considerably from person to person, methods should be employed that can likewise measure and model the specific set of predictors on a person-by-person basis.

The present study used idiographic ESM data to generate personalized prediction models for individual smoking behavior via elastic net regularization (Zou & Hastie, 2005) and naïve Bayes classification (Hand & Yu, 2001)—statistical classification procedures sometimes categorized as machine learning (Christodoulou et al., 2019; Lewis, 1998). Elastic net is a regularized regression method that produces sparse models due to L_1 penalization. In addition to protecting against overfitting, this feature is useful for variable selection—coefficients shrunk to zero are omitted from the model. This may make it an optimal approach for the present aims. Ideally, onboard prediction models for just-in-time interventions will require the fewest possible inputs from participants, lowering participant burden and increasing the likelihood of ongoing engagement. Thus, limiting the number of predictors is paramount to an efficient, real time prediction system. Conversely, while naïve Bayes is one of the oldest and most reliable classification approaches (Hand & Yu, 2001), it does not perform variable selection (i.e. feature selection). Thus, the present study utilized a random forest approach for feature selection prior to performing naïve Bayes classification (see Method for more detail).

Our presupposition was that we would be able to return accurate prediction models for a majority of the participants in the study. With adequate attention given to the particular demands of idiographic time series—namely the idiosyncrasies in temporal fluctuations and psychosocial predictors—the selected prediction algorithms should perform well, given historical precedents. Thus, we expected to return strong accuracy (as reflected by the area under the curve), sensitivity, and specificity for a majority of participants. Finally, we had an exploratory aim to examine the degree to which baseline variables, measured prior to ESM data collection, would predict individual differences in accuracy, sensitivity, or specificity.

Method

Participants

The present data were drawn from individuals enrolled in an ongoing research study examining person-specific predictors and correlates of tobacco use. Seventy participants initiated involvement in the present study, with 52 providing sufficient data for prediction modeling. Retained participants ($N=52$) were adults ($M = 25.63$, $SD = 11.10$, range = 19 - 62) living in the San Francisco Bay Area who self-identified as current tobacco users. Nineteen participants (37%) identified as female and 32 identified as male, with 1 person who preferred not to disclose. Participants were diverse with respect to race/ethnicity (33% white, 2% Black or African American, 38% Asian/Pacific Islander, 23% mixed or ‘other’, with 4% who preferred not to disclose) and sexual orientation (75% heterosexual, 6% homosexual, 12% bisexual/queer, 6% other orientation, with 2% who preferred not to disclose). Participants smoked an average of 4.22 cigarettes per day ($SD = 4.53$, range = 0 - 20) for an average of 8.15 years ($SD = 10.24$, range = 1 – 48). Twenty-seven percent indicated that they were seriously thinking about quitting smoking in the next six months, with 6% indicating they were thinking about quitting in the next

30 days. A comparison of demographic information between retained and excluded participants can be found in Table 1.

Procedure

All study procedures were approved by the University of California, Berkeley Committee for the Protection of Human Subjects and all participants provided informed consent prior to participation

Potential participants were recruited from advertisements on Craigslist, flyers posted in the community, and through an undergraduate research participation pool. Interested individuals were directed to an online screening survey designed to confirm eligibility for study inclusion. To be included in the study, participants were required to be adults with English-language proficiency, have regular access to a web-enabled mobile phone, report having smoked 100 or more cigarettes in their lifetime, and report currently smoking ≥ 1 cigarette per week. Neither desire nor motivation to quit smoking was required to participate.

Eligible participants were invited to present to our lab at the University of California, Berkeley for enrollment and study baseline procedures. After completing consent procedures, participants' self-reported smoking status was biochemically verified using exhaled CO. Participants then completed a computer-based baseline survey, and were introduced to and instructed to complete 30 days of web-based surveys designed to be completed on the participants' mobile phone. After completing the mobile phone assessment period, participants completed a web-based survey assessing their experience completing the mobile phone surveys. Participants who completed at least 80% of the mobile phone surveys were reimbursed with \$25 or course credit.

Baseline Measures

Measures from the baseline survey relevant to the present analysis included participant demographics; measures of current and past tobacco use, cigarette dependence, smoking expectancies, motivation to quit, and internalized smoking stigma; DSM-IV psychiatric diagnosis, and severity of depression, anxiety, and stress.

Demographics. Demographic measures included participant age, gender identity, biological sex, racial/ethnic identity, perceived socioeconomic status, income, romantic partner status, and housing status.

Tobacco Use. Participants were asked to report on the number of cigarettes they had smoked in the past 24 hours, the past week, and their usual number of cigarettes per day. Information about the number of years participants had smoked and previous quit attempts were also collected. The Fagerström Test for Cigarette Dependence (FTCD; Heatherton, Kozlowski, Frecker, & Fagerström, 1991) assessed severity of cigarette dependence. Participants were asked to specify any goals regarding changing their cigarette use. The Stages of Change Scale (DiClemente et al., 1991) assessed intention to quit smoking, classifying smokers into precontemplation (no intention to quit), contemplation (intention to quit in the next 6 months), or preparation (intention to quit in the next 30 days, with at least one quit attempt lasting a minimum of 24 hours in the past year). Participants' beliefs about the effects of smoking and smoking cessation were assessed using the Brief Smoking Consequences Questionnaire (Rash & Copeland, 2008), the Smoking Abstinence Questionnaire (Hendricks, Wood, Baker, Delucchi, & Hall, 2011), and the Smoking Expectancies Questionnaire (Branstetter, Mercincavage, Dino, & Horn, 2015). The extent to which participants identified as a smoker was assessed using the

Smoker's Identity Scale (Dupont et al., 2015), and the Internalized Stigma of Smoking Inventory (Brown-Johnson et al., 2015) assessed smoking stigma.

