Analysis of Error Produced by Truncated SVD and Tikhonov Regularization Methods *

Irina F. Gorodnitsky and Bhaskar D. Rao
Dept. of Electrical and Computer Engineering
University of California, San Diego
La Jolla, CA 92093-0407

Abstract

The paper considers regularization of inverse solutions to linear problems. We derive the relationship between the error due to regularization and the energy distribution in the observation vector across the width of the singular value spectrum of the linear model used. One result of the analysis is that conditions exist under which no close approximation to the true solution can be found. It follows that a check of these conditions to determine the success of the regularization is important and we feel it should be included in the regularization process. The presentation includes a series of simulations to illustrate the analysis.

1 Introduction

Finding solutions to linear problems typically involves inversion of the forward model, which maps the unknown signal to the data. Common types of solutions include least squares (LS), minimum norm, and weighted minimum norm solutions. An acceptable solution must be stable, i.e. continuously dependent on the data, since the data inevitably contains noise. The inverse operator for discrete problems, which we consider in this paper, is always bounded, but it can be numerically ill-conditioned. Inversions involving ill-conditioned operators can produce large magnifications of small perturbations in the data or model. This often leads to unacceptable results with the size of the error in the solution sometimes exceeding the size of the true signal.

In these cases *regularization* is used to produce stable estimates. Regularization techniques have been extensively investigated and two techniques are currently universally used, Tikhonov regularization and Truncated Singular Value Decomposition (TSVD). The application of regularization requires selection of a *regularization parameter*, which is not trivial to identify. A regularization parameter is chosen to achieve a compromise between a close fit to the data and closeness to the true solution.

Regularization methods in effect replace the original forward operator with a well-conditioned approximation that is close to it. The small changes in the forward model are assumed to produce only small errors in a regularized estimates. Therefor the emphasis in study of regularization has been predominantly in the area of optimization of regularization parameters as a function of noise. Ad hoc methods for parameter selection are sometimes still used, based on the assumption that the the error in a solution is determined by the width of the singular value spectrum of the forward operator. Analytical methods for *optimal regularization parameter* determination have been derived using MSE minimization criteria, but all of the them require some knowledge of the statistical properties of noise, which is unlikely in real applications. In [1, 2] expressions for the optimal regularization parameter are given for the case when the noise is white. In [3] and [4] criteria for the optimal regularization parameter in a more general case are given. The criteria in the two papers can be shown to be equal, and they link the error in a regularized solution to the distribution of the energy in the data vector among the singular vectors of the forward operator. Sano [4] has proposed an alternative iterative scheme for determining optimal regularization parameters where no knowledge of the noise is required, but it did not work consistently in our trials.

Here we re-examine the stability of regularized approximations by considering the error in an estimate due to the regularization process itself, using the two popular regularization techniques. Our analysis provides the following results. We show that the error in a regularized estimate is a function of the energy distribution of the data vector across the singular subspaces of the forward operator, relative to the size of the corresponding singular values. From this result we formulate conditions under which successful regular-

*This work was supported by the NSF grant no. MIP-9220550

1058-6393/95 $4.00 © 1995 IEEE

ization can be achieved and identify conditions under which no solution approximation can be found. The word *approximation* is used here in the mathematical sense to mean an estimate that is sufficiently close to the true solution, i.e., the errors in it are comparable to the size of the noise in the observed data. We also use the term *well-conditioned subspaces* of the forward model matrix to mean those singular subspaces that form the closest well-conditioned approximation to the forward matrix. The ill-conditioned subspaces are the remaining subspaces of the forward matrix after extraction of the well-conditioned ones.

The implications of the analysis are as follows. The width of the singular value spectrum of the forward operator has little to do with the error in regularized solutions. Severely ill-conditioned forward operators can yield good regularized approximations, given favorable energy distribution in the data vector. At the same time, mildly ill-conditioned forward operators can lead to no close approximation, even when an optimal regularization parameter is used, for a certain group of data inputs. The results indicate that energy distribution of the data vector is a determining factor in the success of regularization and should be considered in any regularization process.

