

# Pathwise Least Angle Regression and a Significance Test for the Elastic Net

Muhammad Naveed Tabassum and Esa Ollila

Aalto University, Dept. of Signal Processing and Acoustics, P.O. Box 15400, FI-00076 Aalto, Finland

Email: {muhammad.tabassum, esa.ollila}@aalto.fi

**Abstract**—Least angle regression (LARS) by Efron et al. (2004) is a novel method for constructing the piece-wise linear path of Lasso solutions. For several years, it remained also as the de facto method for computing the Lasso solution before more sophisticated optimization algorithms preceded it. LARS method has recently again increased its popularity due to its ability to find the values of the penalty parameters, called *knots*, at which a new parameter enters the active set of non-zero coefficients. Significance test for the Lasso by Lockhart et al. (2014), for example, requires solving the knots via the LARS algorithm. Elastic net (EN), on the other hand, is a highly popular extension of Lasso that uses a linear combination of Lasso and ridge regression penalties. In this paper, we propose a new novel algorithm, called pathwise (PW-)LARS-EN, that is able to compute the EN knots over a grid of EN tuning parameter  $\alpha$  values. The developed PW-LARS-EN algorithm decreases the EN tuning parameter and exploits the previously found knot values and the original LARS algorithm. A covariance test statistic for the Lasso is then generalized to the EN for testing the significance of the predictors. Our simulation studies validate the fact that the test statistic has an asymptotic  $\text{Exp}(1)$  distribution.

## I. INTRODUCTION

In this paper, we consider a linear model, where the  $n$ -vector  $\mathbf{y} \in \mathbb{R}^n$  of observations is modeled as

$$\mathbf{y} = \mathbf{X}\boldsymbol{\beta} + \boldsymbol{\varepsilon}, \quad (1)$$

where  $\mathbf{X} \in \mathbb{R}^{n \times p}$  is a known predictor matrix,  $\boldsymbol{\beta} \in \mathbb{R}^p$  is the unknown vector of regression coefficients and  $\boldsymbol{\varepsilon} \in \mathbb{R}^n$  is the noise vector. For ease of exposition, we consider the centered linear model (i.e. we assume that the intercept is equal to zero).

Elastic net (EN) of [1] is a superset of the popular Lasso (Least absolute shrinkage and selection operator) [2] that is also termed as basis pursuit denoising (BPDN) in the literature. EN has recently been employed, for example, in a single snapshot DoA (direction-of-arrival) finding application in [3]. The EN estimator is defined as the solution of the following penalized residual sum of squares (RSS) optimization problem,

$$\hat{\boldsymbol{\beta}}(\lambda, \alpha) = \arg \min_{\boldsymbol{\beta} \in \mathbb{R}^p} \frac{1}{2} \|\mathbf{y} - \mathbf{X}\boldsymbol{\beta}\|_2^2 + \lambda P_\alpha(\boldsymbol{\beta}) \quad (2)$$

where  $\lambda \geq 0$  is the *EN penalty parameter* and the EN penalty term  $P_\alpha(\boldsymbol{\beta})$ , defined as

$$P_\alpha(\boldsymbol{\beta}) = \alpha \|\boldsymbol{\beta}\|_1 + \frac{(1-\alpha)}{2} \|\boldsymbol{\beta}\|_2^2,$$

is a convex combination of  $\ell_1$ -norm and  $\ell_2$ -norm penalties of the Lasso and ridge regression. The *EN tuning parameter*

$\alpha \in [0, 1]$ , which is chosen by the user, determines the mix between ridge regression and the Lasso. The Lasso is obtained for  $\alpha = 1$ , and will be denoted shortly as  $\hat{\boldsymbol{\beta}}(\lambda)$  instead of  $\hat{\boldsymbol{\beta}}(\lambda, 1)$ . The EN penalty has singularities at the vertexes like Lasso, which is a necessary property for sparse estimation. It also has strictly convex edges that then help in selecting variables as a group, which is a useful property when high correlations exists between predictors.

For a given fixed EN tuning parameter  $\alpha$ , the path of solutions  $\hat{\boldsymbol{\beta}}(\lambda, \alpha)$  indexed by  $\lambda$  are also piece-wise linear as in the case of Lasso ( $\alpha = 1$ ). It is then of interest to find the values of the penalty parameters,  $\lambda_0, \lambda_1, \dots, \lambda_K$ , called the *knots*, at which a new parameter enters the active set of non-zero coefficients. In the special case of Lasso, this can be achieved via the famous *least angle regression (LARS) algorithm* [4]. In this paper, we propose a pathwise (PW-)LARS-EN algorithm that computes the knots of EN over a grid of  $\alpha$  values. The developed PW-LARS-EN algorithm starts with  $\alpha = 1$  and computes the knots via the LARS algorithm. We then decrease the EN tuning parameter and exploit the knowledge of the previously found values of the knots and the original LARS algorithm (using an augmented form of the EN regularization problem) to compute the knots of EN at current  $\alpha$  in the grid. We decrease  $\alpha$  again, and repeat the procedure. In this way we can efficiently compute the solutions over a grid of  $\alpha$  values.