Psychological and Emotional. An electronic version of the Mini International Neuropsychiatric Interview (MINI; Sheehan et al., 1998) assessed whether participants met diagnostic criteria for DSM-IV (APA, 1994) psychiatric disorder and screened for substance use and past exposure to trauma. The personality domains of neuroticism, extraversion, consciousness, agreeableness, and openness to experience were assessed using the NEO Personality Inventory: Five Factor Inventory (NEO-FFI; Costa & McCrae, 1992). Past week experiences of mood- and anxiety-related symptoms were assessed using the Depression Anxiety and Stress Scale-42 Item (DASS; Lovibond & Lovibond, 1995) and the Overall Anxiety Severity and Impairment Scale (OASIS; Norman et al., 2006).

Mobile Phone Survey Measures

Survey development. Prior to the initiation of the present study, we pursued a systematic procedure for creating mobile phone surveys, utilizing input from the extant literature and community stakeholders (Soyster & Fisher, 2019). First, our research team completed a systematic review of the literature on predictors of smoking behavior in order to identify commonly targeted constructs. Then, we reviewed research on tobacco use that utilized ESM and related methodologies, in order to generate an exhaustive list of potential items and item wordings. Next, a series of focus groups with current smokers were conducted in order to solicit stakeholder feedback on the number, content, and specific wording of potential survey items. Focus group discussions were transcribed and transcripts were subjected to qualitative and quantitative analyses, in order to identify convergent themes and word usage. Finally, a set of

survey items was selected that best-fit the extant literature and stakeholder feedback. For greater detail and explication, please see Soyster and Fisher (2019).

Survey procedure. At study baseline, participants' personal mobile phone numbers were entered into a secure web-based mobile survey system. Following the baseline appointment, participants were repeatedly assessed for 30 days via mobile phone survey, completed at 4 random intervals during their reported wakeful hours. Notifications were sent to participant mobile phones as a text message containing a hyperlink to a web-based survey. Each survey asked participants to rate their current experience of a number of variables using a visual analog slider ranging from 0 (*not at all*) to 100 (*as much as possible*).

Each survey assessed the number of cigarettes smoked since the previous survey, as well as symptoms of nicotine withdrawal (*down/depressed, angry, irritable, frustrated, nervous or tense, restless, difficulty concentrating, hunger, tired*); additional measures of positive affect (*happy, calm/relaxed, enthusiastic*) and negative affect (*stressed, bored*); tobacco craving (*urge or craving to smoke*); smoking expectancies (*a cigarette would improve my mood or make me feel better, my health would improve if I quit smoking, smoking is hurting my health*); smoking evaluations (*I enjoyed my last cigarette*); smoking stigma (*I am embarrassed/ashamed that I am a smoker*); motivation and desire to quit (*I am motivated to quit smoking, I want to quit smoking*); impulsivity (*I can delay gratification, I feel impulsive*); smoking-related self-efficacy and locus of control (*If I tried to quit smoking right now I would be successful, The amount I smoke is within my control*); and social and situational context (*people or situations are triggering me to smoke, I feel comfortable in my current location/situation, I am enjoying my interactions with other people*). The first survey of each day included four additional items assessing sleep (*how many hours did you sleep last night, difficulty falling or staying asleep,*

restless or unsatisfying sleep) and predicted cigarette consumption (*how many cigarettes do you think you'll smoke today*). The mean number of completed surveys across all participants was 106.60 ($SD = 9.99$), with a minimum of 87 and a maximum of 128.

Post Mobile Phone Survey Measures

Upon completion of the 30 days of daily phone survey, participants were sent a link to complete a final web-based survey, assessing their experience completing the mobile phone surveys. This survey assessed the perceived costs and benefits of the mobile assessment procedure, the amount of time it took participants to complete the surveys, and included an open-response section where participants could provide feedback regarding their experience.

Data Preparation and Analysis

Each individual's time series of ESM data was prepared for analysis case by case. First, the time stamps accompanying each outgoing survey were used to calculate continuous time from the cumulative sum of time. Before continuing with temporal variable construction, we ensured that the data frame had a uniform structure that reflected four surveys per day, starting with the morning survey. For most data sets, this simply required inserting rows at the beginning or end of the time series. However, for any participants completing surveys on October 19th, 2017, an internal missing row was required due to a dropped survey from the outgoing server. It was imperative to have a uniform data structure of pings and days, because we next created variables for ping number (i.e. morning, midday, evening, and night), day of the week, and cycles of 12-hr and 24-hr frequency. The latter variables were constructed from sine and cosine terms, consistent with methods provided by Flury and Levri (1999).

The participant's smoking data was then assessed for the frequency and distribution of smoking events. Participants were omitted from analyses for insufficient frequency ($N = 8$) and

nonstationarity ($N = 2$). In the latter case, one participant quit smoking during data collection and one participant dramatically increased their rate of smoking. Global trends such as these make dividing the time series into training and testing sets infeasible (see, Fisher & Bosley, 2019). Smoking data were then transformed from counts to dichotomies. For 69% of the sample, this simply required recoding any number of cigarettes smoked greater than 0 as 1, yielding a dichotomous variable that reflected whether or not the participant had smoked since the last ping. However, the remaining 31% ($N = 16$) of the sample smoked too frequently to yield a dummy-coded variable with enough hits and misses to analyze. For these individuals, we coded any number of cigarettes above the median number smoked as 1 and any number at or below the median as 0, reflecting heavy smoking events.

