Trends in Molecular Medicine

CellPress REVIEWS

Opinion

Biosignature Discovery for Substance Use Disorders Using Statistical Learning

James W. Baurley,[1,2,*] Christopher S. McMahan,[2,3] Carolyn M. Ervin,[1]
Bens Pardamean,[1,2] and Andrew W. Bergen[1,4]

There are limited biomarkers for substance use disorders (SUDs). Traditional statistical approaches are identifying simple biomarkers in large samples, but clinical use cases are still being established. High-throughput clinical, imaging, and 'omic' technologies are generating data from SUD studies and may lead to more sophisticated and clinically useful models. However, analytic strategies suited for high-dimensional data are not regularly used. We review strategies for identifying biomarkers and biosignatures from high-dimensional data types. Focusing on penalized regression and Bayesian approaches, we address how to leverage evidence from existing studies and knowledge bases, using nicotine metabolism as an example. We posit that big data and machine learning approaches will considerably advance SUD biomarker discovery. However, translation to clinical practice, will require integrated scientific efforts.

Biosignatures in Substance Use Disorders

Biomarkers for substance use disorders (SUDs) are available for drug use based on detection of the substance or its metabolites, for example, ethyl glucuronide for alcohol [1], tetrahydrocannabinol for marijuana [2], and cotinine for tobacco [3]. They are not, however, readily available for the neurobiological modifications that result in the maladaptive behaviors we describe as addiction [4]. Clinical phenotyping has been used to assess the presence and severity of SUDs and comorbid psychiatric disease and to evaluate treatment options. We now have massive data on patients with SUDs [e.g., genomics and other **omics** (see Glossary), and imaging on the structure and function of the brain]. How can we use this data in biomarker/biosignature discovery? The ability to combine omics with each other and with complex neurocognitive or imaging data promises to deliver biosignatures that will reflect the behavioral and biological modifications that occur in addiction. Standard biostatistical analysis that has been so useful in clinical research does not perform well in this high-dimensional environment where variables vastly outnumber patients [5]. Studies of SUDs and treatments are beginning to use more comprehensive modeling approaches. In one example, data from diverse domains (brain, personality, cognition, demographics, and genetics) were incorporated into a highly predictive model of current and future alcohol abuse in adolescents [6]. In another recent example, researchers performed an integrative analysis to link genomic variation with expression changes in the brains of alcohol-dependent individuals [7]. While these studies point to an encouraging trend, there still appears to be a gap between the massive data available and the routine use of computational and statistical tools in biomarker/biosignature discovery in SUDs.

Highlights

Substance use disorders are complex diseases with significant heritability and psychosocial susceptibility factors that develop with neurobiological and neurocognitive adaptations after chronic exposure to drugs of abuse. Individual effects of genes and psychosocial factors are modest in most cases. Recent large sample collections have began to identify single factors with confidence.

New molecular, imaging and environmental collection technologies have vastly increased the depth of data per individual. New computation and data management technologies no longer limit complex statistical modeling routines.

Machine learning algorithms are beginning to be applied to SUD data. New or streamlined artificial intelligence algorithms are disrupting all industries. These have direct applications to biomarker/biosignature development.

[1]BioRealm, Culver City, CA, USA
[2]Bina Nusantara University, Jakarta, Indonesia
[3]Clemson University, Clemson, SC, USA
[4]Oregon Research Institute, Eugene, OR, USA

*Correspondence:
baurley@biorealm.ai (J.W. Baurley).

https://doi.org/10.1016/j.molmed.2017.12.008

Biosignature discovery provides a way of combining many variables (e.g., genetic effects, **voxels** in neuroimaging) into meaningful models. Without models of net effects, it is difficult to interpret many small effects, especially given complex correlations in data. **High-dimensional data** is also becoming quite common on large populations, while measures of molecular phenotypes, which may not be cost effective or safely accessed, are less commonly collected. Biobanks of large cohorts with genomic data are becoming available, such as the Million Veteran Program, the All of Us Research Program, UK Biobank, GenomeAsia 100K, and even direct-to-consumer services such as 23andMe and Helix. While there is recognition that additional omics are necessary to understand the influence of genotype on phenotype, the most commonly available data will remain genotypic. Using a biosignature approach, there are opportunities to gain new biological insights and assess multiple predicted phenotypes. This is one of the premises behind transcriptome-wide association studies [8]. Using studies where both genomic and transcriptomic data are available, the tissue-specific relationship between DNA (genotypes) and RNA (gene expression) can be modeled to generate predictive models [9]. These models are then used in genomic data to predict expression and association to diseases or quantitative traits [10,11].

Using high-dimensional data to profile and predict SUD molecular phenotypes and outcomes is a new development path (see Figure 1, Key Figure). Discovery is driven by **statistical learning algorithms** suited for detection of biosignatures from high-dimensional data. The learned biosignatures are then validated and applied in various use cases (e.g., to identify subgroups at risk of SUDs or as tools to optimally select treatments). We advocate propagating the data from new observations and predictions back into the development cycle to allow for continuous improvement of the models (Figure 1). We see great promise in statistical learning to discover and validate biosignatures, but recognize that the translational path into clinical settings will have some unique challenges. In the next sections, we illustrate the workflow presented for this model (Figure 1) by learning biosignatures of nicotine metabolism in high-dimensional genomic data.

Application to Nicotine Metabolism

We use the nicotine metabolism pathway as a motivating example for the utility of statistical learning to discover and use SUD biosignatures from high-dimensional data. Nicotine metabolism is strongly influenced by genes; the majority of variance (74%) is due to additive genetic influence [12,13]. The cytochrome P450 monooxygenase 2A6 (CYP2A6) is the dominant but not exclusive metabolic enzyme in nicotine metabolism [14]. Early work established that the ratio of the first two major metabolites of nicotine [*trans*-3′-hydroxycotinine:cotinine, known as the nicotine metabolite ratio (NMR)] can serve as a biomarker of nicotine metabolism [15]. The NMR is estimated biochemically [16] or via prediction using *CYP2A6* genotypes [17,18]. Genes coding for numerous additional oxidases (FMO3, AO, CYP2B6, POR, and AKR1D1) and the uridine diphosphate glycosyltransferases (UGTs) have been found to be associated with nicotine metabolism through individual genetic variants, or, less commonly, gene/protein expression, or protein activity analyses [19–27]. Moreover, in diverse populations and using blood, saliva, and urine biospecimens from smokers, or using labeled nicotine and cotinine in clinical laboratory studies using blood and urine, nicotine metabolism has been reported to vary by ancestry [28–31], age, gender, body mass index, estrogenic hormones, and alcohol and cigarette consumption [32].

In addition to pharmacologic investigations of nicotine metabolism [14], investigators have studied the influence of nicotine metabolism on smoking cessation retrospectively, using either the biochemical measure of the NMR [33–35] or genotypes associated with reduced NMR

Glossary

Bayes factors: a statistical measure that quantifies the evidence for a hypothesis relative to an alternative hypothesis.

Bayesian: a method of statistical inference in which Bayes' theorem is used to update the probability for a hypothesis as more evidence or information becomes available.

Bupropion: an FDA approved non-nicotine smoking cessation pharmacotherapy, most commonly prescribed in an extended-release formulation.

Edge: an element that connects two nodes in a graph, sometimes defining relationships or assigning weights.

Genome-wide association scan: an epidemiological study designed to evaluate the statistical association among genotypes and traits or diseases of interest.

High-dimensional data: data characterized by a large number of dimensions. The refinement of 'large' within this context is still a topic that is debatable. Many define 'large' relative to the size of the available data; that is, the data is high dimensional when the number of dimensions *(P)* exceeds the number of observations *(N)*.

