Interactions of Genetic and Environment Scores: Alternating Lasso Regularization Avoids
Overfitting and Finds Interpretable Scores

Philipp Doebler

TU Dortmund University

Anna Doebler

University of Duisburg-Essen

Philip Buczak & Andreas Groll

TU Dortmund University

Author Note

Corresponding author contact information: Prof. Dr. Philipp Doebler, TU Dortmund University, Department of Statistics, Chair of Statistical Methods in the Social Sciences, Room M/737, Vogelpothsweg 87, D-44227 Dortmund, Germany; Email: doebler@statistik.tu-dortmund.de; Phone: +49 (0) 231/755-8259; Fax: +49 (0) 231/755-3918.

The work of the second author was supported by DFG-grant Do 2035/1-1.

Abstract

Regression models with interaction terms are common models for moderating relationships. When several predictors from one group, e.g., genetic variables, are potentially moderated by several predictors from another, e.g., environmental variables, many interaction terms result. This complicates model interpretation, especially when coefficient signs point in different directions. By first forming a score for each group of predictors, the interaction model's dimension is severely reduced. The hierarchical score model is an elegant one step approach: Score weights and regression model coefficients are estimated simultaneously by an alternating optimization (AO) algorithm. Especially in high dimensional settings, scores remain an effective technique to reduce interaction model dimension, and we propose regularization to ensure sparsity and interpretability of the score weights. A non-trivial extension of the original AO algorithm is presented, which adds a lasso penalty, resulting in the alternating lasso optimization algorithm (ALOA). The hierarchical score model with ALOA is an interpretable statistical learning technique for moderation in potentially high dimensional applications, and encompasses generalized linear models for the main interaction model. In addition to the lasso regularization, a screening procedure called regularization and residualization (RR) is proposed to avoid spurious interactions. ALOA tuning parameter choice and the RR screening procedure are investigated by simulations, and an illustrative application to lifetime depression risk and gene $\times$ environment interactions is provided.

Keywords: gene $\times$ environment interactions; regularization; Lasso; pre-screening of genetic variables; model selection; gene scores; alternating optimization

Interactions of Genetic and Environment Scores: Alternating Lasso Regularization Avoids Overfitting and Finds Interpretable Scores

Draft version, 09/07/20. This paper has not been peer reviewed.

Introduction

Interaction terms are included in statistical models when the effect of one covariate on the outcome depends on the value of another covariate, e.g., to model group-specific effects of a continuous variable. Important motivating instances are effects of (potentially malleable) environmental factors moderated by stable genetic factors. Genotype specific effects of environmental conditions have become an active research field within psychology (Manuck & McCaffery, 2014), not limited to clinical research questions (e.g. Haefel et al., 2008), but also in personality research and developmental psychology (e.g. Fraley et al., 2013). A common statistical approach to describe differential effects of an environment is via regression models with interaction terms. When the set of genomic variables (G) and environmental variables (E) is guided by strong theory, the dimension of G and E might be low. However, the falling cost of genome wide analysis (Mardis, 2011) is an opportunity to answer novel, broad, or exploratory research questions by combining social science experimental designs with genomic assays (Fletcher & Conley, 2013), leading to large number of G and E variables. A high dimensionality of G and E variables poses severe problems for regression models with many interactions (Jolicœur-Martineau et al., 2019). Another challenge are false positive results arising when large models are not held in check by appropriately low p -values, which should depend on the study design (e.g. Fadista et al., 2016). The interpretation of models with many parameters is difficult and, hence, strategies for model selection are valuable tools (e.g. Waldmann et al., 2013), i.e., irrelevant variables are removed from a model.

Another way to forestall interpretation problems is to form a univariate score from several or many genetic variables, a so called *genetic (risk) score*, for which Belsky and

Israel (2014) mention various alternative labels like “multi-locus genetic profile”, “polygenic risk score”, “allelic score”, “SNP score”, “genotype score” or “genetic prediction score”.

Regardless of the label, genetic scores are formed based on biological understanding of the studied phenotype, or based on a genome-wide search, and while weights of each genetic variable can be equal, a wide range of statistical methods is available (Jolicoeur-Martineau et al., 2019).

Gene $\times$ environment score models

In the motivating situation of high-dimensional $G \times E$ interactions the following properties are desirable for a quantitative method: (i) Being able to check whether any interaction is present at all. (ii) Selection of only relevant genetic and environmental variables, and, related to this, (iii) dimension reduction of genetic and environmental variables to ease interpretation. All these characteristics are part of a hierarchical score model recently proposed by Jolicoeur-Martineau et al. (2019). The basic strategy is to form a gene and an environment score of the form $G = \sum_{u=1}^{k_G} p_u g_u$ and $E = \sum_{v=1}^{k_E} q_v e_v$, i.e., the scores are weighted sums and, ideally, only a few of the weights $p_1, \dots, p_{k_G}, q_1, \dots, q_{k_E}$ are non-zero (sparsity). At the score level, a regression model is built, e.g., a generalized linear regression model (McCullagh & Nelder, 1989). Since scores are seen as indicators or measurements of a latent genetic/environmental trait, the hierarchical score model is named latent environment gene interaction (LEGIT) model by Jolicoeur-Martineau et al. (2019), but the approach is much more general and not restricted to G and E variables only.

Why do score models with sparse scores help? The number of parameters k can approach or exceed the sample size n when all interaction terms of the form $g_u e_v$ were included systematically in a model, invalidating strategies like two-way ANOVA or other classic linear models to address (i). When only scores are entered in a regression model, the coefficient of the score-level interaction is the only quantity governing interactions. In score models, precautions have to be made to obtain sparse scores. One strategy is to pair a

hierarchical score model with a model selection strategy (e.g. Goldberg, 1989). Abstracting a bit, the model selection strategy controls the false discovery rate to ensure (ii) and (iii).

Regularization strategies are often computationally advantageous in comparison to sequential model choice, and also work in case of $k \gg n$ (Hastie et al., 2009).

Regularization is also attractive, since in its original formulation, the LEGIT-methodology can lead to spurious selection of main effects that are not part of the interaction, and currently available implementations (Jolicoeur-Martineau, 2019) build on computationally intensive variable selection rationales. The main aims of the current work are to describe the spurious selection in more detail, and to employ lasso regularization (Tibshirani, 1996) to tackle the challenges in a high dimension. To this end, we start by reviewing the hierarchical score model of Jolicoeur-Martineau et al. (2019) in the following section. A regularization and residualization technique to avoid spurious interactions with main-effect-only variables is evaluated with the help of simulations. Subsequently, a novel Alternating Lasso Optimization Algorithm (ALOA) is presented in detail. A simulation experiment evaluates ALOA’s model selection properties based on cross-validation and a variant of the Bayesian Information Criterion (BIC; Schwarz, 1978). The technique is illustrated with an application to $G \times E$ data from depression research.

$G \times E$ interactions with and without scores

Let $g_1, \dots, g_{k_G}$ denote genetic variables, which might be 0-1-coded, but could also be continuous, and let $e_1, \dots, e_{k_E}$ denote environmental variables, not necessarily standardized. We denote vectors of length n of these and other variables in boldface, e.g., $\mathbf{g}_i = (g_{1i}, \dots, g_{ni})'$. We aim to predict an outcome Y conditional on genes, environmental variables and potential interactions. As a model with many individual interaction terms is hard to interpret, the strategy of first forming scores for the genetic and environmental variables and then studying the interaction of the scores is a way to reduce the complexity of the interaction model. Scores are weighted sums of the original variables, so the weights

estimated in the following will reflect the relative influence of each variable. That is why standardization of continuous G - or E -variables might help to assess the relative weight of the variables when they enter scores.