Finally, all temporal variables (time of day, day of the week, 12- and 24-hr cycles) and the dichotomous smoking variable were set forward in time by one observation. The resulting data structure provided row-wise relationships wherein all temporal variables were aligned with smoking events in time and all remaining predictor variables were lagged predictors. Thus, in the k-fold cross-validations to follow (below), the random sub-selection of rows did not betray the contiguity of the data. All prediction models to follow were lagged models – predicting smoking events roughly four hours in the future. Table 2 presents the complete feature space for prediction models to follow (number of possible variables = 43).

Constructing prediction models

An attempt was made to utilize the first 75% of the time series for training and the remaining 25% for testing. However, for some of the less frequent smokers, this division point was shifted up or down in order to maximize the distribution of smoking events across both training and testing sets. Thus, although the majority of time series were split 75/25, the divisions

ranged from 60/40 to 85/15. Following the division of the time series into training and testing sets, two statistical classification methods were employed for constructing prediction models, elastic net regularization (Zou & Hastie, 2005) and naïve Bayes (Hand & Yu, 2001).

Elastic net regularization. Given its use of the least absolute shrinkage and selection operator (LASSO; Tibshirani, 1996), elastic net is useful for regularization and variable selection. This combination facilitates exploratory analysis of potential predictors, while guarding against overfitting. We used the `cv.glmnet` function in the `glmnet` package in R (Hastie & Qian, 2014) for conducting elastic net regularization. The alpha parameter of the elastic net sets the balance between LASSO and ridge regression (Friedman, Hastie, & Tibshirani, 2001). Values between zero (pure ridge regression) and one (pure LASSO regression) blend the two approaches, with higher values emphasizing LASSO regression and lower values emphasizing ridge regression. All models were initially run with $\alpha = 0.50$, reflecting an equal blend of the two regularization approaches. In the event that no variables were selected, alpha was decreased incrementally by values of 0.05 until a set of nonzero coefficients were returned. A k-fold cross-validation with 10 folds was used to select the optimal model. Consistent with best-practice standards for this approach, the lambda threshold with the minimum mean cross-validated error was used to select the final model.

Naïve Bayes (with random forest feature selection). The naïve Bayes is a probabilistic classifier based on Bayes' theorem (Lewis, 1998). It calculates the maximum a posteriori probability (MAP) for a given hypothesis (H), based on a set of evidence (E), such that $\text{MAP} = \max(P(E|H) * P(H))$. In the present case, H is the hypothesis that a smoking event did or did not occur and E is the feature space containing the temporal variables and lagged predictors. The relative probabilities that an event did or did not occur are calculated, and class (smoked versus

did not smoke) is assigned based on the MAP (see, Zhang, 2004, for further detail). The naïve Bayes classifier assumes conditional independence among predictor values, given class membership. Although this assumption is rarely met, violations typically do not undermine prediction accuracy (Hand & Yu, 2001; Lewis, 1998). In fact, Rish (2001) found that the naïve Bayes classifier performed best when features were either perfectly independent or functionally dependent (i.e. correlated).

The naïve Bayes classifier is a useful and accurate classification approach for prediction modeling. In a large study examining the accuracy of nine different prediction algorithms, naïve Bayes and random forest (Ho, 1995) were the most accurate prediction algorithms across seven separate experiments (Catal & Diri, 2009). Notably, naïve Bayes does not perform variable selection and the inclusion of redundant or irrelevant features can lead to biased posterior probabilities (Zhang, Peña, & Robles, 2009). However, random forest algorithms produce variable importance metrics (Breiman, 2001), and recent studies have suggested employing these metrics for screening and pre-selecting features (c.f. Lunetta, Hayward, Segal, & Van Eerdewegh, 2004). Thus, we employed a random forest feature selection step prior to applying the naïve Bayes. We used the `cforest` function in the `party` package in R (Hothorn, Hornik, Strobl, & Zeileis, 2010). The number of randomly preselected predictor variables (i.e. `mtry`) was set to two and the number of trees run was 50. The `mtry` tuning parameter was set to a relatively low value in order not to discriminate against the selection of correlated predictor variables (Strobl, Boulesteix, Kneib, Augustin, & Zeileis, 2008). Variable importance was then used to select variables for the successive naïve Bayes model. The `varimpAUC` function was used to select all variables with importance values greater than zero. The `naiveBayes` function from the

e1071 package in R (Dimitriadou et al., 2006) was used for conducting naïve Bayes classification.

Results

Complete data, code, analyses, and results for all 52 datasets are available on the Open Science Framework at <https://osf.io/nkemg/>.

Intraindividual results

Elastic net regularization. The results of all 52 elastic net prediction models are presented in Table 3. The mean area under the curve was 0.75 ($SD = 0.12$), and the range was from 0.50 to 1.00. The mean sensitivity was 0.79 ($SD = 0.17$; range = 0.35, 1.00), and the mean specificity was 0.78 ($SD = 0.19$; range = 0.25, 1.00). Figure 1 presents the distributions for AUC, sensitivity, and specificity (top row).

Naïve Bayes Classification (with random forest variable selection). Results for all 52 prediction models are again provided in Table 3. The mean area under the curve was 0.78 ($SD = 0.11$), and the range was from 0.52 to 1.00. The mean sensitivity was 0.81 ($SD = 0.18$; range = 0.30, 1.00), and the mean specificity was 0.79 ($SD = 0.16$; range = 0.45, 1.00). Figure 1 presents the distributions for AUC, sensitivity, and specificity (bottom row).