Learned weights: the estimated model parameters, such as regression coefficients.

Measures of fit: statistical techniques designed to assess how well a model fits a data set.

Model complexity: typically refers to the size of the model; that is, the number of free parameters.

Nicotine replacement therapy: smoking cessation therapy providing the patient with nicotine; there are five FDA-approved modes of administration of NRTs: gum, patch, lozenge, spray, and inhaler.

Node: a point at which lines or pathways intersect or branch.

Omics: the exhaustive and systematic study of a molecular analyte (DNA, RNA, protein, or metabolites, and modifications to the same) from one or multiple ('meta') species, sometimes in relation to a disease or trait, for example, substance-use behavior.

Operators: a mapping that takes as inputs elements of a space and

Key Figure

Biosignature Development Workflow

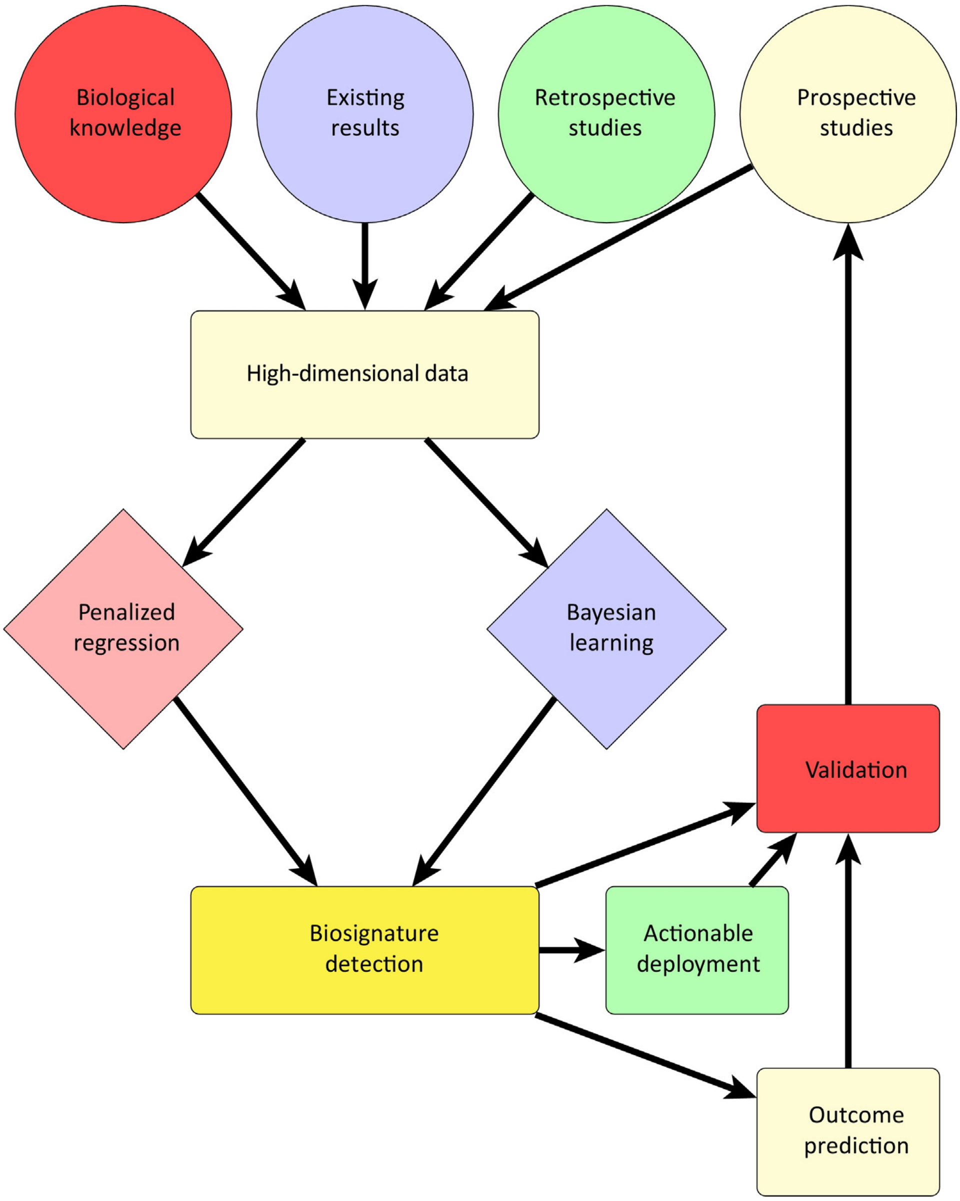

Trends in Molecular Medicine

Figure 1. Data where the number of variables vastly outnumbers the number of samples (high dimensional data) are becoming commonplace in studies of substance abuse disorders and treatment approaches. We present two approaches (penalized regression and Bayesian learning) for detecting the combination of variables (biosignatures) predictive of SUD phenotypes (e.g., nicotine metabolism). Biosignature detection is followed by validation, then prospective assessment of utility for translation to clinical practice.

returns other elements of the same space.

Penalized regression: a regression method that makes use of a penalty structure to regularize regression coefficients, also known as regularized and shrinkage regression.

Posterior predictive distribution: the distribution of possible unobserved values conditional on the observed values.

Posterior probability: the conditional probability of an event after observed/known evidence is taken into account.

Prediction bias: the bias associated with a prediction, with bias referring to the tendency of a measurement process to over- or under-estimate the value of a population parameter.

Prediction variance: the variability associated with a prediction, with variance being a measure of dispersion/deviation from a mean.

Priors: probability models which are meant to reflect a modeler's *a priori* knowledge of the parameters in the data model before some evidence is taken into account.

Single nucleotide polymorphism: a variation in a single nucleotide that occurs at a specific position in the genome.

Statistical learning algorithms: algorithms which implement different statistical learning techniques. That is, algorithms which complete function estimation from a given collection of data.

Varenicline: a smoking cessation aid that is a partial agonist of $\alpha4\beta2$ nicotinic acetylcholine receptors.

Versioning: the variables selected in a model and its estimated weights may change as more data becomes available. Careful tracking of each iteration of the data, algorithms, models, predictions, and validations allows biomarkers and biosignatures to continually improve.

Voxels: much like a pixel in a two-dimensional image, a voxel is a tiny cube that contains information and is used to build a three-dimensional image.

[33,36,37]. There has been one prospective analysis of the influence of the NMR on smoking cessation, examining the efficacy of **nicotine replacement therapy** (NRT), **varenicline** and placebo in slow and normal nicotine metabolizers, with the NMR determined from direct biochemical measurement [38]. Note that biomarkers of nicotine metabolism, as studied in the literature, have differed somewhat depending on the genotyping approach, biochemical ratios, and cutoffs selected, as well as on the clinical or population samples used to establish the biomarker; one common dichotomization stratifies individuals with normal metabolism versus slow metabolism.

In general, retrospective studies of smokers, randomized to NRT or placebo, have shown that individuals with biomarkers of slow metabolism, whether defined by genotype or biochemical ratio, were significantly more likely to remain abstinent than individuals with normal metabolism. In one retrospective analysis [37], individuals with slow nicotine metabolism did not benefit (no reduction in relapse proportions) from active treatment (NRT, **bupropion**, or combined active treatment) compared to placebo treatment, while individuals with fast nicotine metabolism did benefit from active treatment. In the prospective trial stratified by the NMR, individuals with normal nicotine metabolism responded significantly better to active treatment than placebo, and those randomized to varenicline responded significantly better than those randomized to NRT [38]. Together, these findings suggest that treatment success can be optimized by assigning treatment to patients by their NMR status, for example, assigning more active pharmacotherapy to normal metabolizers and less active pharmacotherapy to slow metabolizers. A clinical trial using the NMR to assess nicotine metabolism and provide metabolism-matched pharmacotherapy is in progress [39]. Biosignatures for predicting nicotine metabolism in clinical trials of smoking cessation therapies, and in cohorts being studied for tobacco-related behaviors, diseases, and exposures, will be useful to characterize the role of nicotine metabolism in these complex outcomes [40].