Hierarchical score model with alternating optimization (LEGIT)

For dimension reduction with gene and environment scores, the hierarchical score model (Jolicoeur-Martineau et al., 2019) is assumed in the following:

$$\mathbf{Y} = \gamma_0 + \gamma_1 \mathbf{G} + \gamma_2 \mathbf{E} + \gamma_3 \mathbf{G} \odot \mathbf{E} + \boldsymbol{\epsilon}, \quad \text{where} \quad (1)$$

$$\mathbf{G} = \sum_{u=1}^{k_G} p_u \mathbf{g}_u \quad \text{and} \quad (2)$$

$$\mathbf{E} = \sum_{v=1}^{k_E} q_v \mathbf{e}_v, \quad (3)$$

for (partly unknown) regression coefficients $\gamma_0, \dots, \gamma_3$, unknown weights $(p_u)_u, (q_v)_v$, and a vector of errors $\boldsymbol{\epsilon} = (\epsilon_1, \dots, \epsilon_n)'$. Here, $\odot$ denotes the elementwise multiplication of two vectors (Hadamard product). One can set $\gamma_1 = \gamma_2 = 1$ for identifiability (or alternatively $\sum_u |p_u| = 1$ and $\sum_v |q_v| = 1$). When one identifies the model by assuming that $\gamma_1 = \gamma_2 = 1$, the sums of the absolute values of $\mathbf{p} = (p_u)_{u=1, \dots, k_G}$ and $\mathbf{q} = (q_v)_{v=1, \dots, k_E}$ are unconstrained. Note that this hierarchical model is not linear in its parameters, as products of regression coefficients and weights appear in the equation for $\mathbf{Y}$ after substituting for $\mathbf{G}$ and $\mathbf{E}$. This entails that standard techniques for linear models cannot be applied directly. Similarly, the corresponding least squares criterion is not convex, limiting the usefulness of techniques for convex optimization. It is however biconvex (Section 5.9 Hastie et al., 2015), and by exploiting this property Jolicoeur-Martineau et al. (2019) show that an alternating optimization (AO) scheme finds a local minimum. In combination with multiple random starting values, AO allows to obtain least squares estimates of the hierarchical score model, when k_E and k_G are smaller than the sample size n .

AO proceeds iteratively, starting from some initial values by fitting the linear model

in Equation 1 for the current values of $(p_u)_u$ and $(q_v)_v$. This updates the estimates of the regression coefficients γ_i , $i = 0, 1, 2, 3$. To update one of the sets of weights, the other parameters are considered fixed, the equation for $\mathbf{G}$ (or $\mathbf{E}$, respectively) is plugged into Equation 1 and the resulting *linear* model (with certain offsets) is fit. This updates the estimates of $(p_u)_u$ and $(q_v)_v$, and decreases the sum of squares until convergence. More details are provided by Jolicoeur-Martineau et al. (2019), who also provide a reference implementation in the R-package **LEGIT** (Jolicoeur-Martineau, 2019) for an even more general model with interactions.

Spurious interactions and how to avoid them

When a genetic variable does not interact with any of the environment variables, it could nevertheless explain variance in the outcome, i.e., a main effect of the genetic variable is present, but no interaction. The same could be true for any environment variable. The model in Equation 1 is on the level of scores, so by intention it is not fine grained enough to discern main and interaction effects on the variable level. However, variable selection methods might choose a genetic or an environment variable with a main effect on $\mathbf{Y}$ but without any interactions to contribute to the G - or E -score, respectively. We interpret this as a spurious and, hence, undesirable selection in a situation where researchers are interested in interactions, since a G -score (or E -score, respectively) is prone to misinterpretations in terms of the moderating effects of main-effects-only variables. Consequently, it is important to be able to identify variables without any interactions. The implementation in the R-package **LEGIT** builds on natural evolutionary variable selection (NEVS) if the number of variables is large (Jolicoeur-Martineau, 2019). NEVS is a genetic algorithm (Goldberg, 1989; Holland, 1975) employing heuristics inspired by biological evolution to find a model.

We propose a regularization and residualization (RR) technique to avoid spurious interactions: To identify main-effect-only variables, an initial variable screening step with

lasso regularization is proposed (regularization). Instead of the original $\mathbf{Y}$, a partial variable that is uncorrelated to variables without interactions is then employed (residualization). The following will illustrate the benefits of RR. Because the reference implementation employs NEVS, we evaluate this variable selection technique with and without RR.

Prescreening with lasso regularization. The linear model for prescreening has $k := k_G + k_E + k_G k_E$ covariates and an intercept:

$$\mathbf{Y} = \beta_0 + \sum_{u=1}^{k_G} \beta_u^G \mathbf{g}_u + \sum_{u=1}^{k_E} \beta_u^E \mathbf{e}_u + \sum_{u=1}^{k_G} \sum_{v=1}^{k_E} \beta_{uv} \mathbf{g}_u \mathbf{e}_v + \boldsymbol{\epsilon}. \quad (4)$$

In general, this model has many more coefficients than the LEGIT model, and there are no scores. To avoid confusion, we do not use γ , p nor q to denote the model coefficients. By $\boldsymbol{\beta}$ we denote a vector of length k of all regression coefficients apart from the intercept β_0 and by $\mathbf{X}$ we denote an appropriate design matrix with rows $\mathbf{x}_i$, $i = 1, \dots, n$. The model can hence be written in a more familiar fashion as $\mathbf{Y} = \beta_0 + \mathbf{X}\boldsymbol{\beta} + \boldsymbol{\epsilon}$. Unfortunately, for realistic sample sizes n , $k \gg n$ is possible, so that least squares is not an appropriate estimation approach. We hence propose a lasso regularization approach (Tibshirani, 1996), i.e., estimates are the minimum of the objective function

$$\frac{1}{2} \sum_{i=1}^n (y_i - \beta_0 - \boldsymbol{\beta}' \mathbf{x}_i)^2 + \lambda \sum_{u=1}^k |\beta_u| = \frac{1}{2} \|\mathbf{y} - \beta_0 - \boldsymbol{\beta} \mathbf{X}\|_2^2 + \lambda \|\boldsymbol{\beta}\|_1. \quad (5)$$

Here, λ is a tuning parameter, which controls the amount of shrinkage of the regression parameters towards 0. The larger λ , the stronger the shrinkage. The approach allows to shrink coefficients to zero, essentially removing them from the model. Model selection and sparsity are built into lasso estimation, making it a one step procedure in comparison to least squares estimation in combination with a variable selection method. Instead of fixing λ a priori, we employ leave one out cross-validation (LOO-CV) to choose λ . We choose the largest λ for which the cross-validated mean squared error (MSE) is within one standard

error of the minimal cross-validated MSE as implemented in `glmnet` (Friedman et al., 2010). This model is hence more restrictive than the model with minimal cross-validated MSE and typically more sparse. In this context, note that the penalized criterion in Equation 5 can result in a whole group of coefficients being shrunk to zero, so especially all interaction terms could be removed.

Assume a lasso fit of the model in Equation 4 has been obtained, i.e., estimates $(\hat{\beta}_u^G)_u$, $(\hat{\beta}_v^E)_v$ and $(\hat{\beta}_{uv})_{uv}$ with maybe some of the regression coefficients shrunk to zero. The sets

$$M_G := \{u : \hat{\beta}_u^G \neq 0 \text{ and for all } v, \hat{\beta}_{uv} = 0\} \quad (6)$$

$$M_E := \{v : \hat{\beta}_v^E \neq 0 \text{ and for all } u, \hat{\beta}_{uv} = 0\} \quad (7)$$

contain the indices of genetic and environment variables which only contribute a main effect.

We propose to partial out the main effects with indices in $M_G \cup M_E$ of $\mathbf{Y}$, so that a subsequent investigation of the $G \times E$ -score interactions does not spuriously select such variables to one of the scores. A variant would be to remove the variables entirely, but other variables covarying with the removed variables might then be selected instead. Assuming that the number of indices in $M_G \cup M_E$ is small, obtain for the linear model of the form

$$\mathbf{Y} = \delta_0 + \sum_{u \in M_G} \delta_u^G \mathbf{g}_u + \sum_{v \in M_E} \delta_v^E \mathbf{e}_v + \boldsymbol{\epsilon} \quad (8)$$

least squares estimates $\hat{\delta}_0$, $(\hat{\delta}_u^G)_{u \in M_G}$, $(\hat{\delta}_v^E)_{v \in M_E}$. The residuals

$$\mathbf{Y}_{\text{partial}} := \mathbf{Y} - \hat{\delta}_0 + \sum_{u \in M_G} \hat{\delta}_u^G \mathbf{g}_u + \sum_{v \in M_E} \hat{\delta}_v^E \mathbf{e}_v \quad (9)$$

are empirically uncorrelated with all variables with indices in $M_G \cup M_E$. If $|M_G| + |M_E|$ is large, again a lasso could be used to find partial variables, though the empirical

uncorrelatedness does not necessarily hold. Employing $\mathbf{Y}_{partial}$ in subsequent steps instead of $\mathbf{Y}$ should hence avoid spurious selection. We confirm this in the following simulation experiment.