Previous work on the subject consists of analysis by Per Hansen [5] in which he develops bounds on regularization error based on an assumed model for the observation vector. From this he designs what he calls the discrete Pickard condition (DPC), which states that the energy in the data vector projected onto singular subspaces of the forward operator must, on average, decay faster than the singular values of the operator in order for a regularized solution to exist. Our analysis and simulations show that this is not the case. The results show that the energy distribution of the data vector among the singular vectors of the operator does not need to decrease with the size of the corresponding singular values, and can in fact increase, as long as the respective subspaces remain well-conditioned.

2 Background

Problem Statement

The discrete linear statistical model is an observation plus noise model which can be written as

$$\mathbf{A}\mathbf{x} = \mathbf{b} + \mathbf{n_O}, \qquad (1)$$

The matrix $\mathbf{A}$ is the true forward model which generates the noise free data signal $\mathbf{b}$. The vector $\mathbf{n_O}$ represents noise. The problem is to identify the vector of interest $\mathbf{x}$ from the observation $\mathbf{b} + \mathbf{n_O}$. The model $\mathbf{A}$ is assumed to be known. The system can be either overdetermined or underdetermined. To encompass both cases in our analysis, we assume, without bearing on the generality of the results, that $\mathbf{A}$ is full rank, rank$=m$, where m is the smallest dimension of $\mathbf{A}$. Thus in the overdetermined case we take $\mathbf{A}$ to be $n \times m$ and in the underdetermined case $\mathbf{A}$ is $m \times n$. Problems described by such linear models include signal estimation, extrapolation, interpolation, reconstruction, classification, and imaging.

Well known solutions to (1) are LS for the overdetermined case and the minimum norm solution in the underdetermined case. Both are described by the same equation

$$\mathbf{x_{ls}} = \mathbf{A}^+(\mathbf{b} + \mathbf{n_O}), \qquad (2)$$

where $\mathbf{A}^+$ denotes the generalized inverse. Alternative solutions optimizing other criteria also exist, and can be expressed using inverses of linearly weighted forward models. The results of our analysis on the effect of the inverse operation on the solutions, are independent of linear transformations on the forward operator. Hence, for the purposes of the paper, we adopt LS/minimum norm as the criteria for the true solutions without affecting the generality of the results.

The exact solution to (1) is then given by

$$\mathbf{x_0} = \mathbf{A}^+\mathbf{b} \;=\; \textstyle\sum_{i=1}^{m} \frac{\mathbf{u_i^H b}}{\sigma_i}\mathbf{v_i}, \qquad (3)$$

where $\mathbf{u_i}$ and $\mathbf{v_i}$ are the left and right singular vectors and $\sigma_i : \sigma_1 \geq \sigma_2 \geq \cdots \sigma_n$ are the singular values. We refer to $\mathbf{u_i}$ as singular subspaces of $\mathbf{A}$.

SVD of $\mathbf{A}$ forms a basis for the study of regularization. The condition number of $\mathbf{A}$ determines the sensitivity of the inverse operation to inaccuracies in the data, and is defined by the ratio σ_1/σ_n. High sensitivity to noise is caused by the presence of noise in the the singular subspaces of $\mathbf{A}$ associated with sufficiently small σ_i.

Regularization of solutions

The Tikhonov regularization method amounts to finding the unique solution $\mathbf{x}_\lambda$ to the least squares problem with a quadratic constraint:

$$min\{\|\mathbf{A}\mathbf{x} - \mathbf{b}\|_2^2 + \lambda^2\|bfx\|_2^2\}. \qquad (4)$$

The regularized solution x_λ is a function of the regularization parameter λ. It can be written in terms of SVD as:

$$\mathbf{x}_\lambda = \textstyle\sum_{i=1}^{m} \frac{\sigma_i^2}{\sigma_i^2 + \lambda_i^2} \frac{\mathbf{u_i^H b}}{\sigma_i}\mathbf{v_i}. \qquad (5)$$

The effect of λ_i is to dampen any terms in the sum with singular values smaller than λ_i. Hence λ_i can be assumed to satisfy $\lambda_i \geq \sigma_n$.