The prior knowledge, results, and data described above can be directly used in biosignature development (inputs in Figure 1). Statistical learning algorithms (described in the next sections) can be applied to detected biosignatures of nicotine metabolism. These biosignatures (once validated) can be used to predict other outcomes (such as smoking cessation) or to personalize treatments (e.g., bupropion, varenicline, or NRT for a smoker; Figure 1).

Biosignature Detection in High-Dimensional Data

The data layout for biosignature learning is shown in Figure 2. One or more SUD studies have both high-dimensional data (e.g., genomic) and molecular phenotypic data (e.g., metabolites). For simplicity, it was assumed that the clinical factors were binary and that the genotypes were **single nucleotide polymorphisms** (SNPs), coded by the number of copies of the minor allele [41]. Millions of genotyped and imputed SNPs can be available. The relevant variables (the biosignature) and its net effects on a predictable molecular phenotype (denoted Z in Figure 2) are then learned using statistical algorithms. Once the relationships are represented in models, they can be applied to new data for prediction and assessment of the outcomes of interest (denoted Y in Figure 2). For nicotine metabolism, these could be smoking cessation, lung cancer risk, treatment response, and so on. Of note, the molecular phenotype does not need to be measured once the model is learned, as it can be predicted from the biosignature, sometimes from summary statistics [10,11].

Simple models consider only one genetic variant at a time. A **genome-wide association scan** (GWAS) represents millions of tests for genetic associations with the trait or outcome [5]. There have been four GWASs using the NMR as the trait to date [13,42–44]. These four GWASs used

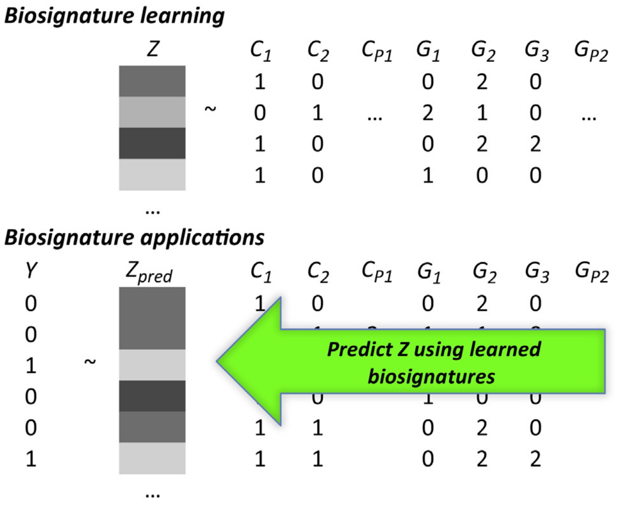

Trends in Molecular Medicine

Figure 2. Biosignatures of Nicotine Metabolism. Nicotine metabolism biosignatures are learned from genotypes *G* and clinical *C* data in laboratory studies of nicotine metabolism. Nicotine metabolism is then predicted (Z_{pred}) in existing or new observations using the biosignatures and corresponding model weights. The predicted nicotine metabolite ratio can them be associated with clinical outcomes *Y*, such as smoking cessation (1 indicates success). Adapted from Gusev *et al.* [10].

readily available GWAS genotyping arrays, as well as typical statistical pipelines for genotype cleaning and imputation via the 1000 Genomes Project resource[i]. Two GWASs represented meta-analyses of single ancestries [13,44] – one was a multiancestry meta-analysis [42] and the other was a multiancestry mega-analysis [43]. As expected, in all four scans, variants in and near *CYP2A6* on human chromosome 19 (the gene encoding the primary nicotine metabolic enzyme) were associated with the NMR at genome-wide significance. In each GWAS, the most significant associations were located proximal, within, and distal to *CYP2A6*, with individual SNPs, significance ranks, and span of association being dependent upon study, sample size, and ancestry composition. We and others [13,42,43] have noted complex patterns of association with the NMR that span into nearby genes, including *CYP2B6* (Figure 3) [42]. Given that there are complex patterns of marginal associations and that the number of variables exceed the sample size (known in statistics as $P > N$), how does one define a biosignature of nicotine metabolism based on genomic data?

The dimensionality problem can be addressed in part by using existing knowledge and results to reduce the model space (Figure 1). From available genome-wide data [42], we selected 11 genomic regions implicated in nicotine metabolism for modeling. We identified the relevant regions using a combination of knowledge bases (such as PharmGKB [45]) and recent literature [19–27]. We then used algorithms to reduce model complexity. The rationale here was that there can be too many variables for a human to decide which should go in a model.

Two classes of statistical learning algorithms can be used for selecting which variables should be in a model (Box 1) and estimating their joint effects on an outcome [46]. These algorithms explore a trade-off between **model complexity**, **prediction bias**, and **prediction variance**. That is, prediction bias can be reduced by increasing model complexity. Alternatively, one can also trade prediction bias for reduced variance using approaches to reduce complexity. The

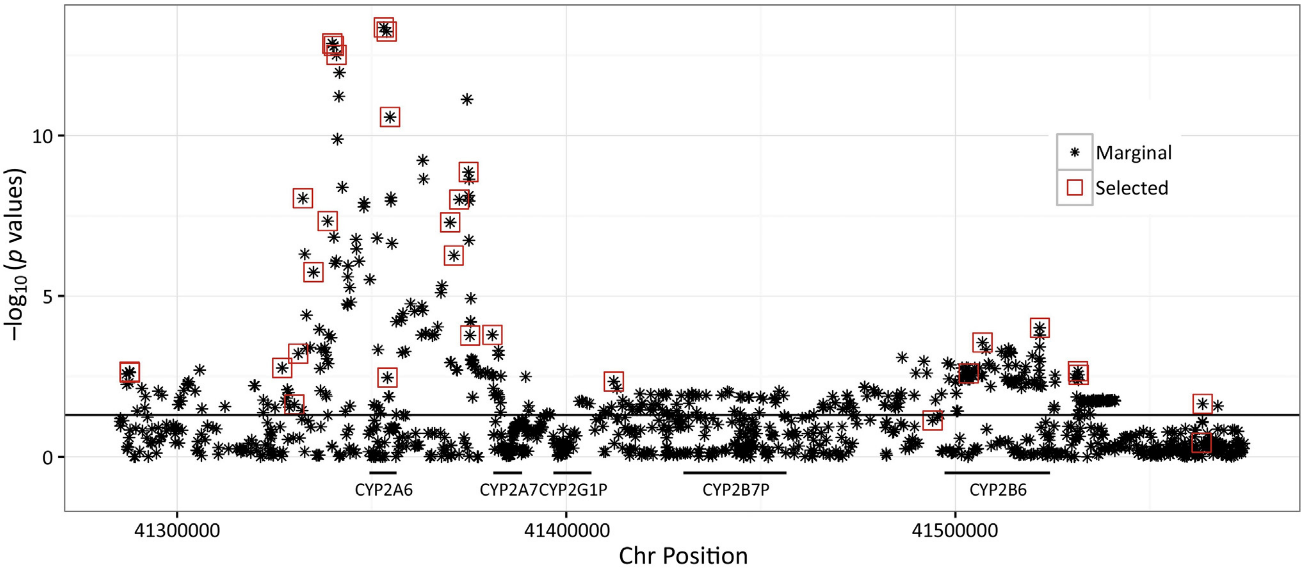

Trends in Molecular Medicine

Figure 3. Genetic Associations with Nicotine Metabolism in the *CYP2A6* Region of Human Chromosome 19. The variants selected using penalized

goal is to find a sweet spot, selecting the right model complexity to minimize prediction error [47,48].