Simulation experiment 1: Spurious interactions and lasso based interaction prescreening

Data generation. In all scenarios, we generated data for $n = 400$ simulees on a single (i.e., $k_E = 1$) standard normally distributed environment variable. The genes were assumed to be independent standard normal variables. The following two “factors” were fully crossed: $k_G \in \{4, 100\}$ and the data generating model, which was one of:

1. $Y = \epsilon$ (intercept only),
2. $Y = (g_1 + g_2) \cdot e + \epsilon$ (interactions only),
3. $Y = g_3 + g_4 + \epsilon$ (main effects only), and
4. $Y = g_3 + g_4 + (g_1 + g_2) \cdot e + \epsilon$ (interactions and main effects),

where ϵ was standard normal in all cases. This resulted in eight scenarios. If a $G \times E$ -interaction is the target of the investigation, selection of g_1, g_2 is desirable in models 2. and 4, but not the selection of g_3 and g_4 . If g_3 or g_4 are selected, the contribution to the score might be misinterpreted as an interaction with e . Pseudo-random sampling was done in R 3.6.1 (R Core Team, 2019) with the help of the package `rlecuyer` (Sevcikova & Rossini, 2015) to avoid overlap of pseudo-random number streams on multiple cores.

Analysis of simulated data. NEVS as implemented in the function `nes_var_select` in the R-package LEGIT (Jolicoeur-Martineau, 2019) with default tuning parameters was performed. NEVS builds on the AO scheme of Jolicoeur-Martineau et al. (2019), which is repeatedly performed for different variable subsets. The model with the smallest corrected Akaike Information Criterion (AICc, the default method; Hurvich & Tsai, 1989) value is finally selected. It was recorded which genetic variables were selected.

In each replication, the same simulated data was also subjected to the RR procedure, followed by NEVS, so that the effect of RR could be studied. Though NEVS is computationally demanding (especially when $k_G = 100$), 1000 replications of all scenarios were simulated.

Simulation results. NEVS with and without RR converged in all 1000 replications. Tables 1 and 2 contain the selection rates (out of 1000 replications) for genetic variables in the model equation. Note that the columns labeled “other” will include up to 100 variables (columns $g_1, \dots, g_4$ are only used if the variable appears in the model equation). In data generating model 4 all genetic variables are in the model equation when $k_G = 4$, and there is no entry in the “other” column in that case. As all genetic variables are independent standard normal, another summary would not have been more informative, i.e., because of the exchangeability of all genetic variables not in the model equation, individual descriptive statistics are not more meaningful than the mean selection rate in the “other” column.

Table 1

Results of Simulation Experiment 1 (natural evolutionary variable selection (NEVS) without regularization and residualization (RR)).

Design			Selection rate				
k_G	n	Model equation	g_1	g_2	g_3	g_4	other
4	400	$Y = \epsilon$	-	-	-	-	0.32
100	400	$Y = \epsilon$	-	-	-	-	0.18
4	400	$Y = (g_1 + g_2) \cdot e + \epsilon$	1.00	1.00	-	-	0.15
100	400	$Y = (g_1 + g_2) \cdot e + \epsilon$	1.00	1.00	-	-	0.16
4	400	$Y = g_3 + g_4 + \epsilon$	-	-	1.00	1.00	0.14
100	400	$Y = g_3 + g_4 + \epsilon$	-	-	1.00	1.00	0.16
4	400	$Y = g_3 + g_4 + (g_1 + g_2) \cdot e + \epsilon$	0.84	0.84	0.87	0.87	-
100	400	$Y = g_3 + g_4 + (g_1 + g_2) \cdot e + \epsilon$	0.73	0.74	0.80	0.80	0.15

For the studied scenarios, Table 1 indicates that NEVS without RR will select each genetic variable not in the model equation in up to 32% of cases ($k_G = 4$, intercept only) to around 14% ($k_G = 4$, main effects only). NEVS was always correct in the interaction model

w.r.t. g_1 and g_2 , but did include other genetic variables in 15% and 16% of replications. Unfortunately, in model 4 with spurious main effects and interaction effects, occasionally the interacting genetic variables are not selected (they were selected only in 84% ($k_g = 4$) and 73%-74% ($k_G = 100$)), with the main-effects-only variables having even slightly higher selection probabilities (87% and 80%).

Table 2

Results of Simulation Experiment 1 (natural evolutionary variable selection (NEVS) with regularization and residualization (RR)).

Design			Selection rate				
k_G	n	Model equation	g_1	g_2	g_3	g_4	other
4	400	$y = \epsilon$	-	-	-	-	0.32
100	400	$y = \epsilon$	-	-	-	-	0.18
4	400	$y = (g_1 + g_2) \cdot e + \epsilon$	1.00	1.00	-	-	0.15
100	400	$y = (g_1 + g_2) \cdot e + \epsilon$	1.00	1.00	-	-	0.16
4	400	$y = g_3 + g_4 + \epsilon$	-	-	0.14	0.20	0.41
100	400	$y = g_3 + g_4 + \epsilon$	-	-	0.09	0.08	0.18
4	400	$y = g_3 + g_4 + (g_1 + g_2) \cdot e + \epsilon$	1.00	1.00	0.16	0.16	-
100	400	$y = g_3 + g_4 + (g_1 + g_2) \cdot e + \epsilon$	1.00	1.00	0.14	0.16	0.15

Table 2 indicates no difference for models 1 (intercept only) and 2 (interaction only) between NEVS with and without RR, with no differences in selection rate up to the second decimal place. Differences emerge in the other two data generating models: In the main effects only model, RR reduces the selection of the main-effects-only variables from 100% to values between 8% and 20%, which was desired. In the spurious main effects and interaction model 4, selection of the spurious main effects is reduced to 14%-16%, while the interacting variables are always selected, though there is no main effect of them in the data generating model. So while still somewhat liberal, NEVS's behaviour can be modified by RR in the intended direction. We note that next to AICc, more conservative criteria are available in `nes_var_select`. Regardless of the way NEVS is tuned, it is a computationally demanding algorithm.

Alternating optimization with the lasso

We now turn to a different approach intended to find sparse interaction models. We modify AO by introducing lasso constraints on the regression parameters for the weights of the score equations. This way, sparse sets of weights $(p_u)_u$ and $(q_v)_v$ can be obtained, i.e., some of the weights are zero. Similar to model selection strategies employed by Jolicoeur-Martineau et al. (2019), this leads to scores with easier interpretation, but is less burdensome in terms of computation and handles the $k_G + k_E > n$ case.

Model specification and identification

Assume the model from Equations 1, 2 and 3 holds, which we reiterate here for convenience:

$$\mathbf{Y} = \gamma_0 + \gamma_1 \mathbf{G} + \gamma_2 \mathbf{E} + \gamma_3 \mathbf{G} \odot \mathbf{E} + \boldsymbol{\epsilon} \quad \text{where} \quad (10)$$

$$\mathbf{G} = \sum_{u=1}^{k_G} p_u \mathbf{g}_u \quad \text{and} \quad (11)$$

$$\mathbf{E} = \sum_{v=1}^{k_E} q_v \mathbf{e}_v, \quad (12)$$

for regression coefficients $\gamma_0, \dots, \gamma_3$ and unknown weights $(p_u)_u, (q_v)_v$. In the following, we identify the model by assuming that $\gamma_1 = \gamma_2 = 1$, so none of the weights $\mathbf{p} = (p_u)_{u=1, \dots, k_G}$ and $\mathbf{q} = (q_v)_{v=1, \dots, k_E}$ are fixed. This is slightly different from Jolicoeur-Martineau et al. (2019), but the original parametrization is easy to obtain.