Recently, in [5], a covariance test statistic  $T_k$  was proposed for testing the significance of predictors in the context of the Lasso. The covariance test statistic requires solving the knots via the LARS algorithm. In this paper, we generalize the covariance test statistic for the EN, which we denote as  $T_k(\alpha)$ . This test statistic then requires computing the knots of EN solution for a given fixed  $\alpha \in [1, 0)$ . It was postulated in [5, Sec. 8] that covariance test statistic for the EN also follows standard exponential distribution,  $\text{Exp}(1)$ , when the null hypothesis that all signal variables are in the model holds true. The authors, in [5], proved this fact in the orthonormal case, i.e., when  $\mathbf{X}^T \mathbf{X} = \mathbf{I}_p$  and  $n = p$ . Thanks to the proposed PW-LARS-EN algorithm, we are able to compute the empirical distribution of  $T_k(\alpha)$  also in the general non-orthogonal case. Our simulation studies then confirm that the  $\text{Exp}(1)$  approximation for the covariance test statistic  $T_k(\alpha)$  is valid in the general case as well.

The paper is organized as follows. In Section II, we describe and derive the pathwise LARS-EN algorithm which computes

the knots of EN estimators over a grid of EN tuning parameter values. The covariance test statistic for the EN is described in Section III, and Section IV presents the simulation study. A real data example is given in Section V while Section VI concludes the paper.

*Notations:* Uppercase boldface letters are used for matrices and lowercase boldface letters for vectors. The vector space  $\mathbb{R}^n$  is equipped with the usual inner product,  $\langle \mathbf{a}, \mathbf{b} \rangle = \mathbf{a}^\top \mathbf{b}$ , where  $(\cdot)^\top$  denotes the transpose. This induces the conventional  $\ell_2$ -norm  $\|\mathbf{a}\|_2 = \sqrt{\mathbf{a}^\top \mathbf{a}}$ . The  $\ell_1$ -norm is defined as  $\|\mathbf{a}\|_1 = \sum_{i=1}^n |a_i|$ , where  $|a|$  denotes the absolute value of  $a \in \mathbb{R}$ . The  $\ell_0$ -norm of a vector is defined as  $\|\mathbf{a}\|_0 = \#\{i | a_i \neq 0\}$ , which is equal to the total number of non-zero elements in it. The support of  $\mathbf{a} \in \mathbb{R}^p$  is the index set of its non-zero elements, i.e.,  $\text{supp}(\mathbf{a}) = \{j \in \{1, \dots, p\} : a_j \neq 0\}$ .

## II. PATHWISE LARS-EN ALGORITHM

It is well-known that the paths of solutions of  $\hat{\beta}(\lambda)$  as a function of the regularization parameter  $\lambda$  are piece-wise linear in each coefficient. The values of  $\lambda$  at which a new predictor variable becomes active (and hence a change in the slope occurs) or leaves the active set are referred to as *knots* and are denoted as  $\lambda_0 > \lambda_1 > \dots > \lambda_K$ . The values of the knots are not fixed, but depend on the data  $(\mathbf{y}, \mathbf{X})$ . LARS-Lasso delivers the entire solution path as a function of the regularization parameter  $\lambda$ . For a detailed discussion of LARS-Lasso algorithm, we refer the reader to [4] or [6, Sect. 5.6]. In essence, the LARS-Lasso finds the knots. The knowledge of the knots then allows to solve the whole coefficient paths as the coefficients either increase or decrease in the interval  $(\lambda_k, \lambda_{k+1})$ .

It is also well-known that the EN estimator  $\hat{\beta}(\lambda, \alpha)$  has piece-wise linear solution paths for a given fixed  $\alpha \in [1, 0)$ . We now develop the PW-LARS-EN algorithm for finding the knots  $\lambda_0(\alpha), \lambda_1(\alpha), \dots, \lambda_K(\alpha)$  of the EN solution in (2) for fixed  $\alpha$ . Let  $\lambda_0(\alpha)$  denotes the smallest value of  $\lambda$  such that all estimated coefficients are zero, i.e.,  $\hat{\beta}(\lambda_0(\alpha), \alpha) = \mathbf{0}$ . This value is easily shown to be

$$\lambda_0(\alpha) = \frac{\max_j |\langle \mathbf{x}_j, \mathbf{y} \rangle|}{\alpha},$$

where maximum is over  $j \in \{1, \dots, p\}$ . Let  $\mathcal{A}(\lambda, \alpha)$  denotes the *active set* at  $\lambda < \lambda_0(\alpha)$ , i.e., the index set of predictors with nonzero coefficients values:

$$\mathcal{A}(\lambda, \alpha) = \text{supp}\{\hat{\beta}(\lambda, \alpha)\}$$

Thus  $[\hat{\beta}(\lambda, \alpha)]_j \neq 0$  for all  $j \in \mathcal{A}(\lambda, \alpha)$ . Each knot is a border value, after which a change in the set of active predictors occurs. The active set at the knot  $\lambda_k(\alpha)$  is denoted by  $\mathcal{A}_k(\alpha) = \mathcal{A}(\lambda_k, \alpha)$ . The active set  $\mathcal{A}_1(\alpha)$  thus contains a single index  $\mathcal{A}_1(\alpha) = \{j_1\}$ , where  $j_1$  is the predictor that becomes active first and is known to be  $j_1 = \arg \max_j |\langle \mathbf{x}_j, \mathbf{y} \rangle|/\alpha$ . By definition of the knots, one has that  $\mathcal{A}_k(\alpha) = \{\text{supp}\{\hat{\beta}(\lambda, \alpha)\}\}$  for all  $\lambda_{k-1}(\alpha) < \lambda \leq \lambda_k(\alpha)$  and  $\mathcal{A}_k(\alpha) \neq \mathcal{A}_{k+1}(\alpha)$  for all  $k = 0, 1, \dots, K-1$ . We denote by  $\mathbf{X}_{\mathcal{A}}$  the matrix  $\mathbf{X}$  restricted to the columns of the active set  $\mathcal{A}$ .

The PW-LARS-EN algorithm finds the knots  $\lambda_0(\alpha) > \lambda_1(\alpha) > \lambda_2(\alpha) > \dots > \lambda_K(\alpha)$  for given fixed  $\alpha$  in a grid

$$[\alpha] = \{\alpha_i \in [1, 0) : 1 = \alpha_1 < \alpha_2 < \dots < \alpha_m < 0\}.$$

The success of our proposed algorithmic scheme is due to the usage of a dense grid, and therefore we use  $[\alpha] = \{1, 0.99, 0.98, \dots\}$  as the default grid. This grid is used in the simulations and data analysis examples of the paper. It should be noted that  $\alpha_m$  is the smallest value of EN parameter  $\alpha$  for which the user wishes to obtain the EN solution paths. Often the regime  $[0.5, 0)$  may not be very interesting as those EN solutions are closer to the Ridge regression estimator ( $\alpha = 0$ ) than the Lasso estimator ( $\alpha = 1$ ). In our examples,  $\alpha_m$  is often chosen as  $\alpha_m = 0.5$ . The PW-LARS-EN algorithm first finds the knots for the Lasso (i.e. the case of  $\alpha = 1$ ) using the LARS-Lasso. We then decrease  $\alpha$  and consider the first value  $\alpha_1$  in the  $\alpha$ -grid, and exploit the augmented form in (3) for subsequent values of  $\alpha$ .

Let  $\text{LARS}(\mathbf{y}, \mathbf{X})$  denotes the LARS-Lasso algorithm that computes the knots  $\{\lambda_k\}_{k=0}^K$  for the Lasso along with corresponding solutions at the knots  $\{\hat{\beta}(\lambda_k)\}_{k=0}^K$ . Let

$$\{\lambda_k, \hat{\beta}(\lambda_k)\} = \text{LARS}(\mathbf{y}, \mathbf{X})|_k$$

denotes the case that the  $k^{\text{th}}$  knot and the respective solution is extracted from a sequence of the knot-solution pairs found by the LARS-Lasso algorithm. Next note that we can write the EN objective function in *augmented form* as follows:

$$\frac{1}{2} \|\mathbf{y} - \mathbf{X}\beta\|_2^2 + \lambda P_\alpha(\beta) = \frac{1}{2} \|\mathbf{y}_a - \mathbf{X}_a(\eta)\beta\|_2^2 + \gamma \|\beta\|_1 \quad (3)$$

where

$$\gamma = \lambda\alpha \quad \text{and} \quad \eta = \lambda(1 - \alpha), \quad (4)$$

are new parameterizations of the tuning and shrinkage parameter pair  $(\alpha, \lambda)$ , and

$$\mathbf{y}_a = \begin{pmatrix} \mathbf{y} \\ \mathbf{0} \end{pmatrix} \quad \text{and} \quad \mathbf{X}_a(\eta) = \begin{pmatrix} \mathbf{X} \\ \sqrt{\eta} \mathbf{I}_p \end{pmatrix}$$

are the *augmented* forms of the response vector  $\mathbf{y}$  and the predictor matrix  $\mathbf{X}$ , respectively. Note that (3) resembles the Lasso objective function with  $\mathbf{y}_a \in \mathbb{R}^{n+p}$  and that  $\mathbf{X}_a(\eta)$  is an  $(n+p) \times p$  matrix. Next consider the case of finding the knot  $\lambda_k(\alpha_i)$  for the  $i^{\text{th}}$  value,  $\alpha_i$ , in the grid  $[\alpha]$ . Equivalently stated, we wish to find the pair