The First Approach: Penalized Regression

Penalized regression approaches add a penalty term to the typical optimization problem [47,48]. That is, many penalized regression methods estimate the regression coefficients by minimizing a penalized residual sum of squares which is given by:

$$\hat{\boldsymbol{\beta}} = \text{argmin}_{\boldsymbol{\beta}} \sum_{i=1}^{N} \left(y_i - \beta_0 - \sum_{j=1}^{P_1} \beta_{1j} C_{ij} - \sum_{j=1}^{P_2} \beta_{2j} G_{ij} - \sum_{j=1}^{P_3} \beta_{3j} Z_{ij} \right)^2 + P(\boldsymbol{\beta}, \boldsymbol{\lambda}),$$

where $\boldsymbol{\lambda}$ represents a collection of tuning parameters which implicitly controls the model complexity, $\boldsymbol{\beta}$ denotes the collection of all the regression coefficients, and $P(\cdot, \cdot)$ is a penalty function [47,48]. For example, a penalty function of the form

$$P(\boldsymbol{\beta}, \boldsymbol{\lambda}) = \lambda_1 \alpha \sum_{k=1}^{3} \sum_{j=1}^{P_k} |\beta_{kj}| + \lambda_2 (1 - \alpha) \sum_{k=1}^{3} \sum_{j=1}^{P_k} \beta_{kj}^2,$$

could be considered, where setting $\alpha = 1$ results in the least absolute shrinkage and selection operator (LASSO) [49], $\alpha = 0$ provides the usual ridge estimator [50], and $\alpha \in (0, 1)$ results in the elastic net [51]. In general, through regularization, the ridge estimator is able to provide better prediction performance by exploiting the so-called bias versus variance trade-off and can be used (unlike standard ordinary least squares) to uniquely estimate the regression coefficients when $P > N$ [50]. Unlike ridge, LASSO provides a parsimonious model through automatic variable selection, though it has been empirically shown that, in comparison to LASSO, ridge maintains a higher level of prediction accuracy in the face of correlated predictors [49]. The elastic net is a blend of ridge and LASSO, which attempts to gain from their strengths and overcome their individual weaknesses. In particular, the elastic net makes use of a linear combination of the ridge and LASSO penalties and can therefore complete automatic variable

Box 1. Model Specification

Generalized linear models can be used as a foundation for biosignature discovery and predicting nicotine metabolism from genomic and clinical data [88]. The conditional mean of Y_i, the observed molecular phenotype or outcome of individual i, is related to P explanatory variables through the link function $g(\cdot)$:

$$g(\mu_i) = \beta_0 + \sum_{j=1}^{P_1} \beta_{1j} C_{ij} + \sum_{j=1}^{P_2} \beta_{2j} G_{ij} + \sum_{j=1}^{P_3} \beta_{3j} Z_{ij},$$

where there are P_1 clinical factors (e.g., age, sex, BMI), P_2 genetic markers, and P_3 derived variables, denoted by C_{ij}, G_{ij}, and Z_{ij}, respectively, with $P = P_1 + P_2 + P_3$. β_{kj} represents the usual regression coefficients. The derived variables can be interaction terms of other functions that combine sets of other variables. In the analysis described herein, the natural logarithm of NMR is used as the response variable so that $g(\cdot)$ can be taken to be the identity link.

selection while maintaining a high degree of prediction accuracy in the face of correlated predictors [51].

There are numerous algorithmic varieties of this general approach, with different properties in terms of handling correlation, sparsity, and picking out features in the data [48]. Moreover, extensive work has shown that none of the penalized regression procedures are universally better in all situations [48,52]. So it is natural to posit the question, ‘What penalized regression method should I use?’ Our retort, ‘Why choose one?’

Different penalized regression procedures can be applied to the motivating nicotine metabolism data; for details on these algorithms, their implementation, and their penalty structures see Table 1. In particular, Table 1 provides nine different penalized regression methods, citations that present the relevant background on each procedure, and R-packages that can be utilized to implement each of the methods. As discussed above, these algorithms have different properties and therefore provide diverse insights into the data. We believe that the collection of models (the ensemble) characterizes the biosignature of this molecular phenotype. This is demonstrated in Figure 4, where the rows represent the biosignatures learned by the different algorithms applied to the nicotine metabolism data [53]. The unshaded portions represent the genetic signature identified by each approach. As expected, a handful of SNPs (out of 3752) were selected in all the models [53]. The largest model included 63 SNPs, but more parsimonious models explained NMR just as well with fewer SNPs (58–62% NMR) [53]. The genetic biosignature found by these methods in human chromosome 19 are overlaid on the marginal

Table 1. Penalized Regression Algorithms Applied to the Nicotine Metabolism Dataset[a]

Method	R-package	Refs
LASSO	glmnet	[49]
Elastic net	glmnet	[51]
Adaptive LASSO	parcor	[89]
SCAD	ncvreg	[90]
MCP	ncvreg	[91]
SCAD-Ridge	ncvreg	[92]
MCP-Ridge	ncvreg	[92]
Dantzig selector	flare	[93]
lq LASSO	flare	[94]

[a] The table includes the R packages implementing it, and the primary research articles describing the algorithms.

Trends in Molecular Medicine

Figure 4. An Ensemble of Models to Define Biosignatures. The rows highlight (unshaded) the sets of SNPs selected by different penalized regression algorithms

results in Figure 3. This highlights additional signals near the *CYP2A6* and *CYP2B6* genes that may have been missed in a more traditional approach. The collection of variables that is predictive of an individual's nicotine metabolism is more than those discoverable using standard genetic association scans. To predict nicotine metabolism with new genotypes, one simply applies the **learned weights** to the SNP biosignatures [53]. The predicted nicotine metabolism can be then applied in additional association or clinical studies.

The differences in the SNPs selected by the penalized regression algorithms (Figure 4) suggest that there are multiple ‘good’ models and one may want to average over the strengths of a set of models when making predictions [54,55]. This leads to the second approach where model uncertainty can be quantified [54].

The Second Approach: Bayesian Model Averaging

In the previous section we discussed the uncertainty with which SNPs belong in a biosignature of nicotine metabolism. **Bayesian** approaches can account for this uncertainty in the model specification. Here, the **posterior probability** for a given model is:

$$p(M|D) = \frac{p(D|M)p(M)}{\sum_{m\in M} p(D|m)p(m)}$$

Obtaining the denominator would involve exploring all possible biosignatures which may not be computationally feasible. Thus the posterior probability is usually approximated by assessing the relative merit of a subset of models [56].

The likelihood above is actually marginalized over the parameters in the model and again is often approximated.

$$p(D|M) = \int_{\beta} p(D|\beta, M)p(\beta)d\beta$$

The **priors** on the model *p(M)* and its parameters *p(β)* can give us the opportunity to formally introduce existing results, knowledge bases, and assumptions into the modeling [57]; for example, what variables are biologically important, the direction and magnitude of their effects, and the certainty with which they are involved. In fact, most penalized regression approaches can also be represented by specifying priors on *p(β)* [48].