The lasso estimator as a constrained optimization problem. Before we state the algorithm and discuss its convergence properties, we discuss the constraints. A key fact from the lasso literature is the following equivalence (e.g., Hastie et al., 2009, p. 68): Given a linear model $y_i = \beta_0 + \beta_1 x_{i1} + \dots + \beta_p x_{ip} + \epsilon_i$, $i = 1, \dots, n$, and a tuning parameter λ , the

lasso estimate $\hat{\beta}^{lasso}$ solves the lasso problem in Lagrangian form, i.e.,

$$\hat{\beta}^{lasso} = \operatorname{argmin}_{\beta} \left\{ \frac{1}{2} \sum_{i=1}^n (y_i - \beta_0 - \sum_{j=1}^p \beta_j x_{ij})^2 + \lambda \sum_{j=1}^p |\beta_j| \right\}. \quad (13)$$

The above form of the lasso estimator might be most familiar and is also stated above in Equation 5. A constrained optimization problem is equivalent: A tuning parameter t exists, such that

$$\hat{\beta}^{lasso} = \operatorname{argmin}_{\beta} \left\{ \sum_{i=1}^n (y_i - \beta_0 - \sum_{j=1}^p \beta_j x_{ij})^2 \right\}, \quad (14)$$

$$\text{subject to } \sum_{j=1}^p |\beta_j| < t. \quad (15)$$

For our purposes, this second version is more convenient for two reasons: (i) As $\sum_{u=1}^p |\beta_u|$ will correspond to the sum of the score weights, the interpretation of t is clearer than that of λ in the Lagrangian version. (ii) The design matrices of the lasso problems change from one iteration of the proposed algorithm to the next. Consequently, a fixed value of λ would change its interpretation. Hence, we prefer to find for a given t_0 a λ_0 satisfying the constraint with the following method: Determine lasso estimates for a sequence of λ -values, starting with a very low value of $10^{-5}\lambda_{\max}$, where $\lambda_{\max} > 0$ is so large as to produce an intercept only model¹. Finding such a solution path can be numerically advantageous over fitting the lasso for single values because of warm starts (Friedman et al., 2010). From the resulting coefficients $\beta_0(\lambda), \dots, \beta_p(\lambda)$ for each λ in the sequence, calculate $t(\lambda) = \sum_{u=1}^p |\beta_u(\lambda)|$. Use cubic spline interpolation of the resulting $t(\lambda)$ -sequence to find λ_0 such that $\sum_{u=1}^p |\beta_u(\lambda_0)| \approx t_0$ or, if t_0 is large, pick the minimal λ , for which $\sum_{u=1}^p |\beta_u(\lambda)| < t_0$. This trick allows us to switch from the constrained version of the lasso to the Lagrangian version, which is typically implemented.

¹Starting with $\lambda = 0$ gives the least squares solution, which however in the presence of many covariates might not exist, while the value $10^{-5}\lambda_{\max}$ will typically lead to estimates close to least squares.

Minimizing the MSE subject to constraints for the genetic and environment score. For full generality, let $(w_u^g)_u$ and $(w_v^e)_v$ denote a priori weights ≥ 0 for the G- and E-variables, such that $\sum_u w_u^g = k_G$ and $\sum_v w_v^e = k_E$. These weights will govern the strength of the penalty on the variable level. The larger the weight, the stronger the penalty. One can of course use weights equal to one for each variable, but if the variables differ in variance, variables with large variance can dominate the resulting penalty, so z -standardization of all variables is recommended (and the default in many lasso implementations). If a weight is equal to zero, the variable is not penalized at all, which is important for cases with strong a priori support for a variable. The constraint is placed on the weighted sums of the absolute values of the score weights, i.e., there are constants $t_g \geq 0$ and $t_e \geq 0$, such that

$$0 \leq \sum_u w_u^g |p_u| \leq t_g \quad \text{and} \quad 0 \leq \sum_v w_v^e |q_v| \leq t_e. \quad (16)$$

The following algorithm minimizes the mean squared error of $\mathbf{Y}$ in Equation 10 subject to the above constraints.

Alternating Lasso Optimization Algorithm

For the basic algorithm presented in the following, only one value each of t_g and t_e is considered, so initially, the aim is not to present an algorithm for a whole sequence or grid of t -values. The final choice of (t_g, t_e) is informed by the Bayesian Information Criterion (BIC; Schwarz, 1978) or a cross validation procedure outlined further below.

The Alternating Lasso Optimization Algorithm (ALOA) is assumed to have converged when parameter estimates (including those for γ_i , $i = 0, 1, 2, 3$) do not change more than a (small) constant $\Delta > 0$ from one iteration to the next. So Δ can be seen as an additional tuning parameter, which governs the number of steps, very similar to the setup in Jolicoeur-Martineau et al. (2019). We use $\Delta = 10^{-5}$ as a default. The algorithm proceeds as follows:

0. Random Initialization. Randomly generate $(p_u^{(0)})_u$ and $(q_v^{(0)})_v$ subject to the

constraint (16). Set $\mathbf{G}^{(0)} = \sum_{u=1}^{k_G} p_u^{(0)} \mathbf{g}_u$ and $\mathbf{E}^{(0)} = \sum_{v=1}^{k_E} q_v^{(0)} \mathbf{e}_v$. Find least squares estimates $\gamma_0^{(0)}$ and $\gamma_3^{(0)}$ of the linear model $\mathbf{Y} = \gamma_0 + \mathbf{G}^{(0)} + \mathbf{E}^{(0)} + \gamma_3 \mathbf{G}^{(0)} \odot \mathbf{E}^{(0)} + \epsilon$.

Let $\boldsymbol{\theta}^{(0)}$ denote a vector containing all elements of $\boldsymbol{\gamma}^{(0)}$, $\mathbf{p}^{(0)}$ and $\mathbf{q}^{(0)}$ and let formally $\boldsymbol{\theta}^{(-1)}$ be a vector of the same length with all elements equal to ∞ .

Loop While $\max |\boldsymbol{\theta}^{(l)} - \boldsymbol{\theta}^{(l-1)}| \geq \Delta$, repeat the following steps:

1. Update gene score model. Substituting (11) into (10) and re-arranging yields

$$\mathbf{Y} = \gamma_0^{(l)} + \sum_{u=1}^{k_G} p_u \mathbf{g}_u + \mathbf{E}^{(l)} + \gamma_3^{(l)} \mathbf{E}^{(l)} \odot \sum_{u=1}^{k_G} p_u \mathbf{g}_u + \epsilon \quad (17)$$

$$= \underbrace{\gamma_0^{(l)} + \mathbf{E}^{(l)}}_{\text{offset}} + \sum_{u=1}^{k_G} p_u \underbrace{\mathbf{g}_u \odot (1 + \gamma_3^{(l)} \mathbf{E}^{(l)})}_{=:\mathbf{g}_u^{(l)}} + \epsilon, \quad (18)$$

where $(p_u)_u$ are the only free parameters of this model. Note there is no intercept but an offset and the synthetic covariates $\mathbf{g}_u^{(l)}$ depend on l . Obtain an unstandardized lasso fit with tuning parameter λ_G of model (18) subject to the constraint (16) for the G variables². Note that the constraint for the E variables is met (by choice of the starting values or the previous iteration). Calculate $\mathbf{G}^{(l+1)}$ from (11) using these weights³.

2. Update of $\boldsymbol{\gamma}$. Obtain a least squares fit of model (10) using $\mathbf{G}^{(l+1)}$ and $\mathbf{E}^{(l)}$ for $\mathbf{G}$ and $\mathbf{E}$, respectively. Call the resulting estimates $\boldsymbol{\gamma}^{(l+0.5)}$.

3. Update of environment score. Use the very same steps as for the update of the gene score using $\mathbf{G}^{(l+1)}$ and $\boldsymbol{\gamma}^{(l+0.5)}$ for the calculation of offset and synthetic covariates. Call the resulting estimates $(q_v^{(l+1)})_v$ and calculate $\mathbf{E}^{(l+1)}$ from them.

²Note that when weights are used, the reference implementation in `glmnet` rescales penalization factors to sum to k_G .

³Due to the lasso penalty, it could be that all weights trivialize in this step, i.e. $\mathbf{p}^{(l+1)} = \mathbf{0}$. Consequently $\mathbf{G}^{(l+1)} = \mathbf{0}$ and, hence, the main model (10) reduces to a main effects model for the E-variables only. In case additionally $\mathbf{E}^{(l^*)}$ trivializes for some l^* , the main model becomes an intercept only model.

4. **Update of γ .** Obtain a least squares fit of model (10) using $\mathbf{G}^{(l+1)}$ and $\mathbf{E}^{(l+1)}$ for $\mathbf{G}$ and $\mathbf{E}$, respectively. Call the resulting estimates $\gamma^{(l+1)}$.
5. Let $\boldsymbol{\theta}^{(l+1)}$ denote the vector with all elements from $\gamma^{(l+1)}$, $\mathbf{p}^{(l+1)}$ and $\mathbf{q}^{(l+1)}$ and increase l by one.