Comparisons. Two additional comparisons were carried out during the course of our analyses. First, given recent arguments that traditional logistic regression models outperform classification approaches labeled as ‘machine learning’ (Jie, Collins, Steyerberg, Verbakel, & van Calster, 2019; Makridakis, Spiliotis, & Assimakopoulos, 2018a, 2018b), we used the variables retained for elastic net predictions in logistic regression models via the glm function in the stats package of R. That is, elastic net provided variable selection and the traditional logistic regression model provided the regression coefficients used for prediction. Second, although

arguments can be found in the literature for performing feature selection prior to naïve Bayes classification, the use of random forests as carried out in the present study is not widely applied. Thus, we compared this approach to a naïve Bayes classification that simply used the full feature set, rather than pruning via random forest feature selection beforehand. Both comparison approaches underperformed relative to the elastic net and naïve Bayes with feature selection. For the traditional logistic regression models, the mean area under the curve was 0.72 ($SD = 0.12$), and the range was from 0.50 to 0.96. The mean sensitivity was 0.75 ($SD = 0.18$; range = 0.15, 1.00), and the mean specificity was 0.76 ($SD = 0.16$; range = 0.39, 1.00). For naïve Bayes without feature selection, the mean area under the curve was 0.69 ($SD = 0.12$), and the range was from 0.54 to 1.00. The mean sensitivity was 0.76 ($SD = 0.18$; range = 0.33, 1.00), and the mean specificity was 0.72 ($SD = 0.17$; range = 0.36, 1.00).

Exemplars. Of the 52 participants, two—participants 38 and 76—produced perfect prediction models. Of note, in each case, both the elastic net and naïve Bayes approaches yielded perfect predictions, with AUC, sensitivity, and specificity all equal to 1.00. However, in both cases, the two approaches to variable selection selected a different set of variables, indicating that perfect prediction may not require an equally perfect dimension reduction process.

For participant 38, the elastic net selected six variables: *irritable, difficulty concentrating, calm, smoking is hurting my health, Tuesday, and Thursday*. The random forest variable importance metric selected 14 variables: *people or situations are triggering me to smoke, difficulty concentrating, enthusiastic, smoking is hurting my health, motivated to quit, I can delay gratification, my health would improve if I quit, I am enjoying my interactions with other people, midday ping, evening ping, night ping, Saturday, 24-hour cycle, 12-hour cycle*.

For participant 76 the elastic net selected 10 variables: *enjoyed my last cigarette, people or situations are triggering me to smoke, difficulty concentrating, hungry, happy, nervous, tired, a cigarette would improve my mood or make me feel better, I can delay gratification, Thursday*. The random forest variable importance metric selected 25 variables: *craving, enjoyed my last cigarette, people or situations are triggering me to smoke, difficulty concentrating, restless, bored, happy, enthusiastic, nervous, tired, smoking is hurting my health, a cigarette would improve my mood or make me feel better, If I tried to quit smoking right now, I could be successful, I can delay gratification, impulsive, I feel comfortable in my current situation, I am enjoying my interactions with other people, midday, evening, night, Tuesday, Thursday, Friday, 24-hour cycle, 12-hour cycle*.

In both cases, the elastic net selected fewer variables. One interesting point of contrast is that the prediction models for participant 38 were fairly independent. Only two of the variables selected by the elastic net were included in the 14 variables selected by random forest. However, for participant 76, the naïve Bayes prediction model included all but one of the 10 elastic net variables.

Between-subject results

We then examined whether baseline variables collected prior to ESM sampling could predict the sensitivity, specificity, or accuracy (AUC) of individual models. Given that there were a number of candidate predictor variables ($N = 55$), and little *a priori* rationale for subsetting the predictor set, we again employed a random forest variable selection step to select potential predictors. We then used the selected variables as predictors in a ridge regression to assess the degree to which the model explained variation in AUC, sensitivity, and specificity.

Ridge was selected to guard against overfitting via penalization, without further pruning the set of predictors. Complete models are provided in Supplemental Materials.

The models for elastic net AUC, sensitivity, and specificity included 14, 17, and 7 variables, respectively, and explained 39%, 61%, and 17% of the variance. The models for naïve Bayes with feature selection AUC, sensitivity, and specificity included 10, 14, and 13 variables, respectively, and explained 33%, 43%, and 57% of the variance. The most consistent predictors were neurotic features (e.g. anxiety, depression), and relationship status. Greater anxiety and depression predicted more accurate predictions, whereas being in a committed relationship predicted less accurate predictions.

Discussion

The present study sought to apply statistical classification methods to idiographic time series data in order to make accurate future predictions of behavior. We selected cigarette smoking as an exemplar behavior for which the current approach could be applied. We recruited 70 individuals who presented as regular smokers; 52 completed ESM data collection and provided sufficient time series data. We were interested in building prediction models that accurately reflected the precise timing and psychosocial predictors of smoking behavior. Each individual's time series was split into sequential training and testing sections, so that trained models could be tested on future observations. Two prediction algorithms were employed, elastic net regularization and naïve Bayes classification. Sample-wide prediction accuracy was nearly 80%, with some models demonstrating perfect prediction accuracies. Moreover, these were *lagged* analyses. Thus, not only did we support the prediction accuracy of our models with out-of-sample tests in time-forward data, the models themselves were time-lagged, such that each prediction was for the subsequent measurement. Such a system could be the basis for mobile,

just-in-time interventions for substance use. Models that accurately predict future behavior could ostensibly be used for delivering personalized interventions at empirically-indicated moments of need.

Overall, both approaches produced accurate predictions of future smoking behavior. The naïve Bayes procedure—with a preliminary random forest step for feature selection—returned the highest area under the curve (AUC), sensitivity, and specificity. However, whereas the top end of the performance range included prediction models with perfect performance, it should also be acknowledged that both classification approaches yielded models with prediction accuracy at chance. Moreover, 18 participants who enrolled in the study did not produce analyzable data. Whereas, eight of these individuals failed to provide data (either by dropping out or simply failing to respond to survey pings), 10 individuals who provided adequate data were excluded because they either did not smoke often enough or significantly changed their rate of smoking during the survey period. Thus, although the present study provides promising evidence for the ability to make accurate future predictions of smoking behavior, and points towards an application in just-in-time interventions, it is important to emphasize that this approach is not universally applicable and further work is needed to better understand for whom this will or will not work *a priori*.