More complex relationships among variables can also be learned from the data. Combinations of SNPs or other factors can be condensed into new derived variables Z_j. For example, tree structures, denoted Λ, can be considered where the output of each **node** is determined by its input values and a set of **edge** parameters [58]. One such tree structure is shown in Figure 5 and represents the following system of equations:

$$Z_1 = \theta_{1,1}G_1 + \theta_{1,2}G_2 + (1 - \theta_{1,1} - \theta_{1,2})G_1G_2$$

$$Z_2 = \theta_{2,1}G_3 + \theta_{2,2}G_4 + (1 - \theta_{2,1} - \theta_{2,2})G_3G_4$$

$$Z_3 = \theta_{3,1}Z_1 + \theta_{3,2}Z_2 + (1 - \theta_{3,1} - \theta_{3,2})Z_1Z_2$$

These tree-based derived variables provide a very flexible way of representing interactions [58]. For example, given binary inputs, different edge parameters can represent different **operators**. If $\theta_1 = \theta_2 = 0.5$, the effects are additive; if $\theta_1 = \theta_2 = 0$ there is an effect only when both variants are present (logical AND); and if $\theta_1 = \theta_2 = 1$ there is an effect if either variant is present (logical OR). The effects of the derived variables represent the net effect of the tree [58].

Under this theme, every pairwise SNP effect on nicotine metabolism can be considered (over 6 million derived variables). The evidence for or against each can be computed using **Bayes factors**, the ratio of posterior to prior odds. The top associations are shown in Figure 6. Seven genomic regions showed evidence of interactions. The *CYP2A6* region is an important hub, with other genetic variants near *UGT2B10*, *UGT1A4/A9*, *POR*, *NR1I3*, *NFE2L2*, and *HNF4A*, contributing to its effects on nicotine metabolism (Figure 6).

More complex combinations (starting with one tree and then modifying it to find better combinations) can be sought. This process can be repeated hundreds of thousands of times. Combinations of variants that had large impacts on nicotine metabolism were found. For example, the best learned model had a posterior probability of 0.90%, explained 43% of nicotine metabolism, and had a rather large effect on natural log NMR ($\beta_\Lambda = -1.35$). As in penalized regression, multiple models fit the data well, but the posterior probabilities provide an intuitive way to average models [54,55]. Trees with the highest posterior probabilities contained a handful of SNPs and had strikingly similar **measures of fit** [53]. This suggests that one should make nicotine metabolism predictions by averaging over the collection of models. The **posterior predictive distribution** allowed us to generate predictions of new observations using the entire distribution of explored biosignatures [59].

Incorporating These Approaches into Research and Clinical Translation

Statistical learning algorithms can help identify biosignatures of SUD outcomes. Once learned, these models can provide biological insights on their own as well as be applied to existing or

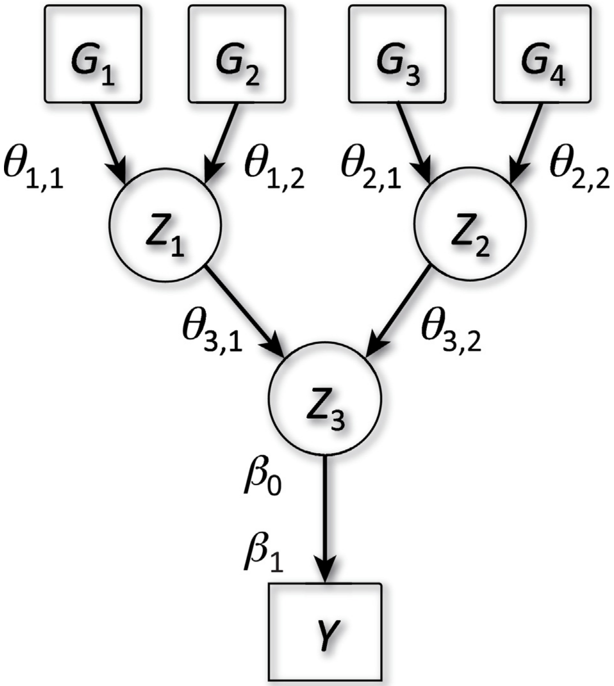

Trends in Molecular Medicine

Figure 5. Tree-Based Structures Can Represent Complex Relationships in Sets of Variables. Here each derived variable Z is computed from its inputs (genetic variants, clinical factors, or other derived variables) and a pair of edge parameters θ. The regression coefficient β_1 represents the net effect of the entire combination of variables on the outcome of interest Y. These structures were explored using Bayesian algorithms to learn biosignatures of nicotine metabolism.

new data to generate predictions. While the focus of the nicotine metabolism example relied on genomics, the approach can be applied to other data commonly available or becoming available in studies of SUDs, such as metabolomic [60,61], personality assessment [62,63], neuroimaging [64,65], and mobile health applications [66]. The first approach described involves learning using penalized regression algorithms. These algorithms select variables while simultaneously estimating their effects. As demonstrated in the nicotine metabolism application, they extract different features from the data. Thus we advocate using the entire ensemble to characterize the biosignatures. The second approach involves learning a distribution of models and leveraging that distribution in prediction [54]. These Bayesian algorithms allow us to consider more complex relationships among variables [58]. In the nicotine metabolism example, variants near *CYP2A6* and other parts of the pathway jointly influenced nicotine metabolism (Figure 6). This suggests that learning algorithms can identify combinations of genetic variances that explain molecular phenotypes that may be missed using traditional analyses. The next steps involve validating the biosignatures in other datasets with both genomic data and nicotine metabolites, and then applying the learned biosignatures and weights in other datasets, to observe how the biosignatures influence smoking-related outcomes.

Many statistical learning approaches have been around for some time, but are just now being applied to SUDs. Computing used to be a major bottleneck in applying these algorithms. But with new acceleration computing and software stacks, algorithms are being retooled to handle

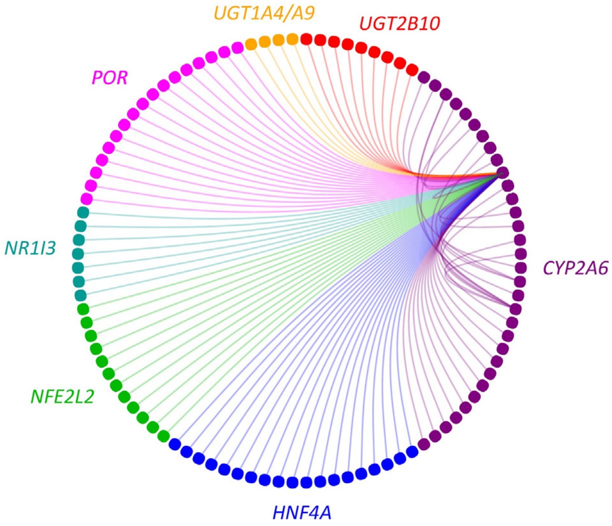

Figure 6. Joint SNP Effects on Nicotine Metabolism. The effects of combination of genetic variant on nicotine metabolism can be explored using Bayesian algorithms [58]. This plot shows that many genetic variants (dots) in different genes (color) can modify the effects of *CYP2A6* variants on nicotine metabolism. This presents another way of defining biosignatures from a collection of models for use in prediction or generating new hypotheses.

Trends in Molecular Medicine

much larger and more complex datasets. As in all industries, data is now in abundance. There are more data on individuals with SUDs available now than ever before [67]. Several groups have begun to explore applying these algorithms to learn predictors of response to treatment for SUDs [68–70], but this is just the beginning of a new wave of discovery. There has been a recent increase of deep learning algorithms being applied to health applications: automatic detection of new tumors from imaging data [71], discovery of new drug targets [72], and precision treatments in cancer patients[ii]. The availability of data, computation, and algorithms have profound implications for the future of biosignature and biomarker development in SUD screening, diagnosis, and treatment.