$$\gamma_k = \lambda_k(\alpha_i) \cdot \alpha_i \quad \text{and} \quad \eta_k = \lambda_k(\alpha_i) \cdot (1 - \alpha_i)$$

using the alternative parametrization given in (4). Obviously, the knowledge of either  $\gamma_k$  or  $\eta_k$  for any fixed  $\alpha_i$  would allow us to solve for  $\lambda_k(\alpha_i)$ . In our algorithm, we exploit the following property: If we would know  $\eta_k$ , then (3) is simply a Lasso problem and due to equivalence (3),  $\gamma_k$  is the  $k^{\text{th}}$  knot of the Lasso problem for the augmented data  $(\mathbf{y}_a, \mathbf{X}_a(\eta))$ . Thus we can find  $\gamma_k$  and  $\hat{\beta}(\lambda_k, \alpha_i)$  via the LARS-Lasso algorithm as

$$\{\gamma_k, \hat{\beta}(\lambda_k, \alpha_i)\} = \text{LARS}(\mathbf{y}_a, \mathbf{X}_a(\eta_k))|_k.$$

The knot  $\lambda_k(\alpha_i)$  is then found simply by scaling the found knot  $\gamma_k$  as  $\lambda_k(\alpha_i) = \gamma_k/\alpha_i$ . Naturally, we do not know  $\eta_k$ , but fortunately a good approximation can be found using the  $k^{\text{th}}$  knot that was found previously (for  $\alpha_{i-1}$ ). That is, we use

$$\tilde{\eta}_k = \lambda_k(\alpha_{i-1}) \cdot (1 - \alpha_i) \approx \eta_k.$$

Algorithm 1 outlines the pseudo code for this procedure for any given value of  $k$ . It presumes that predictors are standardized ( $\|\mathbf{x}_j\|^2 = 1$ ).

---

**Algorithm 1:** PW-LARS-EN algorithm.

---

**input :**  $\mathbf{y} \in \mathbb{R}^n$ ,  $\mathbf{X} \in \mathbb{R}^{n \times p}$ ,  $[\alpha] \in \mathbb{R}^m$  (recall  $\alpha_1 = 1$ ).

**output:**  $\{\lambda_k(\alpha_i), \hat{\beta}(\lambda_k, \alpha_i)\}_{k=1, i=1}^{K, m}$ .

- 1  $\{\lambda_k(\alpha_1), \hat{\beta}(\lambda_k, \alpha_1)\}_{k=1}^K = \text{LARS}(\mathbf{y}, \mathbf{X})$
  - 2 **for**  $i = 2$  **to**  $m$  **do**
  - 3     **for**  $k = 2$  **to**  $K$  **do**
  - 4          $\tilde{\eta}_k = \lambda_k(\alpha_{i-1}) \cdot (1 - \alpha_i)$
  - 5          $\{\gamma_k, \hat{\beta}(\lambda_k, \alpha_i)\} = \text{LARS}(\mathbf{y}_a, \mathbf{X}_a(\tilde{\eta}_k))|_k$
  - 6          $\lambda_k(\alpha_i) = \gamma_k/\alpha_i$
- 

Let us illustrate the Algorithm on an simulated data set. We calculate the EN solution paths for simulated data from  $\mathbf{y} \sim \mathcal{N}_n(\mathbf{X}\beta^*, \mathbf{I})$ , where  $n = 100$ ,  $\mathbf{X} \in \mathbb{R}^{100 \times 10}$  is an orthogonal predictor matrix,  $\beta^* \in \mathbb{R}^{10}$  contains  $k = 5$  non-zero regression coefficients. Fig. 1 depicts the EN coefficients paths along with the knots (i.e.  $\lambda_k(\alpha)$  values in which a new variable enters the active set) that were found via our PW-LARS-EN algorithm. We consider 3 different values of EN tuning parameter  $\alpha = 1$  (Lasso),  $\alpha = 0.9$  and  $\alpha = 0.5$ , which are used throughout the paper for reporting the results. It can be observed that the proposed PW-LARS-EN algorithm is able to correctly find the EN solution path for each  $\alpha$ -value, which is evident from the Fig. 1.




Fig. 1. Elastic net coefficient profiles for three different values of  $\alpha$  and the found knots via PW-LARS-EN algorithm on an simulated data ( $n = 100$ ,  $p = 10$ ), where the true coefficient vector  $\beta^*$  has 5 non-zeros. Note that the left panel corresponds to Lasso ( $\alpha = 1$ ).