The present study represents the first application of idiographic machine learning to the prediction of future substance use behavior. It has been argued that fully idiographic approaches—those that model intraindividual data, one participant at a time—are required for a maximally accurate understanding of individual phenomenology and behavior (Fisher et al., 2018; Medaglia, Jeronimus, & Fisher, 2019). It should be acknowledged that the application of ESM paradigms to tobacco use research is not a new idea. A number of studies have used

ecological and ambulatory data to measure cigarette smoking *in vivo*, and to better-understand the real-world predictors of smoking behavior (Serre, Fatseas, Swendsen, & Auriacombe, 2015). However, multilevel modeling has typically been employed for analyzing these data (Rendina, Moody, Ventuneac, Grov, & Parsons, 2015; Serre et al., 2015), and the homogeneity requirements for such procedures are likely to be violated, as random effects can only deviate from common structures in the data. As we note in the introduction, smoking is a behavior that (a) is enacted by individuals, (b) occurs within the context and timing of those individuals' lives and (c) is dictated by idiosyncratic predictor variables that vary person by person.

We presupposed that these idiosyncrasies in predictor variables—the individual differences in psychosocial influences and temporal fluctuations—would be vital to making accurate predictions of smoking behavior. We have argued on several occasions previously that time-varying, intraindividual processes are likely to be nonergodic (Fisher et al., 2018; Medaglia et al., 2019), and that statistical models interested in describing the functioning and behavior of individuals should be carried out on a person by person basis (Fisher, 2015; Fisher et al., 2019; Fisher, Reeves, Lawyer, Medaglia, & Rubel, 2017). Our results appear to have substantiated these assumptions. Individual prediction models were both highly heterogeneous in their composition, and largely accurate, exhibiting strong AUC, sensitivity and specificity. Sixty percent of naïve Bayes prediction models ($N = 31$) and 56% of elastic net prediction models ($N = 29$) exhibited prediction accuracy of 75% or greater. Greater than half of all prediction models exhibited sensitivity and specificity at or above 80%.

Taken together, the current study presents a set of methods for accurately classifying smoking events based on emotional, behavioral, and temporal features. Moreover, these methods could be leveraged to provide predictions for just-in-time interventions on mobile devices,

providing the appropriate treatment content, at the right time, based on an individual's changing internal and contextual state. However, such an implementation would require ongoing collection of data and continued engagement from the individual. Thus, it is imperative to restrict the feature space as much as possible in order to streamline data collection and limit participant burden. Although elastic net produces sparse models as a function of L_1 penalization, naïve Bayes does not perform variable selection and requires a preliminary step for pruning potential features. The present study employed a random forest algorithm for determining variable importance and selecting variables. In addition to reducing the number of variables, we found that this step also increased prediction accuracy.

The present study had two notable limitations. First, we employed a fixed measurement interval for data collection and a fixed time lag for prediction models, regardless of the frequency of the smoking behavior. For some individuals, this required us to predict *heavy* smoking events, rather than presence versus absence. In order to do the latter, a shorter measurement interval and higher frequency sampling paradigm would be called for. However, this would potentially shorten the length of the lag interval. It is an empirical question whether the predictive relationships found in the present study are robust to differential lag times. Further research is required to fine tune the relationships between the measurement frequency, the frequency of the behavior, and the predictive power of features across varying lag times. Second, the sample size was modest and it remains to be seen how well the current method will scale. Future research should endeavor to replicate the current findings across larger samples.

Nevertheless, the current findings are promising and point toward a repeatable and reliable method by which to predict individual behavior from idiographic ESM data. It is our argument that, for problems such as substance use occurrence, it is unlikely that there is a single

unified predictive model for all individuals. Instead, we propose that methods such as these are generalizable as *processes*—pipelines for progressing data from collection to prediction, case by case. Replicating the current results in larger and potentially more diverse samples could support the notion that this process can be applied generally to individuals who smoke.

REFERENCES

- Ashley, E. A. (2015). The precision medicine initiative: a new national effort. *JAMA*, 313(21), 2119-2120.
- Bae, S., Chung, T., Ferreira, D., Dey, A. K., & Suffoletto, B. (2018). Mobile phone sensors and supervised machine learning to identify alcohol use events in young adults: Implications for just-in-time adaptive interventions. *Addictive Behaviors*, 83, 42-47. doi: <https://doi.org/10.1016/j.addbeh.2017.11.039>
- Border, R., Johnson, E. C., Evans, L. M., Smolen, A., Berley, N., Sullivan, P. F., & Keller, M. C. (2019). No Support for Historical Candidate Gene or Candidate Gene-by-Interaction Hypotheses for Major Depression Across Multiple Large Samples. *American Journal of Psychiatry*, appi.ajp.2018.18070881. doi: 10.1176/appi.ajp.2018.18070881
- Breiman, L. (2001). Random Forests. *Machine Learning*, 45(1), 5-32. doi: 10.1023/A:1010933404324
- Catal, C., & Diri, B. (2009). Investigating the effect of dataset size, metrics sets, and feature selection techniques on software fault prediction problem. *Information Sciences*, 179(8), 1040-1058. doi: <https://doi.org/10.1016/j.ins.2008.12.001>
- Christodoulou, E., Ma, J., Collins, G. S., Steyerberg, E. W., Verbakel, J. Y., & Van Calster, B. (2019). A systematic review shows no performance benefit of machine learning over logistic regression for clinical prediction models. *Journal of Clinical Epidemiology*, 110, 12-22. doi: <https://doi.org/10.1016/j.jclinepi.2019.02.004>
- Dimitriadou, E., Hornik, K., Leisch, F., Meyer, D., Weingessel, A., & Leisch, M. F. (2006). The e1071 package. *Misc Functions of Department of Statistics (e1071)*, TU Wien.