Retrospective biomarker/biosignature discovery has strengths (available data) and weaknesses (older, possibly less-relevant trials, biospecimen availability) or biases (lack of biospecimens, older molecular datasets). Similarly, hypothesis-driven prospective biomarker discovery and validation has strengths (ability to define variables/study domains, state-of-the-art biospecimen and biomarker data collection) and weaknesses (candidate hypotheses may miss predictive biomarker variables/domains). The choice of a retrospective versus a prospective approach may depend upon currently available resources or the ability to leverage existing public datasets. A prospective discovery and validation design offers the theoretical ability to include all domains or selected (hypothesis-driven) domains; the former is limited by practical considerations and the latter may unfortunately restrict variable discovery. The necessity of validation for replication, and clinical utility analysis to fulfill regulatory and reimbursement requirements, means that designing discovery and validation studies will involve both retrospective and prospective designs. Excluding logical incompatibilities, practical and contingent limitations are more likely to limit or slow biosignature development and translation to practice than the choice of study design or whether either type of design is hypothesis driven.

Guidelines for biomarker development and translation to treatment of omics-based tests has been a hot topic for almost a decade [73,74]. Omics-based tests are defined as an assay composed of or derived from many molecular measurements and interpreted by a fully

Box 2. Clinician's Corner

Biomarkers of SUDs are now used clinically to detect substance use and relapse from abstinence. Emerging biomarkers will enable stratification of diagnosed SUD patients for greatest therapy efficacy.

Future SUD biosignatures will aggregate multiple predictors through omic analysis and statistical learning and dramatically expand diagnostic and predictive utility.

Translation of SUD biosignatures into clinical care requires support for clinical biospecimen testing and clinical counseling.

Improving the effectiveness of SUD biosignatures will be facilitated by integration of SUD biosignature assessment and counseling into clinical care – as smoking status assessment has become routine in multiple medical specialties.

Outstanding Questions

What tools can be developed to help researchers accelerate SUD biomarker/biosignature innovation? Learning algorithms need high-quality data to grow models. Current challenges include culture (data sharing), data merging, and patient consent and regulatory issues.

How will predictive biosignatures or learning algorithms be clinically validated? Traditional biomarker development emphasizes a discovery to clinical validation path, followed by evaluation of a fixed model for clinical utility and use. How will this be adapted for versioned biosignatures in a continuous improvement system? What best practices should there be to encourage innovation, yet maintain patient safety?

What are the challenges of deploying these models into healthcare systems? Predictive biosignatures should be applicable across demographics and environments by qualification across diverse populations. Enabling comparative effectiveness evaluation will require data sharing across regions and practice environments.

specified computational model to produce a clinically actionable result [74]. There are currently no regulatory guidelines on **versioning** of biosignatures whose specifications may adapt and improve with more data (Figure 1). Current guidelines require that biomarker tests for any application be fixed before moving into a clinical trial for assessment of their utility [74,75]. Additionally, evidence is evaluated based on study design, the complexity of the bioassay, and the nature of the mathematical model [76,77]. Many challenges remain in the translation of biosignatures to clinical care; the guidelines, roadmap, and regulatory ecosystem will need to be recalibrated as models and predictions become more dynamic.

Some progress has been made in developing dynamic systems that collect and analyze data from its own processes in order to improve outcomes. Recommendations for development of a rapid learning system for biomarkers encompassing policy, data infrastructure, and patient care are part of an evolutionary process of the biomedical enterprise [74]. This represents an extension of older ideas of a learning [78], continuously improving [79], and genomics-enabled learning [80] healthcare system. Adding biosignature discovery and translation to these ideas implies much greater efforts to align patient care and provider and healthcare system practice than introduction of a single biomarker with a single context of use. The challenges of biosignature translation have perhaps been most thoroughly addressed in oncology [81,82]. However, guidance is emerging on the multiple domain challenges [83,84], the general pathway for biomarker qualification [85], and biomarkers for specific SUDs [86,87].

Concluding Remarks

Here, new strategies for biosignature development have been described that acknowledge the complexities of disease and data and touch on the ongoing challenges of translation to clinical care. In both biosignature development and in translation to clinical care, complex challenges require comprehensive, integrated solutions. We encourage addiction researchers to share data, organize themselves to enable secondary data analyses, and consider applying these and other learning algorithms to their data to generate new biological insights and prediction models. We realize there are some remaining big questions on how the strategy presented fits into existing biomarker development, clinical translation, and regulation paradigms (see Outstanding Questions and Box 2). There is a delicate balance between encouraging standardization and enabling a learning healthcare system that requires scientific and regulatory leadership to advance biosignatures into clinical care.

Acknowledgments

We thank the participants of the PKPD of Nicotine, PKTWIN, and SMOFAM studies and Drs Benowitz, Hops, and Swan. We acknowledge NIH/NIDA grant R43DA041211 for funding support and Amazon Web Services and NVIDIA for computing support. We also acknowledge the work of the entire BioRealm team (https://biorealm.ai/team).