### III. COVARIANCE TEST FOR THE ELASTIC NET

Recently, in [5], a covariance test statistic was proposed for testing the significance of predictors in the context of the Lasso. The covariance test statistic requires solving the knots via the LARS algorithm and it tests the significance of the predictor that has entered the Lasso solution path in the interval  $(\lambda_k, \lambda_{k+1})$ . Thus we wish to test if the predictor that has entered the active set  $\mathcal{A}_{k+1}$  is just noise or a statistically significant predictor. As earlier, let  $\lambda_k$  and  $\mathcal{A}_k$  denote the  $k$ -th knot and the active set for the interval  $(\lambda_{k-1}, \lambda_k)$ , respectively, for the Lasso. The *covariance test statistic* is defined as [5]

$$T_k = \frac{1}{\sigma^2} \left( \langle \mathbf{y}, \mathbf{X}\hat{\beta}(\lambda_{k+1}) \rangle - \langle \mathbf{y}, \mathbf{X}_{\mathcal{A}_k}\hat{\beta}_{\mathcal{A}_k}(\lambda_{k+1}) \rangle \right) \quad (5)$$

where  $\sigma^2$  is the noise variance. Above in (5), the notation  $\hat{\beta}_{\mathcal{A}_k}(\lambda_{k+1})$  refers to the Lasso estimate that is computed using the penalty parameter  $\lambda_{k+1}$  and using just the variables in  $\mathcal{A}_k$ . In other words,  $T_k$  measures how much of the covariance between  $\mathbf{y}$  and obtained fit can be attributed to the predictor which has just entered the model. Under the null hypothesis that all  $k$  signal variables are in  $\mathcal{A}_k$ , i.e.,  $H_0 : \mathcal{A}_k \supseteq \mathcal{A}^*$  where  $\mathcal{A}^* = \text{supp}(\beta^*)$  is the underlying true support set, one has that  $T_k \rightarrow_d \text{Exp}(1)$  as  $n, p \rightarrow \infty$ . This result is valid under the assumption that the noise terms are i.i.d. Gaussian,  $\varepsilon \sim \mathcal{N}_n(\mathbf{0}, \sigma^2 \mathbf{I})$  and under general conditions on the predictor matrix  $\mathbf{X}$ .

The *covariance test statistic* for the EN estimator that uses the EN tuning parameter  $\alpha$ -value is defined as

$$T_k(\alpha) = \frac{1 + \eta_{k+1}}{\sigma^2} \left( \langle \mathbf{y}, \mathbf{X}\hat{\beta}(\lambda_{k+1}, \alpha) \rangle - \langle \mathbf{y}, \mathbf{X}_{\mathcal{A}_k}\hat{\beta}_{\mathcal{A}_k}(\lambda_{k+1}, \alpha) \rangle \right)$$

where  $\lambda_{k+1} \equiv \lambda_{k+1}(\alpha)$  denotes the knot of EN estimator for fixed  $\alpha$  and  $\eta_{k+1} = \lambda_{k+1} \cdot (1 - \alpha)$ . Similarly,  $\hat{\beta}_{\mathcal{A}_k}(\lambda_{k+1}, \alpha)$  refers to an EN estimate for fixed  $\alpha$ , which is computed using the  $(k+1)^{\text{th}}$  knot as the penalty parameter and just the variables in  $\mathcal{A}_k \equiv \mathcal{A}_k(\alpha)$ .

In the simulations that follow, we wish to test validity of the conjecture that EN covariance test statistic  $T_k(\alpha)$  for any fixed  $\alpha$  also converges asymptotically to standard exponential

random variable, i.e.,  $T_k(\alpha) \rightarrow_d \text{Exp}(1)$  as  $n, p \rightarrow \infty$ . Indeed this was postulated in [5, Sec. 8] where the authors only showed that  $\text{Exp}(1)$ -distribution holds true for  $T_k(\alpha)$  when the predictor matrix is orthonormal ( $\mathbf{X}^\top \mathbf{X} = \mathbf{I}_p$ ). The authors stated in [5, Sec. 8.1] that “one is tempted to use this approximation beyond the orthogonal setting as well”. Indeed our simulation results reported below confirm that the  $\text{Exp}(1)$  approximation for  $T_k(\alpha)$  is valid in the general case as well.

#### IV. SIMULATION STUDY

We generate an  $n \times p$  predictor matrix  $\mathbf{X} = (\mathbf{x}_1 \ \mathbf{x}_2 \ \dots \ \mathbf{x}_n)^\top$ , where vectors  $\mathbf{x}_i \in \mathbb{R}^p$ ,  $i = 1, 2, \dots, n$ , are i.i.d. random vectors distributed as  $\mathbf{x}_i \sim \mathcal{N}_p(\mathbf{0}, \Sigma(\sigma^2, \rho, \text{ST}))$ , where  $\Sigma(\sigma^2, \rho, \text{ST}) \in \mathcal{S}_{++}^p$  is a  $p \times p$  positive definite covariance matrix parametrized by marginal variance  $\sigma^2 = \text{var}(x_i) > 0$  and  $\rho \in (-1, 1)$  which determines the correlation coefficient  $\text{corr}(x_i, x_j)$  for all  $i, j \in \{1, \dots, p\}$ . The parameter  $\text{ST}$  determines the structure which can be a compound symmetry (CS) covariance structure,

$$\Sigma(\sigma^2, \rho, \text{CS}) = \sigma^2 \{ (1 - \rho)\mathbf{I} + \rho \mathbf{1}\mathbf{1}^\top \}$$