- Dzau, V. J., Ginsburg, G. S., Van Nuys, K., Agus, D., & Goldman, D. (2015). Aligning incentives to fulfil the promise of personalised medicine. *The Lancet*, 385(9982), 2118-2119.
- Fisher, A. J. (2015). Toward a dynamic model of psychological assessment: Implications for personalized care. *Journal of Consulting and Clinical Psychology*, 83(4), 825-836. doi: 10.1037/ccp0000026
- Fisher, A. J., & Bosley, H. G. (2019). Identifying the Presence and Timing of Discrete Mood States Prior to Therapy. doi: <https://doi.org/10.31219/osf.io/2jrhf>
- Fisher, A. J., Bosley, H. G., Fernandez, K. C., Reeves, J. W., Soyster, P. D., Diamond, A. E., & Barkin, J. (2019). Open trial of a personalized modular treatment for mood and anxiety. *Behaviour Research and Therapy*, 116, 69-79. doi: <https://doi.org/10.1016/j.brat.2019.01.010>
- Fisher, A. J., Medaglia, J. D., & Jeronimus, B. F. (2018). Lack of group-to-individual generalizability is a threat to human subjects research. [10.1073/pnas.1711978115]. *Proceedings of the National Academy of Sciences*.
- Fisher, A. J., Reeves, J. W., Lawyer, G., Medaglia, J. D., & Rubel, J. A. (2017). Exploring the idiographic dynamics of mood and anxiety via network analysis. *Journal of Abnormal Psychology*, 126(8), 1044-1056. doi: <http://dx.doi.org/10.1037/abn0000311>
- Flury, B. D., & Levri, E. P. (1999). Periodic logistic regression. *Ecology*, 80(7), 2254-2260.
- Friedman, J., Hastie, T., & Tibshirani, R. (2001). *The elements of statistical learning* (Vol. 1): Springer series in statistics New York, NY, USA:.
- Hamburg, M. A., & Collins, F. S. (2010). The path to personalized medicine. *The New England Journal of Medicine*, 363(4), 301-304. doi: 10.1056/NEJMp1006304

- Hand, D. J., & Yu, K. (2001). Idiot's Bayes—not so stupid after all? *International statistical review*, 69(3), 385-398.
- Hastie, T., & Qian, J. (2014). Glmnet vignette. Retrieve from http://www.web.stanford.edu/~hastie/Papers/Glmnet_Vignette.pdf. Accessed September, 20, 2016.
- Ho, T. K. (1995). *Random decision forests*. Paper presented at the Proceedings of 3rd international conference on document analysis and recognition.
- Hoepfner, B. B., Schick, M. R., Kelly, L. M., Hoepfner, S. S., Bergman, B., & Kelly, J. F. (2017). There is an app for that—Or is there? A content analysis of publicly available smartphone apps for managing alcohol use. *Journal of substance abuse treatment*, 82, 67-73.
- Hothorn, T., Hornik, K., Strobl, C., & Zeileis, A. (2010). Party: A laboratory for recursive partytioning.
- Hyman, S. E. (2018). The daunting polygenicity of mental illness: making a new map. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 373(1742), 20170031. doi: 10.1098/rstb.2017.0031
- Insel, T. R. (2017). Digital phenotyping: Technology for a new science of behavior. *JAMA*, 318(13), 1215-1216. doi: 10.1001/jama.2017.11295
- Iyengar, R., Altman, R. B., Troyanskaya, O., & FitzGerald, G. A. (2015). Personalization in practice. *Science*, 350(6258), 282-283.
- Jie, M., Collins, G. S., Steyerberg, E. W., Verbakel, J. Y., & van Calster, B. (2019). A systematic review shows no performance benefit of machine learning over logistic regression for clinical prediction models. *Journal of clinical epidemiology*.

- Lewis, D. D. (1998, 1998//). *Naive (Bayes) at forty: The independence assumption in information retrieval*. Paper presented at the Machine Learning: ECML-98, Berlin, Heidelberg.
- Lunetta, K. L., Hayward, L. B., Segal, J., & Van Eerdewegh, P. (2004). Screening large-scale association study data: exploiting interactions using random forests. *BMC genetics*, 5(1), 32.
- Makridakis, S., Spiliotis, E., & Assimakopoulos, V. (2018a). The M4 Competition: Results, findings, conclusion and way forward. *International Journal of Forecasting*, 34(4), 802-808. doi: <https://doi.org/10.1016/j.ijforecast.2018.06.001>
- Makridakis, S., Spiliotis, E., & Assimakopoulos, V. (2018b). Statistical and Machine Learning forecasting methods: Concerns and ways forward. *PLOS ONE*, 13(3), e0194889. doi: [10.1371/journal.pone.0194889](https://doi.org/10.1371/journal.pone.0194889)
- Medaglia, J. D., Jeronimus, B. F., & Fisher, A. J. (2019). Reply to Adolf and Fried: Conditional equivalence and imperatives for person-level science. *Proceedings of the National Academy of Sciences*, 201820221. doi: [10.1073/pnas.1820221116](https://doi.org/10.1073/pnas.1820221116)
- Nahum-Shani, I., Hekler, E. B., & Spruijt-Metz, D. (2015). Building health behavior models to guide the development of just-in-time adaptive interventions: A pragmatic framework. *Health Psychology*, 34(S), 1209.
- Rendina, H. J., Moody, R. L., Ventuneac, A., Grov, C., & Parsons, J. T. (2015). Aggregate and event-level associations between substance use and sexual behavior among gay and bisexual men: comparing retrospective and prospective data. *Drug and alcohol dependence*, 154, 199-207.