The TSVD method is the other standard regularization techniques in which a number of subspaces of $\mathbf{A}$ with the smallest singular values are truncated. The truncation parameter k, also known as the order of approximation, is the number of subspaces used to compute the estimate. The *kth order TSVD solution* is defined by

$$\mathbf{x_k} = \sum_{i=1}^{k} \frac{\mathbf{u_i^H b}}{\sigma_i} \mathbf{v_i}. \tag{6}$$

3 Analysis

We examine how the behavior of the right hand side in (1) affects the error in regularized solutions. For this purpose we examine the error in an estimate due to the regularization process itself. We define

$$\gamma_i = \frac{\mathbf{u_i^T b}}{\sigma_i} \tag{7}$$

as a measure of the energy of b in a singular subspace of A relative to the power of that subspace. We show that the shape of the function γ is the determining factor in regularization error. It is also useful to define

$$\gamma_l = \max_{1<i<m} \gamma_i. \tag{8}$$

The error due to regularization is given in the following two theorems.

Theorem 1 *Let* $\mathbf{x}$ *and* $\mathbf{x}_\lambda$ *be the exact and the Tikhonov regularized solutions given by (3) and (5) respectively. The following can be shown to hold*

$$\|\mathbf{x_0} - \mathbf{x}_\lambda\|_2 = \left(\sum_{i=1}^{m} \left(\frac{\lambda_i^2}{\sigma_i^2 + \lambda_i^2}\gamma_i\right)^2\right)^{1/2} = \begin{cases} \leq & \sqrt{m}\frac{\lambda_p^2}{\sigma_p^2+\lambda_p^2}\gamma_p \\ \geq & \frac{\lambda_p^2}{\sigma_p^2+\lambda_p^2}\gamma_p \end{cases}, \tag{9}$$

where $\frac{\lambda_p^2}{\sigma_p^2+\lambda_p^2}\gamma_p = \max_{1<i<m} \frac{\lambda_i^2}{\sigma_i^2+\lambda_i^2}\gamma_i$.
and

$$\frac{\|\mathbf{x_0} - \mathbf{x}_\lambda\|_2}{\|\mathbf{x_0}\|_2} = \begin{cases} \leq & \sqrt{m}\frac{\lambda_p^2}{\sigma_p^2+\lambda_p^2}\frac{\gamma_p}{\gamma_l} \\ \geq & \frac{\lambda_p^2}{\sigma_p^2+\lambda_p^2}\frac{\gamma_p}{\gamma_l} \end{cases}, \tag{10}$$

where $l \leq p, \quad \gamma_p \leq \gamma_l$.

Proof: Rewriting (3), the norm of the true solution $\mathbf{x_0}$ is

$$\|\mathbf{x_0}\|_2 = \|\mathbf{V}_m[\gamma_1, \cdots, \gamma_m]^T\|_2 = \begin{cases} \leq & \sqrt{m}(\max_{1<i<m} \gamma_i) \\ \geq & \max_{1<i<m} \gamma_i \end{cases}, \tag{11}$$

where $\mathbf{V_m}$ is the section of $\mathbf{V}$ consisting of m columns. In the overdetermined case $\mathbf{V_m} = \mathbf{V}$.
Using (5) and (11)

$$\|\mathbf{x_0} - \mathbf{x}_\lambda\|_2 = \|V_m \left[\frac{\lambda_1^2}{\sigma_1^2 + \lambda_1^2}\gamma_1, \cdots, \frac{\lambda_n^2}{\sigma_n^2 + \lambda_n^2}\gamma_m\right]^T\|_2 = \begin{cases} \leq & \sqrt{m}\frac{\lambda_p^2}{\sigma_p^2+\lambda_p^2}\gamma_p \\ \geq & \frac{\lambda_p^2}{\sigma_p^2+\lambda_p^2}\gamma_p \end{cases}$$

which gives (9). Combining (9) and (11) gives the desired result (10).