Convergence and local optima. In the Appendix we prove that each iteration decreases the sum of squared errors (SSE) of $\mathbf{Y}$ in Equation 10 until convergence of the sequence

$$\text{SSE}^{(l)} = \sum_{i=1}^n (y_i - \hat{y}_i^{(l)})^2, \quad (19)$$

where $\hat{y}_i^{(l)}$ is obtained by plugging $\boldsymbol{\theta}^{(l)}$ into Equations 10, 11 and 12. As a consequence, a local optimum respecting the constraints exists (cf. Jolicoeur-Martineau et al., 2019). It is advisable to run the algorithm multiple times with different pseudo-random starting values, as it can get stuck in a local minimum⁴.

Numerical stability. As we determine λ_g and λ_e from the given t_g - and t_e -values by spline interpolation, occasionally a λ -value that is a bit too small will result from calculations. This will, on some of the iterations, allow for an even smaller SSE than the final one at convergence, typically by less than by one tenth of a percent. Another subtle numerical issue is the potential absence of main effects: If $(p_u)_u$ and/or $(q_v)_v$ are (close to) 0 in an iteration, the interaction effect estimate $\gamma_3^{(l)}$ can become very large. However, from an applied perspective, it should occur only very rarely that main effects are completely absent.

Choice of tuning parameters. Regardless of distributional assumptions, the tuning parameters t_g , t_e can be chosen with the help of K -fold cross-validation (CV). If a Gaussian error distribution for ϵ in Equation 10 is plausible and if the ALOA converges to the global minimum, the ALOA estimate is the maximum likelihood solution subject to the

⁴A “main effects only”-initialization is an alternative to random starting values: Fit a genetic main effects only model with constraint (16) for $(p_u^{(0)})_u$. Employ $\mathbf{G}^{(0)}$ as an offset in a main effects only lasso fit for the environment model. Finally, initialize γ with least squares and proceed with the loop.

constraints in Equation 16. With this distributional assumption in mind and in analogy to what has been proposed for the lasso for linear (mixed) models (Schelldorfer et al., 2011), we define the BIC criterion as

$$BIC =: -2\ell + \ln(n)k, \quad (20)$$

where ℓ is the log-likelihood evaluated at the ALOA estimate $\hat{\boldsymbol{\theta}}$, n is the sample size and k is equal to the number of non-zero elements of $\hat{\boldsymbol{\theta}}$. We see k as a measure for model complexity (cf. Zou et al., 2007). Specifically,

$$\ell = -\frac{n}{2} \ln(2\pi\hat{\sigma}^2) - \frac{1}{2\hat{\sigma}^2} \text{SSE}, \quad (21)$$

where SSE is calculated from the ALOA estimate $\hat{\boldsymbol{\theta}}$ and $\hat{\sigma}^2 = n^{-1}\text{SSE}$, so that ℓ simplifies to $\ell = -\frac{n}{2} \ln(2\pi\hat{\sigma}^2) - \frac{n}{2}$. As the BIC is computationally far less demanding than CV, computing BIC for a grid of pairs (t_g, t_e) is a faster alternative for determining tuning parameters.

Simulation Experiment 2: ALOA-based model selection

The experiment has the following purpose: Find out whether 10-fold CV is matched by BIC in terms of the tuning parameters determined, and, consequently, the identification of the relevant G - and E -variables. This is a crucial prerequisite for parameter recovery of the model. As CV and BIC will not align perfectly, the correlation of the predictions will quantify the level of agreement of the two methods.

For each replication, $n = 400$ data were generated from a LEGIT model with $k_E = 2$ independent environmental variables and $k_G = 4$ or $k_G = 100$ genetic variables, all of which were standard normally distributed. For different choices of γ_i , $i = 1, 2$, the scores were set as $G = \gamma_1\sqrt{0.5}(g_1 + g_2)$ and $E = \gamma_2 e_1$, so that both were centered with variances γ_i^2 ,

$i = 1, 2$. The simulated outcome was then generated as

$$Y = G + E + G \cdot E + \epsilon, \quad (22)$$

with a standard normally distributed error ϵ . The coefficients γ_i were either “small” ($\gamma_i = 0.5$) or “big” ($\gamma_i = 1$), leading to different weights of the scores relative to the noise variables e_2 and $g_3, \dots, g_{k_G}$. This created eight scenarios ($k_G = 4, 100$ crossed with $\gamma_1 = 0.5, 1$ and $\gamma_2 = 0.5, 1$).

Tuning parameters for the same synthetic data sets were then chosen based on 10-fold CV or BIC, comparing mean values of tuning parameters and (false) selection rates. As the CV procedure was computationally demanding (100 replications taking up to 24h on a single cluster node), 500 replications of each of the eight scenarios were recorded. Table 3 contains the results.

The signal-to-noise ratio is large enough so that even in the case with $k_G = 100$ and small values of γ_1 and γ_2 , the variables in the simulated model are always selected. Spurious selection of e_2 is frequent. Also, other genetic variables obtain non-zero coefficients occasionally. Regardless of whether CV or BIC is preferred, very similar tuning parameters result. Recall that t_g and t_e are (upper bounds of) the sum of the absolute values of the scores, so that they are supposed to increase with the values of γ_i , $i = 1, 2$, though not necessarily linearly. Similarly, also the predicted values found by CV and BIC are highly aligned, with correlations typically very close to 1. In sum, both methods agree, but spurious selection of unrelated variables cannot be entirely avoided. Fortunately, that does not necessarily entail a substantial weight in the G- or E-score.

Extension to generalized linear models

Up to here, the focus was on linear models and least-squares with a lasso penalty. Least squares fitting is inappropriate when outcomes are binary, count data or non-negative data, but generalized linear models (GLMs; McCullagh & Nelder, 1989) are

alternatives for such data. The ALOA is not limited to the least-squares case, similar to what has been implemented by Jolicoeur-Martineau (2019) for AO. We briefly discuss the necessary changes.

In a GLM, the dependent variables $Y_1, \dots, Y_n$ follow some exponential family distribution, a class of probability distributions including among the Gaussian and other continuous distributions the Bernoulli and the binomial distribution for binary data, and the Poisson distribution for count data. Let $E(Y_i) = \mu_i$ for $i = 1, \dots, n$, let $\mathbf{x}_i$ denote the i th row-vector of covariates, and let h be a smooth monotonic “link function”⁵. Then

$$h(\mu_i) = \mathbf{x}_i \boldsymbol{\beta} \quad (23)$$

for a vector of regression coefficients $\boldsymbol{\beta}$; the expectations transformed with h are linear in the covariates. The link function will be chosen in combination with the distribution, and, for example, the logit function $h(x) = \ln x - \ln(1 - x)$ is often chosen in the binomial case, leading to logistic regression as a special case of a GLM. Next to $\boldsymbol{\beta}$, many distributions require the estimation of a scale parameter ϕ , e.g., the error variance σ^2 in a Gaussian model.

The main model from the LEGIT Equation 1 is replaced by a GLM, so that the full model equation becomes

$$h(\boldsymbol{\mu}) = \gamma_0 + \gamma_1 \mathbf{G} + \gamma_2 \mathbf{E} + \gamma_3 \mathbf{G} \odot \mathbf{E} \quad \text{where} \quad (24)$$

$$\mathbf{G} = \sum_{u=1}^{k_G} p_u \mathbf{g}_u \quad \text{and} \quad (25)$$

$$\mathbf{E} = \sum_{v=1}^{k_E} q_v \mathbf{e}_v, \quad (26)$$

where in a slight abuse of notation $h(\boldsymbol{\mu}) = (h(\mu_1), \dots, h(\mu_n))'$ for the expectations

$\mu_i = E(Y_i)$. Note that ϵ is missing from Equation 24 because the mean is given, but a scale

⁵It is more customary in the GLM literature to call the link function g , but we have reserved this symbol for the genetic variables.

parameter might have to be estimated. In the Appendix we discuss why this modification still guarantees convergence of ALOA.

Illustrative application to lifetime depression data

Based on data from Gonda et al. (2018), we apply our method to analyze the effect of several genetic factors and stress exposure, as well as their respective $G \times E$ interactions, on lifetime depression. The (binary) lifetime depression status was assessed through a self-reported background questionnaire by the original authors. Stress exposure was quantified in terms of recent negative life events (RLE) and measured with the *List of Threatening Experiences* questionnaire (Brugha et al., 1985) which concerns twelve threatening medical, personal or financial life events. Gonda et al. (2018) categorized the number of experienced negative life events as low (0-1), moderate (2) or high (3 or more).