- Rish, I. (2001). *An empirical study of the naive Bayes classifier*. Paper presented at the IJCAI 2001 workshop on empirical methods in artificial intelligence.
- Schork, N. J. (2015). Personalized medicine: time for one-person trials. *Nature*, 520(7549), 609-611.
- Serre, F., Fatseas, M., Swendsen, J., & Auriacombe, M. (2015). Ecological momentary assessment in the investigation of craving and substance use in daily life: a systematic review. *Drug and Alcohol Dependence*, 148, 1-20.
- Siegel, R. L., Miller, K. D., & Jemal, A. (2018). Cancer statistics, 2018. *CA: A Cancer Journal for Clinicians*, 68(1), 7-30. doi: 10.3322/caac.21442
- Soyster, P. D., & Fisher, A. J. (2019). Involving stakeholders in the design of ecological momentary assessment research: An example from smoking cessation. doi: <https://doi.org/10.31234/osf.io/mg9tw>
- Spruijt-Metz, D., & Nilsen, W. (2014). Dynamic Models of Behavior for Just-in-Time Adaptive Interventions. *IEEE Pervasive Computing*, 13(3), 13-17. doi: 10.1109/MPRV.2014.46
- Strobl, C., Boulesteix, A.-L., Kneib, T., Augustin, T., & Zeileis, A. (2008). Conditional variable importance for random forests. *BMC Bioinformatics*, 9(1), 307. doi: 10.1186/1471-2105-9-307
- Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society. Series B (Methodological)*, 267-288.
- Warren, G. W., Alberg, A. J., Kraft, A. S., & Cummings, K. M. (2014). The 2014 Surgeon General's report: "The Health Consequences of Smoking—50 Years of Progress": a paradigm shift in cancer care. *Cancer*, 120(13), 1914-1916.

- Xu, X., Bishop, E. E., Kennedy, S. M., Simpson, S. A., & Pechacek, T. F. (2015). Annual healthcare spending attributable to cigarette smoking: an update. *American journal of preventive medicine*, 48(3), 326-333.
- Zhang, H. (2004). The optimality of naive Bayes. *AA*, 1(2), 3.
- Zhang, M.-L., Peña, J. M., & Robles, V. (2009). Feature selection for multi-label naive Bayes classification. *Information Sciences*, 179(19), 3218-3229. doi: <https://doi.org/10.1016/j.ins.2009.06.010>
- Zou, H., & Hastie, T. (2005). Regularization and variable selection via the elastic net. *Journal of the royal statistical society: series B (statistical methodology)*, 67(2), 301-320.

Table 1: Participant characteristics.

	Retained Participants (N=52)	Excluded Participants (N=18)
Age (<i>M, SD</i>)	25.63 (11.1)	22.06 (4.65)
Gender identity (<i>N, %</i>)		
Female	19 (36.54%)	6 (33.33%)
Male	32 (61.54%)	12 (66.66%)
Other	0 (0.00%)	0 (0.00%)
Did not disclose	1 (1.92%)	0 (0.00%)
Racial identity (<i>N, %</i>)		
White	17 (32.69%)	6 (33.33%)
Black or African American	1 (1.92%)	1 (5.56%)
Asian/Pacific Islander	20 (38.50%)	6 (33.33%)
Mixed or 'other'	12 (23.08%)	5 (27.78%)
Did not disclose	2 (3.85%)	0 (0.00%)
Ethnic identity (<i>N, %</i>)		
Hispanic/Latinx	5 (9.62%)	3 (16.67%)
Non-Hispanic/Latinx	45 (86.54%)	15 (83.33%)
Did not disclose	2 (3.85%)	0 (0.00%)
Sexual orientation (<i>N, %</i>)		
Heterosexual	39 (75.00%)	14 (77.78%)
Homosexual	3 (5.77%)	1 (5.56%)
Bisexual/Queer	6 (11.54%)	2 (11.11%)
Other orientation	3 (5.77%)	1 (5.56%)
Did not disclose	1 (1.92%)	0 (0.00%)
MacArthur SES (<i>M, SD</i>)	6.42 (1.80)	6.6 (1.93)
FTND (<i>M, SD</i>)	4.15 (1.66)	2.11 (1.82)
CPD (<i>M, SD</i>)	4.22 (4.53)	3.44 (3.64)
Years smoked (<i>M, SD</i>)	8.15 (10.24)	3.78 (2.92)
Stage of change (<i>N, %</i>)		
Precontemplation	35 (67.31%)	6 (33.33%)
Contemplation	14 (26.92%)	8 (44.44%)
Preparation	3 (5.77%)	4 (22.22%)
Desire to quit (<i>M, SD</i>)	4.76 (3.03)	4.90 (2.70)

Note: MacArthur SES = MacArthur Subjective Socioeconomic Status (scale range 1 [lowest] – 10 [highest]); FTND = Fagerström Test of Nicotine Dependence (scale range 0 [lowest] – 10 [highest]); CPD = average cigarettes per day.