Theorem 2 *Let* $\mathbf{x}$ *and* $\mathbf{x_k}$ *be the exact and the TSVD regularized solutions given by (3) and (6) respectively. Define* $\gamma_p = \max_{k+1<l<m} \gamma_i$ *and* $v = m - (k+1)$ *to be the number of truncated subspaces. Then the following hold*

$$\|\mathbf{x_0} - \mathbf{x_k}\|_2 = \left(\sum_{i=k+1}^{m} \gamma_j^2\right)^{1/2} = \begin{cases} \leq & \sqrt{v}\gamma_p \\ \geq & \gamma_p \end{cases} \tag{12}$$

and

$$\frac{\|\mathbf{x_0} - \mathbf{x_k}\|_2}{\|\mathbf{x}\|_2} \begin{matrix} \leq \begin{cases} \frac{\sqrt{v}\gamma_l}{\gamma_p} & , \max_{1<i<m} \gamma_i = \gamma_l \neq \gamma_p \\ \sqrt{v} & , \max_{1<i<m} \gamma_i = \gamma_p \end{cases} \\ \geq \begin{cases} \frac{1}{\sqrt{m}}\frac{\gamma_l}{\gamma_p} & , \max_{1<i<m} \gamma_i = \gamma_l \neq \gamma_l \\ \frac{1}{\sqrt{m}} & , \max_{1<i<m} \gamma_i = \gamma_p \end{cases} \end{matrix} \tag{13}$$

Proof: Using (6) and (11) we have

$$\|\mathbf{x_0} - \mathbf{x_k}\|_2 = \|\mathbf{V}_m[0, \cdots, 0, \gamma_{k+1}, \cdots, \gamma_m]^T\|_2 \begin{cases} \leq \sqrt{v}\gamma_p \\ \geq \gamma_p \end{cases} \tag{14}$$

Combining this with (11) gives (13).

Note that the bounds given in the theorems are tight, since cases closely satisfying these bounds can be easily designed.

The theorems show how the relative sizes of γ_i affect the regularization error. In the discussion we assume that a reasonable regularization parameter is being used, e.g. $\sigma_m < \lambda_i < \sigma_1$. In Tikhonov regularization the error is shown to depend on two factors. The term $\frac{\lambda_i^2}{\sigma_i^2+\lambda_i^2}$ is a decreasing function of σ_i. When σ_i is large, $\lambda_i < \sigma_i$, but λ_i/σ_i increases with decreasing

σ_i, so that $\lambda_i >> \sigma_i$ for σ_i corresponding to the ill-conditioned subspaces. The value of $\frac{\lambda_i^2}{\sigma_i^2+\lambda_i^2}$ sharply increases for λ_i and σ_i in the transition interval from well-conditioned to ill-conditioned subspaces. Thus regularization produces a disproportionately large error for a component of the data **b** in the ill-conditioned subspaces, compared to an equal sized component in the well-conditioned subspaces. The more energy of b that shifts to the ill-conditioned subspaces, the larger the error due to regularization. This is independent of how close to the optimum the regularization parameter may be.

The equality $\gamma_p = \gamma_l$ also maximizes the error and the term $\frac{\lambda_l^2}{\sigma_l^2+\lambda_l^2}$ determines whether the equality occurs. When γ_l corresponds to an ill-conditioned subspace, the likelihood of the equality $\gamma_p = \gamma_l$ is increased sharply, since $\frac{\lambda_l^2}{\sigma_l^2+\lambda_l^2} \approx 1$ there. In this case the upper bound on the error is $\sqrt{(m)}$, which is independent of the signal and model parameters, and so the regularization error is likely to be sharply higher.

The results for the TSVD method are similar. The component of b that is in the ill-conditioned subspaces is truncated in the process of regularization, which contributes directly to the regularization error in the form of γ_j, $k+1 < j < m$. The result is independent of the choice of the truncation parameter k.

Whether the error approaches the lower or upper bound depends on the relative sizes of the γ_i. The error is closer to the upper bound when all γ_i are of similar size, while the lower bound is approached when the energy of b is concentrated in one or very few singular subspaces. However, when all γ_i are equal, the error due to TSVD is $\frac{\sqrt{v}}{\sqrt{m}}$, which indicates relatively good performance. When the γ_i are not exactly equal, the error worsens until $\gamma_l = \max_{1<i<m} \gamma_i$ corresponding to a well-conditioned subspace becomes sufficiently larger than the rest, i.e. the ratio $\frac{\gamma_p}{\gamma_l} < 1/\sqrt{m}$ is reached. Also, a wider gap between the retained γ_i, $i \leq k$ and the truncated γ_j, $j > k$ improves regularization accuracy, with the error diminishing as the gap widens.