Polymorphisms of seven genes are examined: 5HTTLPR, rs6265 (BDNF), rs6295 (HTR1A), rs6311 (HTR2A), rs7766029 (CB1), rs7958311 (P2RX7) and rs8836 (GALR2). Some of these genes have been linked to major depressive disorder (MDD) in previous studies: While there seems to be general agreement that *5HTTLPR* has no direct effect on MDD, the question whether or not it moderates the relationship between stress and MDD is fairly contested (Culverhouse et al., 2018; Karg et al., 2011; Risch et al., 2009). In their review, Gyekis et al. (2013) conclude that *rs6265 (BDNF)* is not directly linked with MDD. Verhagen et al. (2010) present similar findings, yet their gender stratified analysis reveals a significantly increased risk for male Met-allele (A) carriers. Hosang et al. (2014) find that the Met-allele also significantly moderates the relationship between stressful life events and MDD. Kishi et al. (2013) report a significant link between *rs6295 (HTR1A)* and MDD in their meta-analysis with an increased risk associated with the G allele. For *rs6311 (HTR2A)*, Zhao et al. (2014) find in their review study that the A allele (T) is significantly associated with MDD.

For the remaining three genes included in the analysis of Gonda et al. (2018) no

corresponding review articles could be identified, although the findings of Juhász et al. (2009) and Juhász et al. (2014) indicate that rs7766029 and rs8836 potentially moderate the relationship between recent negative life events and depression. In line with the original analysis, all gene variables were included.

Since lifetime depression is a binary outcome, we use the extension of our method to generalized linear models discussed in the previous section. The original gene variables from Gonda et al. (2018) are re-coded into sets of dummy variables. In the case where a main effect of the respective gene on lifetime depression is plausible based on the reviews discussed above, we use the lowest risk polymorphism as reference. We therefore set CC as the reference category for rs6295 (BDNF) and rs6311 (HTR2A), respectively. For the remainder of genes we use the most common polymorphism from the original data as reference.

Ten-fold CV informs optimal tuning parameter choice for t_g and t_e . We use a random initialization for β_g and β_e with five multistarts in the inner validation loop and ten multistarts for the final model based on the optimal values for t_g and t_e . For interpretation purposes, we also fit the final model to 1000 bootstrap samples of the original data to obtain 20%- and 80%-quantiles of the empirical distribution of all coefficients as well as the *selection rate* measuring how often the corresponding variable is selected, i.e., the stability of the variable selection. To aid the interpretation of the main model's coefficients $\gamma_0, \gamma_1, \gamma_2$ and γ_3 in the final bootstrap conditional on the selected tuning parameters we recorded the proportion out of 1,000 bootstrap samples for which the two-sided t-test for $\gamma_i = 0$, $i = 0, 1, 2, 3$ was significant at $\alpha = 5\%$. Note that this is a statistical description of results rather than inference and that no variable selection was performed in the main model.

Table 4 shows the results of our analysis. No main-effect-only variables have been identified during the pre-screening step. Recalling that the model specified in Equations 24, 25 and 26 is unique barring a change of sign, note that during the fitting process the signs of the gene- and environment score coefficients have been adjusted to always be

non-negative. The direction of the $G \times E$ interaction can, however, still be interpreted. However, the corresponding quantiles suggest that it is inconclusive in this example. With respect to the gene variables, our results indicate that for 5HTTLPR the ss polymorphism may be associated with a higher risk than the heterozygous sl polymorphism. For rs6265, the AG polymorphism appears to lower the risk of depression compared to the GG variant, while for rs8836 the CC polymorphism seems to be associated with a lower risk than the CG polymorphism. The remaining gene variable coefficients are either equal or close to 0. Contrary to the literature review above, this is also the case for rs6295 and rs6311.

Figure 1 visualizes the interaction pattern obtained from the point estimates in Table 4: As there are only three values for RLE, there are three corresponding regression lines. While risk increases with genescore for all three levels, the increase is most pronounced for high RLE. We obtained proportions of $p_{\gamma_0} = .988$ (intercept), $p_{\gamma_1} = .658$ (gene score), $p_{\gamma_2} = 0.971$ (environment score) and $p_{\gamma_3} = .817$ (interaction), showing that evidence for a substantial gene score main effect is weak, but strong for the environment score. This is in line with the analysis of Gonda et al. (2018). Based on our analysis, the pattern for the interaction is inconclusive (with 18.3% of bootstrap resamples without a significant interaction term in the main model and positive as well as negative estimates), but seems to warrant a further investigation in a larger sample. In contrast, Gonda et al. (2018) are more optimistic on the case for interactions based on Bayesian relevance analysis. We did not include age and gender as further covariates in our analysis.

Discussion

Relationship to the original LEGIT model

The LEGIT model (Jolicoeur-Martineau et al., 2019) is originally a linear model for structured interactions of gene, environment and other variables. More generally, the LEGIT model can be seen as a regression model with main effects and interaction terms, that are structured: Two (or even more) scores formed as weighted sums of variables out of

Figure 1. Interaction of genescore and recent negative life events.

predefined sets (e.g., gene- and environment variables). The scores take the place of individual variables in a regular linear model, and main effects and interactions of scores result, shifting the interpretation to the score level. The scores can be seen as estimates of main factors or latent variables. Since the score weights are estimated, the approach bears some similarity to measurement models in confirmatory factor analysis and structural equation models (SEMs), which would typically not include interactions, though latent moderated SEMs are of course an extension for that purpose (Maslowsky et al., 2015).

The current work extends the methodology of the LEGIT model in two important ways: (i) When interactions are central to the investigation, variables without any interactions might be seen as a false positive component of one of the scores in the LEGIT

models. The regularization and residualization techniques allow to partial out variables that only contribute main effects, improving subsequent variable selection. (ii) In an era of big data, the number of variables might easily exceed the number of observations. The LEGIT model cannot be fit without further modification. The proposed lasso penalty leads to the alternating lasso optimization algorithm (ALOA). ALOA avoids computationally heavy variable selection. By exploiting the dual form of the lasso, we have shown that the lasso constraints can be placed on the sum of the absolute values of the score weights, which is related to the total variance of the score. Similar to unpenalized alternating optimization with variable selection, the lasso discards some variables entirely from the score equations, as coefficients can have zero estimates. This built-in sparsity leads to scores with easy interpretation, as potentially only a few variables are included, when the tuning parameter is chosen accordingly.

Jolicoeur-Martineau et al. (2019) discuss for the original alternating optimization technique for the LEGIT model that “variable selection cannot be done with standard regularization techniques because the full model is never seen.” The current work circumvents this difficulty by not placing a lasso penalty on the whole model, but instead the score equations are penalized. This is a crucial step which also avoids penalties conditional on other parameters, which would threaten convergence of the algorithm as remarked by Jolicoeur-Martineau et al. (2019). We assume that our strategy extends to other regularization techniques, e.g., ridge (Hoerl & Kennard, 1970) and elastic net (Zou & Hastie, 2005) penalties, which has to be verified in further work. ALOA is much faster than AO in combination with NES-variable selection or in combination with other variable selection methods requiring fitting the LEGIT model many times, e.g. forward selection. The computational gain could be crucial when big (genomic) data is to be used.

Can alternating optimization schemes be avoided entirely for this class of regression models? We do not have a definitive answer to this question but remark that the geometry of the penalized likelihood surface might be ill-suited for other optimization approaches: In

the hierarchical LEGIT model, the parameters are non-linear and the whole problem is non-convex (only biconvex).

Limitations and cautious remarks

Regularization and Residualization (RR). The lasso-based prescreening model includes interaction terms of all variables. These might have large variance relative to the main effects. Standardization is a default way to handle this, but also penalty factors are strategies from the lasso literature (Friedman et al., 2010). Alternatively, the hierarchical lasso (Lim & Hastie, 2015) could help to screen for interaction effects.

We recommend to use RR, if a priori interactions are to be studied. Note that removing “main effects only variables” without calculating partial variables (residualization) is possible. There is some risk in this simpler strategy: Predictors not removed might share variance with the removed ones, leading to the inclusion of a predictor in the score with only small interactions. Also, the assumption of gene and/or environment scores might not be an option in purely exploratory research, where it is unclear whether interactions exist at all. We hence propose to avoid scores initially in the exploratory case and use lasso-based pre-screening instead to check whether any interactions are present at all.