Resources

[i] www.1000genomes.org

[ii] http://candle.cels.anl.gov

References

1. Boscolo-Berto, R. *et al.* (2013) Ethyl glucuronide concentration in hair for detecting heavy drinking and/or abstinence: a meta-analysis. *Int. J. Legal Med.* 127, 611–619
2. Cone, E.J. *et al.* (2015) Nonsmoker exposure to secondhand cannabis smoke. III. Oral fluid and blood drug concentrations and corresponding subjective effects. *J. Anal. Toxicol.* 39, 497–509
3. Benowitz, N.L. *et al.* (2017) Urine cotinine screening detects nearly ubiquitous tobacco smoke exposure in urban adolescents. *Nicotine Tob. Res.* 19, 1048–1054
4. Volkow, N.D. *et al.* (2015) Biomarkers in substance use disorders. *ACS Chem. Neurosci.* 6, 522–525
5. Yi, H. *et al.* (2015) Penalized multimarker vs. single-marker regression methods for genome-wide association studies of quantitative traits. *Genetics* 199, 205–222
6. Whelan, R. *et al.* (2014) Neuropsychosocial profiles of current and future adolescent alcohol misusers. *Nature* 512, 185–189
7. Mamdani, M. *et al.* (2015) Integrating mRNA and miRNA weighted gene co-expression networks with eQTLs in the nucleus accumbens of subjects with alcohol dependence. *PLoS One* 10, e0137671
8. Gamazon, E.R. *et al.* (2015) A gene-based association method for mapping traits using reference transcriptome data. *Nat. Genet.* 47, 1091–1098
9. Lonsdale, J. *et al.* (2013) The Genotype-Tissue Expression (GTEx) project. *Nat. Genet.* 45, 580–585
10. Gusev, A. *et al.* (2016) Integrative approaches for large-scale transcriptome-wide association studies. *Nat. Genet.* 48, 245–252
11. Mancuso, N. *et al.* (2017) Integrating gene expression with summary association statistics to identify genes associated with 30 complex traits. *Am. J. Hum. Genet.* 100, 473–487
12. Swan, G.E. *et al.* (2009) Genetic and environmental influences on the ratio of 3′hydroxycotinine to cotinine in plasma and urine. *Pharmacogenet. Genomics* 19, 388–398
13. Loukola, A. *et al.* (2015) A genome-wide association study of a biomarker of nicotine metabolism. *PLoS Genet.* 11, e1005498
14. Benowitz, N.L. *et al.* (2009) Nicotine chemistry, metabolism, kinetics and biomarkers. *Handb. Exp. Pharmacol.* 192, 29–60
15. Dempsey, D. *et al.* (2004) Nicotine metabolite ratio as an index of cytochrome P450 2A6 metabolic activity. *Clin. Pharmacol. Ther.* 76, 64–72
16. St Helen, G. *et al.* (2012) Reproducibility of the nicotine metabolite ratio in cigarette smokers. *Cancer Epidemiol. Biomarkers Prev.* 21, 1105–1114
17. Bloom, J. *et al.* (2011) The contribution of common CYP2A6 alleles to variation in nicotine metabolism among European-Americans. *Pharmacogenet. Genomics* 21, 403–416
18. Benowitz, N.L. *et al.* (2006) CYP2A6 genotype and the metabolism and disposition kinetics of nicotine. *Clin. Pharmacol. Ther.* 80, 457–467
19. Haberl, M. *et al.* (2005) Three haplotypes associated with CYP2A6 phenotypes in Caucasians. *Pharmacogenet. Genomics* 15, 609–624
20. Wortham, M. *et al.* (2007) Expression of constitutive androstane receptor, hepatic nuclear factor 4 alpha, and P450 oxidoreductase genes determines interindividual variability in basal expression and activity of a broad scope of xenobiotic metabolism genes in the human liver. *Drug Metab. Dispos.* 35, 1700–1710
21. Ring, H.Z. *et al.* (2007) Gene-gene interactions between CYP2B6 and CYP2A6 in nicotine metabolism. *Pharmacogenet. Genomics* 17, 1007–1015
22. Bloom, A.J. *et al.* (2013) CYP2B6 non-coding variation associated with smoking cessation is also associated with differences in allelic expression, splicing, and nicotine metabolism independent of common amino-acid changes. *PLoS One* 8, e79700
23. Bloom, A.J. *et al.* (2013) The contribution of common UGT2B10 and CYP2A6 alleles to variation in nicotine glucuronidation among European Americans. *Pharmacogenet. Genomics* 23, 706–716
24. Bloom, A.J. *et al.* (2013) Effects upon *in-vivo* nicotine metabolism reveal functional variation in FMO3 associated with cigarette consumption. *Pharmacogenet. Genomics* 23, 62–68
25. Murphy, S.E. *et al.* (2014) Nicotine N-glucuronidation relative to N-oxidation and C-oxidation and UGT2B10 genotype in five ethnic/racial groups. *Carcinogenesis* 35, 2526–2533
26. Patel, Y.M. *et al.* (2015) The contribution of common genetic variation to nicotine and cotinine glucuronidation in multiple ethnic/racial populations. *Cancer Epidemiol. Biomarkers Prev.* 24, 119–127
27. Tanner, J.-A. *et al.* (2017) Predictors of variation in CYP2A6 mRNA, protein, and enzyme activity in a human liver bank: influence of genetic and nongenetic factors. *J. Pharmacol. Exp. Ther.* 360, 129–139
28. Pérez-Stable, E.J. *et al.* (1998) Nicotine metabolism and intake in black and white smokers. *J. Am. Med. Assoc.* 280, 152–156
29. Benowitz, N.L. *et al.* (2002) Slower metabolism and reduced intake of nicotine from cigarette smoking in Chinese-Americans. *J. Natl. Cancer Inst.* 94, 108–115
30. Fagan, P. *et al.* (2015) Biomarkers of tobacco smoke exposure in racial/ethnic groups at high risk for lung cancer. *Am. J. Public Health* 105, 1237–1245
31. Wang, H. *et al.* (2015) Associations between genetic ancestries and nicotine metabolism biomarkers in the multiethnic cohort study. *Am. J. Epidemiol.* 182, 945–951
32. Chenoweth, M.J. *et al.* (2014) Known and novel sources of variability in the nicotine metabolite ratio in a large sample of treatment-seeking smokers. *Cancer Epidemiol. Biomarkers Prev.* 23, 1773–1782
33. Lerman, C. *et al.* (2010) Genetic variation in nicotine metabolism predicts the efficacy of extended-duration transdermal nicotine therapy. *Clin. Pharmacol. Ther.* 87, 553–557
34. Schnoll, R.A. *et al.* (2009) Nicotine metabolic rate predicts successful smoking cessation with transdermal nicotine: a validation study. *Pharmacol. Biochem. Behav.* 92, 6–11
35. Patterson, F. *et al.* (2008) Toward personalized therapy for smoking cessation: a randomized placebo-controlled trial of bupropion. *Clin. Pharmacol. Ther.* 84, 320–325
36. Ho, M.K. *et al.* (2009) Association of nicotine metabolite ratio and CYP2A6 genotype with smoking cessation treatment in African-American light smokers. *Clin. Pharmacol. Ther.* 85, 635–643
37. Chen, L.-S. *et al.* (2014) Pharmacotherapy effects on smoking cessation vary with nicotine metabolism gene (CYP2A6). *Addiction* 109, 128–137
38. Lerman, C. *et al.* (2015) Use of the nicotine metabolite ratio as a genetically informed biomarker of response to nicotine patch or varenicline for smoking cessation: a randomised, double-blind placebo-controlled trial. *Lancet Respir. Med.* 3, 131–138
39. Wells, Q.S. *et al.* (2017) Nicotine metabolism-informed care for smoking cessation: a pilot precision RCT. *Nicotine Tob. Res.* Published online October 14, 2017. http://dx.doi.org/10.1093/ntr/ntx235
40. U.S. Department of Health and Human Services (2010) *How Tobacco Smoke Causes Disease: The Biology and Behavioral Basis for Smoking-Attributable Disease. A Report of the Surgeon General,* Public Health Service, Office of the Surgeon General