(where  $\mathbf{1}$  a  $p$ -vector of 1s) or a first-order autoregressive (AR1) covariance structure in which case the  $(i, j)^{\text{th}}$  element is

$$[\Sigma(\sigma^2, \rho, \text{AR1})]_{ij} = \sigma^2 \rho^{|i-j|}$$

for  $i, j \in \{1, \dots, p\}$ .

##### A. Global Null Hypothesis

Let  $\mathbf{y} \sim \mathcal{N}_n(\mathbf{X}\beta^*, \mathbf{I})$ , where  $\beta^* = \mathbf{0}$ , i.e., all coefficients are zero and the covariance structures for  $\mathbf{x}_i$  confirms with  $\Sigma_1 = \Sigma(1, 0, \text{I}) = \mathbf{I}_p$ ,  $\Sigma_2 = \Sigma(1, 0.25, \text{CS})$  and  $\Sigma_3 = \Sigma(1, 0.25, \text{AR1})$ . For all three cases, we test the global null hypothesis of  $\beta^* = \mathbf{0}$  (i.e.  $\|\beta^*\|_0 = k = 0$ ) via covariance test  $T_0(\alpha)$ , i.e., we test the significance of the predictor that first enters the active set. We generated  $n = 100$  observations and  $p$  varied from  $p = 10$  to  $p = 50$ . We compute the statistic  $T_0(\alpha)$  for 1000 data sets simulated from the model above and report the empirical mean, variance, and 95% quantile. If  $\text{Exp}(1)$  is a good approximation for the empirical distribution of  $T_0(\alpha)$ , we expect these figures to be close to mean, variance and 95% quantile of the  $\text{Exp}(1)$ -distribution. The values reported in the Table I confirm that the empirical distribution of  $T_0(\alpha)$  can be well approximated by  $\text{Exp}(1)$  distribution in all the cases of predictor matrix and EN tuning parameter considered ( $\alpha \in \{1.0, 0.9, 0.5\}$ ).

##### B. General Null Hypothesis

Let the observations vector  $\mathbf{y} \sim \mathcal{N}_n(\mathbf{X}\beta^*, \mathbf{I})$ , the covariance structure for  $\mathbf{x}_i \in \mathbb{R}^p$  conforms with  $\Sigma_1 = \mathbf{I}_p$  and the number of true non-zero coefficients is varying,  $\|\beta^*\|_0 = k = \{1, 2\}$ . The sample size is  $n = 100$  and number of predictors is  $p = 50$ . Note also that the non-zeros values of  $\beta^*$  are equal to 3, which are greater than  $\sigma\sqrt{2}\log p$  as suggested in [5, Theorem 1] for  $k > 1$  case. We compute the EN covariance test statistic  $T_k(\alpha)$  for testing the significance of the predictor

TABLE I  
MEAN, VARIANCE AND  $q_{.95}$  QUANTILE OF THE EMPIRICAL DISTRIBUTION OF  $T_0(\alpha)$  FOR  $\alpha = (1.0, 0.9, 0.5)$  OVER 1000 SIMULATED DATA SETS FROM THE GLOBAL NULL MODEL.

		$n = 100, p = 10$			$n = 100, p = 50$			
		Exp(1)	1.0	0.9	0.5	1.0	0.9	0.5
$\Sigma_1$	Mean	1	0.998	0.994	0.985	1.014	1.008	0.996
	Var	1	1.269	1.255	1.212	1.270	1.248	1.202
	$q_{.95}$	3.0	3.200	3.181	3.091	3.044	3.050	3.045
$\Sigma_2$	Mean	1	0.989	0.974	0.928	0.993	0.973	0.927
	Var	1	1.423	1.408	1.306	1.118	1.078	0.979
	$q_{.95}$	3.0	3.118	3.064	2.917	3.051	2.996	2.897
$\Sigma_3$	Mean	1	0.991	0.988	0.982	1.009	0.998	0.976
	Var	1	1.619	1.618	1.600	1.285	1.258	1.189
	$q_{.95}$	3.0	3.311	3.266	3.261	3.241	3.198	3.031