ALOA in applications. Statistical inference is complicated by the use of the lasso penalty. Especially, p -values for lasso models are subject of current research. If the lasso penalty does not shrink a coefficient to zero, it cannot be inferred that the coefficient is significantly different from zero. Also, p -values for the main model coefficients (γ_0 and γ_3 in the ALOA algorithm above) are not valid, since they are conditional on the scores, which are estimates. We recommend to study the stability of the variable selection and the regression coefficients with the help of a bootstrap. For this, the whole ALOA procedure is ideally bootstrapped, including the choice of the tuning parameters. However, as the sample application shows, this has to be done with some care, as signs of regression

coefficients can change. In addition, it will be desirable in many applications to extend the main model equation to include additional covariates, a modification already discussed by Jolicoeur-Martineau et al. (2019).

References

- Belsky, D. W., & Israel, S. (2014). Integrating genetics and social science: Genetic risk scores. *Biodemography and social biology*, 60, 137–155.
- Brugha, T., Bebbington, P., Tennant, C., & Hurry, J. (1985). The list of threatening experiences: A subset of 12 life event categories with considerable long-term contextual threat. *Psychological Medicine*, 15, 189–194.
- Culverhouse, R. C., Saccone, N. L., C., H. A., Et al. (2018). Collaborative meta-analysis finds no evidence of a strong interaction between stress and 5-HTTLPR genotype contributing to the development of depression. *Molecular Psychiatry*, 23, 133–142.
- Fadista, J., Manning, A. K., Florez, J. C., & Groop, L. (2016). The (in) famous gwas p-value threshold revisited and updated for low-frequency variants. *European Journal of Human Genetics*, 24, 1202–1205.
- Fletcher, J. M., & Conley, D. (2013). The challenge of causal inference in gene–environment interaction research: Leveraging research designs from the social sciences. *American journal of public health*, 103, S42–S45.
- Fraley, R. C., Roisman, G. I., Booth-LaForce, C., Owen, M. T., & Holland, A. S. (2013). Interpersonal and genetic origins of adult attachment styles: A longitudinal study from infancy to early adulthood. *Journal of personality and social psychology*, 104, 817.
- Friedman, J., Hastie, T., & Tibshirani, R. (2010). Regularization paths for generalized linear models via coordinate descent. *Journal of Statistical Software*, 33, 1–22.
<http://www.jstatsoft.org/v33/i01/>
- Goldberg, D. E. (1989). *Genetic algorithms in search, optimization, and machine learning*. Reading, MA, Addison-Wesley.
- Gonda, X., Hullám, G., Antal, P., Eszlari, N., Petschner, P., Hökfelt, T. G. M., Anderson, I. M., Deakin, J. F. W., Juhász, G., & Bagdy, G. (2018). Significance of risk polymorphisms for depression depends on stress exposure, In *Scientific reports*.

- Gyekis, J. P., Yu, W., Dong, S., Wang, H., Qian, J., Kota, P., & Yang, J. (2013). No association of genetic variants in BDNF with major depression: A meta- and gene-based analysis. *American Journal of Medical Genetics Part B: Neuropsychiatric Genetics*, 162, 61–70.
- Haefl, G. J., Getchell, M., Koposov, R. A., Yrigollen, C. M., De Young, C. G., Klinteberg, B. a., Orelan, L., Ruchkin, V. V., & Grigorenko, E. L. (2008). Association between polymorphisms in the dopamine transporter gene and depression: Evidence for a gene-environment interaction in a sample of juvenile detainees. *Psychological science*, 19, 62–69.
- Hastie, T., Tibshirani, R., & Friedman, J. (2009). *The elements of statistical learning: Data mining, inference, and prediction* (2nd) [12th corrected printing 2017]. New York, NY, Springer.
- Hastie, T., Tibshirani, R., & Wainwright, M. (2015). *Statistical learning with sparsity: The lasso and generalizations*. Boca Raton, FL, Chapman & Hall/CRC.
- Hoerl, A. E., & Kennard, R. W. (1970). Ridge regression: Biased estimation for nonorthogonal problems. *Technometrics*, 12, 55–67.
- Holland, J. H. (1975). *Adaptation in natural and artificial systems*. Ann Arbor, MI, University of Michigan Press.
- Hosang, G. M., Shiles, C., Tansey, K. E., McGuffin, P., & Uher, R. (2014). Interaction between stress and the BDNF Val66Met polymorphism in depression: A systematic review and meta-analysis. *BMC Medicine*, 12.
- Hurvich, C. M., & Tsai, C.-L. (1989). Regression and time series model selection in small samples. *Biometrika*, 76, 297–307.
- Jolicoeur-Martineau, A. (2019). *LEGIT: Latent Environmental & Genetic InTeraction (LEGIT) model* [R package version 1.2.3]. R package version 1.2.3.
<https://CRAN.R-project.org/package=LEGIT>

- Jolicoeur-Martineau, A., Wazana, A., Szekely, E., Steiner, M., Fleming, A. S., Kennedy, J. L., Meaney, M. J., Greenwood, C. M., & the MAVAN team. (2019). Alternating optimization for GxE modelling with weighted genetic and environmental scores: Examples from the MAVAN study. *Psychological Methods*, *24*, 196–216.
- Juhász, G., Chase, D. S., Pegg, E., Downey, D., Tóth, Z., Stones, K., Platt, H., Mekli, K., Payton, A., Elliott, R., Anderson, I. M., & Deakin, J. F. W. (2009). CNR1 gene is associated with high neuroticism and low agreeableness and interacts with recent negative life events to predict current depressive symptoms. *Neuropsychopharmacology*, *34*, 2019–2027.
- Juhász, G., Hullám, G., Eszlari, N., Gonda, X., Antal, P., Anderson, I. M., Hökfelt, T. G. M., Deakin, J. F. W., & Bagdy, G. (2014). Brain galanin system genes interact with life stresses in depression-related phenotypes. *PNAS*, *111*, E1666–E1673.
- Karg, K., Burmeister, M., Shedden, K., & Sen, S. (2011). The serotonin transporter promoter variant (5-HTTLPR), stress, and depression meta-analysis revisited: Evidence of genetic moderation. *Archives of General Psychiatry*, *68*, 444–454.
- Kishi, T., Yoshimura, R., Fukuo, Y., Et al. (2013). The serotonin 1A receptor gene confer susceptibility to mood disorders: Results from an extended meta-analysis of patients with major depression and bipolar disorder. *European Archives of Psychiatry and Clinical Neuroscience*, *263*, 105–118.
- Lim, M., & Hastie, T. (2015). Learning interactions via hierarchical group-lasso regularization. *Journal of Computational and Graphical Statistics*, *24*, 627–654.
- Manuck, S. B., & McCaffery, J. M. (2014). Gene-environment interaction. *Annual review of psychology*, *65*, 41–70.
- Mardis, E. R. (2011). A decade’s perspective on DNA sequencing technology. *Nature*, *470*, 198–203.

- Maslowsky, J., Jager, J., & Hemken, D. (2015). Estimating and interpreting latent variable interactions: A tutorial for applying the latent moderated structural equations method. *International journal of behavioral development*, 39, 87–96.
- McCullagh, P., & Nelder, J. A. (1989). *Generalized linear models*. Boca Raton, FL, Chapman & Hall/CRC.
- Nelder, J., & McCullagh, P. (1989). *Generalized linear models* (2nd). Chapman; Hall/CRC Press.
- R Core Team. (2019). *R: A language and environment for statistical computing*. R Foundation for Statistical Computing. Vienna, Austria. <https://www.R-project.org/>
- Risch, N., Herrell, R., Lehner, T., Liang, K.-Y., Eaves, L., Hoh, J., Griem, A., Kovacs, M., Ott, J., & Merikangas, K. R. (2009). Interaction Between the Serotonin Transporter Gene (5-HTTLPR), Stressful Life Events, and Risk of Depression: A Meta-analysis. *Journal of the American Medical Association*, 301, 2462–2471.
- Schelldorfer, J., Bühlmann, P., & van de Geer, S. (2011). Estimation for high-dimensional linear mixed-effects models using ℓ_1 -penalization. *Scandinavian Journal of Statistics*, 38, 197–214.
- Schwarz, G. (1978). Estimating the dimension of a model. *The Annals of Statistics*, 6, 461–464.
- Sevcikova, H., & Rossini, T. (2015). *Rlecuyer: R interface to rng with multiple streams* [R package version 0.3-4]. R package version 0.3-4. <https://CRAN.R-project.org/package=rlecuyer>
- Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society. Series B (Methodological)*, 267–288.
- Verhagen, M., van der Meij, A., van Deurzen, P., Janzing, J., Arias-Vásquez, A., Buitelaar, J., & Franke, B. (2010). Meta-analysis of the BDNF Val66Met polymorphism in major depressive disorder: Effects of gender and ethnicity. *Molecular Psychiatry*, 15, 260–271.