41. Bush, W.S. and Moore, J.H. (2012) Chapter 11: Genome-wide association studies. *PLoS Comput. Biol.* 8, e1002822
42. Baurley, J.W. *et al.* (2016) Genome-wide association of the laboratory-based nicotine metabolite ratio in three ancestries. *Nicotine Tob. Res.* 18, 1837–1844
43. Patel, Y.M. *et al.* (2016) Novel association of genetic markers affecting CYP2A6 activity and lung cancer risk. *Cancer Res.* 76, 5768–5776
44. Chenoweth, M.J. *et al.* (2017) Genome-wide association study of a nicotine metabolism biomarker in African American smokers: impact of chromosome 19 genetic influences. *Addiction* Published online September 16, 2017. http://dx.doi.org/10.1111/add.14032
45. Whirl-Carrillo, M. *et al.* (2012) Pharmacogenomics knowledge for personalized medicine. *Clin. Pharmacol. Ther.* 92, 414–417
46. Neter, J. *et al.* (1996) *Applied Linear Statistical Models* (Vol. 4), Irwin
47. Hesterberg, T. *et al.* (2008) Least angle and l1 penalized regression: a review. *Stat. Surv.* 2, 61–93
48. Kyung, M. *et al.* (2010) Penalized regression, standard errors, and Bayesian lassos. *Bayesian Anal.* 5, 369–411
49. Tibshirani, R. (1996) Regression shrinkage and selection via the lasso. *J. R. Stat. Soc. Series B Stat. Methodol.* 58, 267–288
50. Hoerl, A.E. and Kennard, R.W. (2006) Ridge regression. In *Encyclopedia of Statistical Sciences*. John Wiley http://dx.doi.org/10.1002/0471667196.ess2280.pub2
51. Zou, H. and Hastie, T. (2005) Regularization and variable selection via the elastic net. *J. R. Stat. Soc. Series B Stat. Methodol.* 67, 301–320
52. Friedman, J. *et al.* (2001) *The Elements of Statistical Learning* (Vol. 1), Springer
53. Bergen, A.W. *et al.* (2016) Smokescreen translational analysis platform: prediction of nicotine metabolism and smoking cessation. In *NIDA Genetics Consortium Meeting*. National Institute on Drug Abuse
54. Hoeting, J.A. *et al.* (1999) Bayesian model averaging: a tutorial. *Statist. Sci.* 14, 382–401
55. Caruana, R. *et al.* (2004) Ensemble selection from libraries of models. In *Proceedings of the Twenty-First International Conference on Machine Learning*, pp. 18, ACM
56. Wasserman, L. (2000) Bayesian model selection and model averaging. *J. Math. Psychol.* 44, 92–107
57. Gelman, A. (2006) Prior distribution. *Encyclopedia of Environmetrics* (Vol. 4), Wiley http://dx.doi.org/10.1002/9780470057339.vap039
58. Baurley, J.W. *et al.* (2010) Discovery of complex pathways from observational data. *Stat. Med.* 29, 1998–2011
59. Gelman, A. *et al.* (2003) *Bayesian Data Analysis* (2nd edn), Chapman & Hall
60. Patkar, A.A. *et al.* (2009) Alterations in tryptophan and purine metabolism in cocaine addiction: a metabolomic study. *Psychopharmacology* 206, 479–489
61. Cho, Y.U. *et al.* (2017) Exploratory metabolomics of biomarker identification for the internet gaming disorder in young Korean males. *J. Chromatogr. B Analyt. Technol. Biomed. Life Sci.* 1057, 24–31
62. Oreland, L. *et al.* (2017) Personality as an intermediate phenotype for genetic dissection of alcohol use disorder. *J. Neural Transm. (Vienna)* Published online January 4, 2017. http://dx.doi.org/10.1007/s00702-016-1672-9
63. Foulds, J. *et al.* (2017) Dimensional personality traits and alcohol treatment outcome: a systematic review and meta-analysis. *Addiction* 112, 1345–1357
64. Volkow, N.D. *et al.* (2017) The conception of the ABCD study: from substance use to a broad NIH coll*aboration. *Dev. Cogn. Neurosci.* Published online October 10, 2017. http://dx.doi.org/10.1016/j.dcn.2017.10.002
65. Bogdan, R. *et al.* (2017) Imaging genetics and genomics in psychiatry: a critical review of progress and potential. *Biol. Psychiatry* 82, 165–175
66. Ramsey, A. (2015) Integration of technology-based behavioral health interventions in substance abuse and addiction services. *Int. J. Ment. Health Addict.* 13, 470–480
67. Agrawal, A. *et al.* (2016) Meta-analyses of genome-wide association data hold new promise for addiction genetics. *J. Stud. Alcohol Drugs* 77, 676–680
68. Acion, L. *et al.* (2017) Use of a machine learning framework to predict substance use disorder treatment success. *PLoS One* 12, e0175383
69. Suchting, R. *et al.* (2017) Using elastic net penalized cox proportional hazards regression to identify predictors of imminent smoking lapse. *Nicotine Tob. Res.* Published online September 7, 2017. http://dx.doi.org/10.1093/ntr/ntx201
70. Chih, M.-Y. *et al.* (2014) Predictive modeling of addiction lapses in a mobile health application. *J. Subst. Abuse Treat.* 46, 29–35
71. Vivanti, R. *et al.* (2017) Automatic detection of new tumors and tumor burden evaluation in longitudinal liver CT scan studies. *Int. J. Comput. Assist. Radiol. Surg.* 12, 1945–1957
72. Zong, N. *et al.* (2017) Deep mining heterogeneous networks of biomedical linked data to predict novel drug-target associations. *Bioinformatics* 33, 2337–2344
73. Institute of Medicine (2007) *Cancer Biomarkers: The Promises and Challenges of Improving Detection and Treatment,* The National Academies Press http://dx.doi.org/10.17226/11892
74. National Academies of Sciences, Engineering, and Medicine (2016) *Biomarker Tests for Molecularly Targeted Therapies: Key to Unlocking Precision Medicine,* The National Academies Press http://dx.doi.org/10.17226/21860
75. Institute of Medicine (2012) *Evolution of Translational Omics: Lessons Learned and the Path Forward,* The National Academies Press http://dx.doi.org/10.17226/13297
76. McShane, L.M. *et al.* (2013) Criteria for the use of omics-based predictors in clinical trials. *Nature* 502, 317–320
77. McShane, L.M. *et al.* (2013) Criteria for the use of omics-based predictors in clinical trials: explanation and elaboration. *BMC Med.* 11, 220
78. Institute of Medicine (2007) *The Learning Healthcare System: Workshop Summary,* The National Academies Press http://dx.doi.org/10.17226/11903
79. McGinnis, J.M. *et al.* (2013) *Best Care at Lower Cost: The Path to Continuously Learning Health Care in America,* National Academies Press
80. Institute of Medicine (2015) *Genomics-Enabled Learning Health Care Systems: Gathering and Using Genomic Information to Improve Patient Care and Research: Workshop Summary,* The National Academies Press http://dx.doi.org/10.17226/21707
81. Kranzler, H.R. *et al.* (2017) Precision medicine and pharmacogenetics: what does oncology have that addiction medicine does not? *Addiction* 112, 2086–2094
82. Salgado, R. *et al.* (2017) Societal challenges of precision medicine: bringing order to chaos. *Eur. J. Cancer* 84, 325–334
83. Amur, S. *et al.* (2015) Biomarker qualification: toward a multiple stakeholder framework for biomarker development, regulatory acceptance, and utilization. *Clin. Pharmacol. Ther.* 98, 34–46
84. Sauer, J.-M. and Porter, A.C. *et al.* (2017) Preclinical biomarker qualification. *Exp. Biol. Med.* Published online January 1, 2017. http://dx.doi.org/10.1177/1535370217743949
85. Amur, S.G. *et al.* (2015) Building a roadmap to biomarker qualification: challenges and opportunities. *Biomark. Med.* 9, 1095–1105
86. Bough, K.J. *et al.* (2013) Biomarkers for smoking cessation. *Clin. Pharmacol. Ther.* 93, 526–538
87. Bough, K.J. *et al.* (2014) Biomarkers for the development of new medications for cocaine dependence. *Neuropsychopharmacology* 39, 202–219

88. Dobson, A.J. and Barnett, A. (2008) *An Introduction to Generalized Linear Models* (3rd edn), Taylor & Francis
89. Zou, H. (2006) The adaptive lasso and its oracle properties. *J. Am. Stat. Assoc.* 101, 1418–1429
90. Fan, J. and Li, R. (2001) Variable selection via nonconcave penalized likelihood and its oracle properties. *J. Am. Stat. Assoc.* 96, 1348–1360
91. Zhang, C.-H. *et al.* (2010) Nearly unbiased variable selection under minimax concave penalty. *Ann. Stat.* 38, 894–942
92. Breheny, P. and Huang, J. (2011) Coordinate descent algorithms for nonconvex penalized regression, with applications to biological feature selection. *Ann. Appl. Stat.* 5, 232
93. Candes, E. and Tao, T. (2007) The Dantzig selector: statistical estimation when *p* is much larger than *n*. *Ann. Stat.* 35, 2313–2351
94. Li, X. *et al.* (2015) The flare package for high dimensional linear regression and precision matrix estimation in R. *J. Mach. Learn. Res.* 16, 553–557