Table 2: Complete Feature Space for Personalized Prediction Models

Lagged Predictors:	
1. Smoked (yes/no) autocorrelation	25. The amount I smoke is within my control
2. How strong is your urge/craving to smoke right now?	26. If I tried to quit smoking right now; I would be successful
3. I enjoyed my last cigarette	27. I can delay gratification
4. People or situations are triggering me to smoke	28. My health would improve if I quit smoking
5. I feel irritable	29. I feel comfortable in my current location/situation
6. I feel angry	30. I am enjoying my interactions with other people
7. I am having difficulty concentrating	Temporal predictors
8. I am hungry	<i>Morning ping (reference point, not included)</i>
9. I feel restless	31. Midday ping
10. I feel down/depressed	32. Evening ping
11. I feel bored	33. Nighttime ping
12. I feel stressed	<i>Monday (reference point, not included)</i>
13. I feel calm/relaxed	34. Tuesday
14. I feel happy	35. Wednesday
15. I feel enthusiastic	36. Thursday
16. I feel nervous/tense	37. Friday
17. I feel tired	38. Saturday
18. I feel frustrated	39. Sunday
19. I feel impulsive	40. 12-hour cycle (sine term)
20. My smoking is hurting my health	41. 12-hour cycle (cosine term)
21. I am motivated to quit smoking	42. 24-hour cycle (sine term)
22. I want to quit smoking	43. 24-hour cycle (cosine term)
23. A cigarette would improve my mood or make me feel better	
24. I am embarrassed/ashamed that I am a smoker	

Table 3: Complete prediction model results.

ID	Obs.	Median	EL AUC	EL Spec	EL Sen	NB AUC	NB Spec	NB Sen
2	102	1	0.89	1.00	0.73	0.87	0.78	0.87
3	85	0	0.78	0.82	0.75	0.82	0.86	0.88
4	77	0	0.74	1.00	0.67	0.81	1.00	0.67
7	101	0	0.96	0.88	1.00	0.98	0.88	1.00
8	104	0	0.67	0.63	0.87	0.72	0.63	0.91
9	81	0	0.76	0.88	0.73	0.67	0.88	0.47
10	105	1	0.60	0.71	0.67	0.69	0.71	0.83
12	117	1	0.82	0.73	0.89	0.86	0.82	0.89
13	78	0	0.73	1.00	0.67	0.85	1.00	0.67
14	108	0	0.65	1.00	0.36	0.60	0.50	0.82
15	65	0	0.83	1.00	0.63	0.83	1.00	0.63
16	78	0	0.81	0.84	0.83	0.95	0.88	1.00
17	114	0	0.73	0.60	1.00	0.75	0.60	1.00
18	109	1	0.72	0.53	0.90	0.73	0.84	0.60
20	68	1	0.52	0.71	0.56	0.70	0.75	0.67
21	91	0	0.88	0.80	1.00	0.93	0.87	1.00
22	89	0	0.78	0.62	1.00	0.70	0.62	0.89
23	104	1	0.78	0.89	0.58	0.69	0.78	0.75
24	113	0	0.61	0.39	0.91	0.60	0.56	0.91
25	92	0	0.73	0.55	1.00	0.82	0.70	1.00
29	110	1	0.78	0.90	0.67	0.69	0.60	0.83
33	76	0	0.83	0.83	0.77	0.90	0.83	0.92
34	84	0	0.82	0.81	0.83	0.75	0.63	1.00
35	81	0	0.62	0.67	0.65	0.75	0.67	0.85
36	68	0	0.51	0.36	0.80	0.75	1.00	0.60
37	91	0	0.78	0.73	0.88	0.64	0.91	0.38
38	96	0	1.00	1.00	1.00	1.00	1.00	1.00
40	99	1	0.50	0.25	1.00	0.69	1.00	0.50
41	113	0	0.76	0.86	0.71	0.77	0.86	0.76
42	91	1	0.67	1.00	0.50	0.61	1.00	0.33
44	68	0	0.79	1.00	0.50	0.82	0.68	1.00
45	66	1	0.71	1.00	0.50	0.93	1.00	0.75
46	102	1	0.79	0.78	0.82	0.89	0.89	0.82
47	85	0	0.64	0.85	0.54	0.68	0.75	0.62
48	90	0	0.76	0.67	1.00	0.86	0.67	1.00
49	77	1	0.60	0.71	0.67	0.68	0.70	0.70
52	77	0	0.85	1.00	0.69	0.90	1.00	0.69
53	78	0	0.61	0.82	0.56	0.52	0.45	0.78
56	81	1	0.77	0.67	0.80	0.73	0.50	1.00
57	70	0	0.67	0.67	0.88	0.76	0.78	0.88
58	77	0	0.86	0.71	1.00	0.85	0.86	0.83
59	70	0	0.86	0.80	0.80	0.77	0.80	0.80
60	72	0	0.63	0.75	0.60	0.60	0.50	0.90
61	90	0	0.79	0.62	0.89	0.83	0.77	0.89
62	88	0	0.74	0.91	0.73	0.71	0.91	0.73
63	55	1	0.80	0.91	0.75	0.61	0.64	0.75
64	50	0	0.89	0.82	1.00	0.89	0.73	1.00
65	84	0	0.65	0.60	0.95	0.69	1.00	0.50
71	62	0	0.80	0.80	1.00	0.80	0.70	1.00
73	90	1	0.62	0.36	1.00	0.79	0.79	0.75
75	113	1	0.94	1.00	0.75	0.87	0.67	1.00
76	126	0	1.00	1.00	1.00	1.00	1.00	1.00
Average	88.83		0.75	0.78	0.79	0.78	0.79	0.81

Note: ID = participant identification number; Obs. = number of observations after lagging and removing rows with missing data; EL = elastic net; NB = naïve Bayes; AUC = area under the curve; Spec = specificity; Sen = sensitivity.

Figure 1: Area under the curve, sensitivity, and specificity for 52 elastic net prediction models (top row) and 52 naïve Bayes prediction models (with random forest feature selection; bottom row).