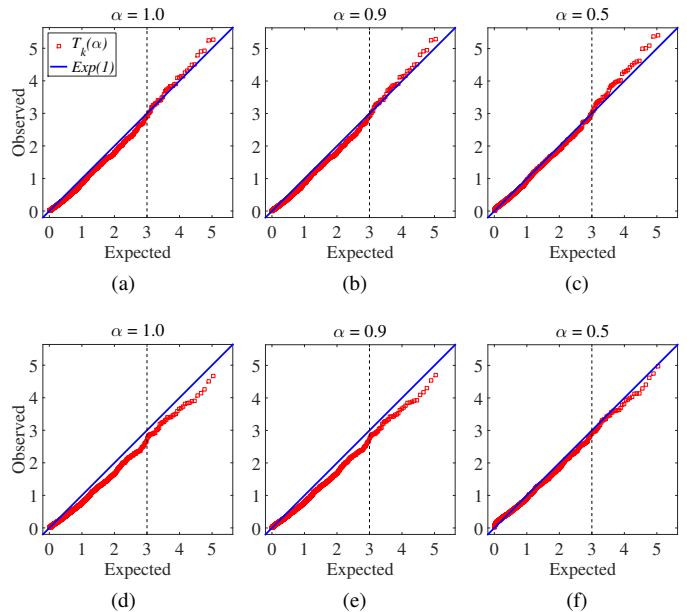


Fig. 2. Quantile-quantile plots, constructed over 1000 simulations, of the EN covariance test statistic  $T_k(\alpha)$  in the model, where  $k = \|\beta^*\|_0 \in \{1, 2\}$  and  $\alpha \in \{1, 0.9, 0.5\}$ . The first panel gives the results for  $k = 1$  and the second panel for  $k = 2$ . The sample size and dimension is  $n = 100, p = 50$ , respectively.

that enters to the active set  $\mathcal{A}_{k+1}$ . Fig. 2 depicts the quantile-quantile (QQ) plots of the EN covariance test statistic  $T_k(\alpha)$ , constructed over 1000 MC simulations, for testing the entry (significance) of the  $(k+1)^{\text{th}}$  predictor, i.e., the predictor that enter the active set in the interval  $(\lambda_k, \lambda_{k+1})$ . The first panel considers the case when  $k = 1$  and the second panel the case when there are two non-zero predictors, i.e.,  $k = 2$ . As can be seen, the QQ plots illustrate that the test statistic  $T_k(\alpha)$  is well approximated by  $\text{Exp}(1)$  distribution.

It has been shown in [5] for the case of Lasso and for orthonormal predictors ( $\mathbf{X}^\top \mathbf{X} = \mathbf{I}$ ) that

$$(T_k, \dots, T_{k+d-1}) \xrightarrow{d} (\text{Exp}(1), \dots, \text{Exp}(1/d)) \quad (6)$$

as  $n, p \rightarrow \infty$ . It is then of interest to investigate if the EN covariance test statistic  $T_k(\alpha)$  admits the property (6) as well but for  $\alpha \neq 1$ . We test this property for  $T_k(\alpha)$  in the case that  $k = 1$ . The simulation set-up is as before except that it

has now an orthonormal predictors matrix. The corresponding QQ plots for different  $\alpha$ -values are shown in Fig. 3, which clearly illustrate that EN covariance test statistic  $T_{k+1}(\alpha)$  is well approximated by  $\text{Exp}(1/2)$  distribution as the result (6) suggests.




Fig. 3. Quantile-quantile plots, constructed over 1000 simulations, of the EN covariance test statistic  $T_{k+1}(\alpha)$  in the model, where  $k = \|\beta^*\|_0 = 1$  and  $\alpha \in \{1, 0.9, 0.5\}$  and  $\mathbf{X}^\top \mathbf{X} = \mathbf{I}$ . The sample size and dimension is  $n = 100$ ,  $p = 50$ , respectively.

## V. REAL DATA EXAMPLE

A common case in practice is that  $\sigma^2$  is unknown and needs to be estimated from the data. If  $n > p$ , one can estimate  $\sigma^2$  via

$$\hat{\sigma}^2 = \|\mathbf{y} - \mathbf{X}\hat{\beta}_{\text{ls}}\|_2^2 / (n - p),$$

where  $\hat{\beta}_{\text{ls}}$  is the least-squares estimate for the full model. When one replaces  $\sigma^2$  with  $\hat{\sigma}^2$  in covariance test statistic  $T_k$  in (5) for the Lasso, its asymptotic distribution changes from  $\text{Exp}(1)$  to  $F_{2, n-p}$ ; see [5, Sec. 6] for details. Since we observed in the previous section that for known  $\sigma^2$ ,  $T_k(\alpha)$  shares the same asymptotic distribution as  $T_k$ , it is safe to presume here that  $F_{2, n-p}$  is an asymptotic distribution of  $T_k(\alpha)$  as well.