- Waldmann, P., Mészáros, G., Gredler, B., Fuerst, C., & Sölkner, J. (2013). Evaluation of the lasso and the elastic net in genome-wide association studies. *Frontiers in genetics*, 4, 270.
- Wood, S. N. (2017). *Generalized additive models: An introduction with R* (2nd ed.). Boca Raton, FL, Taylor & Francis Group/CRC.
- Zhao, X., Sun, L., Sun, Y.-H., Ren, C., Chen, J., Wu, Z.-Q., Jiang, Y.-H., & Lv, X.-L. (2014). Association of HTR2A T102C and A-1438G polymorphisms with susceptibility to major depressive disorder: A meta-analysis. *Neurological Sciences*, 35, 1857–1866.
- Zou, H., & Hastie, T. (2005). Regularization and variable selection via the elastic net. *Journal of the royal statistical society: series B (statistical methodology)*, 67, 301–320.
- Zou, H., Hastie, T., & Tibshirani, R. (2007). On the “degrees of freedom” of the lasso. *The Annals of Statistics*, 35, 2173–2192.

Table 3
Comparison of 10-fold CV and BIC for tuning parameter selection (500 replications).

k_G	γ_1	γ_2	method	t_g	t_e	selection			false selection			correlation		
						g_1	g_2	e_1	e_2	g_{other}	$q_{2.5\%}$	mean	$q_{97.5\%}$	
4	small	small	BIC	0.52	0.39	1.00	1.00	1.00	0.42	0.40	0.99	1.00	1.00	
			CV	0.56	0.40	1.00	1.00	1.00	0.68	0.65				
100	small	small	BIC	1.35	0.37	1.00	1.00	1.00	0.52	0.44	0.94	0.99	1.00	
			CV	1.69	0.40	1.00	1.00	1.00	0.66	0.52				
4	big	small	BIC	0.92	0.32	1.00	1.00	1.00	0.55	0.73	0.99	1.00	1.00	
			CV	0.95	0.31	1.00	1.00	1.00	0.69	0.86				
100	big	small	BIC	1.11	0.30	1.00	1.00	1.00	0.52	0.22	0.99	1.00	1.00	
			CV	1.11	0.31	1.00	1.00	1.00	0.67	0.23				
4	small	big	BIC	0.44	0.66	1.00	1.00	1.00	0.75	0.57	0.99	1.00	1.00	
			CV	0.45	0.66	1.00	1.00	1.00	0.79	0.73				
100	small	big	BIC	1.22	0.62	1.00	1.00	1.00	0.60	0.67	0.97	1.00	1.00	
			CV	1.40	0.65	1.00	1.00	1.00	0.73	0.71				
4	big	big	BIC	0.74	0.52	1.00	1.00	1.00	0.58	0.65	1.00	1.00	1.00	
			CV	0.75	0.52	1.00	1.00	1.00	0.65	0.78				
100	big	big	BIC	1.16	0.48	1.00	1.00	1.00	0.51	0.47	0.99	1.00	1.00	
			CV	1.16	0.51	1.00	1.00	1.00	0.71	0.49				

Notes: (false) selection: proportion of non-zero coefficients in 500 replications.

Table 4

Mean values, quantiles and selection rates of main, gene and environmental coefficients (tuning parameters: $t_g = 2.1$, $t_e = 2.5$).

			mean	quantiles		selection	
				20%	80%	rate	
β_{main}		(Intercept)	-0.30	-0.35	-0.20	1.00*	
		genescore	0.11	0.06	0.21	1.00*	
		envscore	0.42	0.31	0.48	1.00*	
		genescore $\times$ envscore	0.14	-0.19	0.18	1.00*	
β_{g}	5HTTLPR (ref: sl)	ss	0.52	0.00	0.33	0.56	
		ll	-0.04	-0.21	0.00	0.41	
	rs6265 (BDNF; ref: GG)	AA	0.00	0.00	0.03	0.29	
		AG	-0.59	-0.86	0.00	0.70	
	rs6295 (HTR1A; ref: CC)	GG	0.00	0.00	0.00	0.14	
		GC	0.00	-0.42	0.00	0.34	
	rs6311 (HTR2A; ref: CC)	CT	0.00	0.00	0.00	0.16	
		TT	0.00	0.00	0.00	0.13	
	rs7766029 (CB1; ref: CT)	TT	0.00	0.00	0.00	0.20	
		CC	-0.07	0.00	0.00	0.31	
	rs7958311 (P2RX7; ref: GG)	AA	0.00	0.00	0.00	0.23	
		AG	0.00	0.00	0.00	0.16	
	rs8836 (GALR2; ref: CG)	CC	-0.73	-0.69	0.00	0.62	
		GG	0.00	0.00	0.00	0.22	
	β_{e}	RLE	moderate	0.64	0.38	0.85	0.98
			high	1.93	1.66	2.21	1.00

Notes: *: selection rate equal to 1 by design; envscore: environment score; ref: reference polymorphism for dummy coding; RLE: recent negative life events.

Appendix

Proof of convergence to a local minimum

The proof is very similar to arguments from Jolicoeur-Martineau et al. (2019) for the unconstrained case and at its core is the biconvexity (Hastie et al., 2015) of the criterion. For a vector of parameters $\boldsymbol{\theta}$, let $f(\boldsymbol{\theta})$ denote the least squares criterion for the LEGIT model from Equations 1, 2 and 3, i.e., the sum of squares of the residuals $SSE(\boldsymbol{\theta})$ one obtains by plugging $\boldsymbol{\theta}$ into the model equation. If we can show that in all iterations the constraints of the form $\sum_{u=1}^{k_G} w_u^g p_u < t_g$ and $\sum_{v=1}^{k_E} w_v^e q_v < t_e$ are respected and $f(\boldsymbol{\theta}^{(l+1)}) \leq f(\boldsymbol{\theta}^{(l)})$, then a convergent subsequence of $(\boldsymbol{\theta}^{(l)})_l$ exists by the Bolzano-Weierstrass theorem. As the constraints are respected in the initialization step 0., and in the updates 1. and 3. of the gene and environment score models, the constraints hold through all iterations. It suffices to check that in each ALOA step 1., 2., 3. and 4., the SSE decreases. In the $(l+1)$ th iteration in the update of the gene score model (step 1. of the loop), just the $(p_u)_u^{(l)}$ entries of $\boldsymbol{\theta}^{(l)}$ are changed (they are the only free parameters). This ensures that minimizing the SSE of step 1. subject to the constraints for $(p_u)_u$ can only decrease the SSE of the LEGIT model. Similarly the update in step 3. for the environment score decreases the global SSE. The least squares updates in step 2. and 4. of $\boldsymbol{\gamma}$ do neither change $(p_u)_u$ nor $(q_v)_v$, so a further decrease of the LEGIT model's SSE subject to the constraints is guaranteed.

In the GLM case, the argument is almost identical, but the SSE has to be replaced by the deviance. The deviance D of a GLM is given by

$$D = 2[\ell(\hat{\boldsymbol{\beta}}_{max}) - \ell(\hat{\boldsymbol{\beta}})]\phi, \quad (27)$$

where $\ell(\hat{\boldsymbol{\beta}}_{max})$ is the log-likelihood of the saturated model (with one parameter per data point) and $\ell(\hat{\boldsymbol{\beta}})$ is the log-likelihood of the current model (Nelder & McCullagh, 1989).

Because the saturated model is used, $\ell(\hat{\boldsymbol{\beta}}_{max})$ does not depend on the scores $\mathbf{G}$ and $\mathbf{E}$. As

stated above, the deviance seems to depend on ϕ . However, in the exponential family log-likelihoods, ϕ cancels out, so that D is in fact independent of ϕ (p.69 Wood, 2017).