We consider the prostate cancer data set used e.g., in [7] as well as in [5, Sec. 6]. The training data consists of  $n = 67$  observations of male patients who had surgery for prostate cancer and  $p = 8$  predictors, which are clinical measures that are labelled as 1-'lcavol', 2-'lweight', 3-'age', 4-'lbph', 5-'svi', 6-'lcp', 7-'gleason', 8-'pgg45'. The response is the logarithm of PSA (prostate specific antigen) level. A more detailed description of the data set can be found in [8]. Note that rather strong correlations exist between the predictors: correlation between gleason and pgg45 is .757, but high correlations are found between other covariates as well. Correlation between svi and lcp is .673 and .675 between lcavol and lcp. The condition number of 243.30 which indicates a reasonably strong collinearity in the covariates and hence EN can be recommended for this data set.

We calculate the p-values of the EN covariance test statistic  $T_k(\alpha)$  using the  $F$ -distribution. For the data above, there are  $K = p = 8$  knots, i.e., at each active set  $\mathcal{A}_k(\alpha)$ ,  $k = 1, \dots, 8$ , a new variable enters the model. Note that depending on the value of  $\alpha$ , the predictors may enter the model in a different order. Table II lists the p-values along with entering predictor

TABLE II  
EN COVARIANCE TEST APPLIED TO THE PROSTATE CANCER DATA EXAMPLE. THE P-VALUES WITH ENTERING PREDICTOR NUMBER (IN BRACKETS) FOR ALL STEPS.  
(THE P-VALUES ARE ROUNDED TO 3 DECIMAL PLACES)

Step	$\alpha = 1.0$	$\alpha = 0.9$	$\alpha = 0.5$	$\alpha = 0.1$
1	0.000 (1)	0.000 (1)	0.000 (1)	0.000 (1)
2	0.052 (2)	0.464 (5)	0.003 (5)	0.000 (5)
3	0.174 (5)	0.001 (2)	0.053 (2)	0.020 (6)
4	0.930 (4)	0.470 (8)	0.077 (6)	0.000 (2)
5	0.353 (8)	0.047 (4)	0.005 (8)	0.000 (8)
6	0.650 (3)	0.646 (3)	0.011 (4)	0.000 (7)
7	0.051 (6)	0.055 (6)	0.020 (7)	0.000 (4)
8	0.978 (7)	0.978 (7)	0.978 (3)	0.167 (2)

number (in brackets) for all steps. The results are shown for different  $\alpha$ -values. Note that variables enter to the model in different order depending on the value of  $\alpha$ . This feature can probably be attributed to existing strong correlations among the predictors.

## VI. CONCLUSIONS

The EN coefficient path for fixed  $\alpha$  is piece-wise linear as a function of the penalty parameter  $\lambda$ . In this paper, we proposed a novel approach, pathwise (PW-)LARS-EN algorithm that computes the knots of EN over a grid of  $\alpha$  values. Thus the PW-LARS-EN allows to compute the whole EN solution path over a large range of  $\alpha$  values. Furthermore, we illustrated that the found EN knots can be used to construct a significance test for the EN as was done recently in the seminal paper by [5].

## ACKNOWLEDGMENT

The research was partially supported by the Academy of Finland grant no. 298118 which is gratefully acknowledged.

## REFERENCES

- [1] H. Zou and T. Hastie, "Regularization and variable selection via the elastic net," *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, vol. 67, no. 2, pp. 301–320, 2005.
- [2] R. Tibshirani, "Regression shrinkage and selection via the lasso," *Journal of the Royal Statistical Society. Series B (Methodological)*, pp. 267–288, 1996.
- [3] M. N. Tabassum and E. Ollila, "Single-snapshot doa estimation using adaptive elastic net in the complex domain," in *4th International Workshop on Compressed Sensing Theory and its Applications to Radar, Sonar and Remote Sensing (CoSeRa)*. IEEE, 2016, pp. 197–201.
- [4] B. Efron, T. Hastie, I. Johnstone, R. Tibshirani *et al.*, "Least angle regression (with discussion)," *The Annals of statistics*, vol. 32, no. 2, pp. 407–499, 2004.
- [5] R. Lockhart, J. Taylor, R. J. Tibshirani, and R. Tibshirani, "A significance test for the lasso (with discussion)," *Annals of statistics*, vol. 42, no. 2, p. 413, 2014.
- [6] T. Hastie, R. Tibshirani, and M. Wainwright, *Statistical learning with sparsity: the lasso and generalizations*. CRC Press, 2015.
- [7] J. Friedman, T. Hastie, and R. Tibshirani, *The elements of statistical learning*. Springer series in statistics Springer, Berlin, 2001, vol. 1.
- [8] T. A. Stamey, J. N. Kabalin, J. E. McNeal, I. M. Johnstone, F. Freiha, E. A. Redwine, and N. Yang, "Prostate specific antigen in the diagnosis and treatment of adenocarcinoma of the prostate. ii. radical prostatectomy treated patients." *The Journal of urology*, vol. 141, no. 5, pp. 1076–1083, 1989.