We now summarize the results. When the γ_j that correspond to the ill-conditioned subspaces are small relative to the γ_i that correspond to the well-conditioned subspaces, the regularization error is guaranteed to be small, bounded by the parameters of the linear system (1), $\mathbf{u}_1^{\mathbf{H}}\mathbf{b}$ and $1/\sigma_i$, with i such that σ_i are close to $\sigma_1 = \|\mathbf{A}\|_2$. This means that the regularized approximation to x_0 can be found. Otherwise, the regularization error may not be reasonable and the regularized approximation to x_0 is not guaranteed to exist. Noise, which is inevitably present and in general can be assumed to be distributed among all the subspaces of A, induces large errors when no regularization is used in a solution. Yet regularization also produces large errors when the γ_j are sufficiently large. Hence no acceptable estimate can be produced. The chance that regularization will fail increases as the γ_j increase relative to the γ_i.

The analysis shows that the energy distribution of **b** among the singular vectors does not need to decrease with the size of σ_i to obtain a good approximation. In fact, a relatively much larger γ_i, $i \neq 1$, reduces the regularization error, as long as the respective subspace is not in the ill-conditioned part of **A**.

4 Numerical examples

We present three examples to illustrate our results. The examples use varying **b** vectors and an arbitrarily chosen 6×8 **A** matrix whose singular values are $s_i =$ [1 .2573 .1838 .0912 .009 .0009].

The first four singular subspaces of the **A** matrix are well-conditioned while the last two are ill-conditioned. The input data **b** for the three cases are as follows. In the first two examples **b** is chosen to satisfy the favorable conditions for the regularization according to Theorems 1 and 2. The vector **b**, in the first case, is such that γ_i are an increasing function of σ_i. In the second case, γ_i have the highest values at the singular values 3 and 4, which are still well-conditioned. In the third case γ_i are a decreasing function of σ_i, which represents a non-favorable condition for regularization. Noise in the data is generated by adding 10% random gaussian noise to each input.

The noise free solution, the solution with no regularization, and the regularized solutions using Tikhonov regularization and TSVD truncation for the three examples are shown in Figs. 1-3. The L-curve criteria [6] was used to determine the optimal parameters in both regularization procedures. The corresponding γ functions for the three cases are also plotted in the small windows within each figure and are listed in the legends. The projections of the energy of **b** into the singular subspaces, i.e. the **U**′**b** values, are also given.

The first two examples show successfully regularized solutions. In the third example the regularization fails, as predicted by the theory. Note that the energy in **b** is concentrated in the principle subspaces in all three cases, and by itself is not a valid indicator of how successful regularization may be.

References

[1] Y.S. Shim and Z.H. Cho. "SVD pseudoinversion image

reconstruction". *IEEE Trans. Acoust., Speech and Sig. Process.*, 29:904–909, 1981.

[2] B. Sullivan and B. Liu. "On the use of singular value decomposition and decimation in discrete-time band-limited signal extrapolation". *IEEE Trans. Acoust., Speech and Sig. Process.*, 32:1201–1212, 4981.

[3] A. Thorpe and L. Scharf. "Reduced rank methods for solving least squares problems". *Proc. 23rd Asilomar Conf. Signals, Systems and Computers*, pages 609–613, Nov. 1989.

[4] A. Sano. "Optimally regularized inverse of singular value decomposition and application to signal extrapolation". *Signal Processing*, 30:163–176, 1993.

[5] P.C. Hansen. "Truncated singular value decomposition solutions to discrete ill-posed problems with ill-determined numberical rank". *SIAM J. Sci. Statist. Computing*, 11(3):503–518, 1990.

[6] P.C. Hansen. "Analysis of discrete ill-posed problems by means of L-curve". *SIAM Review*, 34:561–580, 1992.

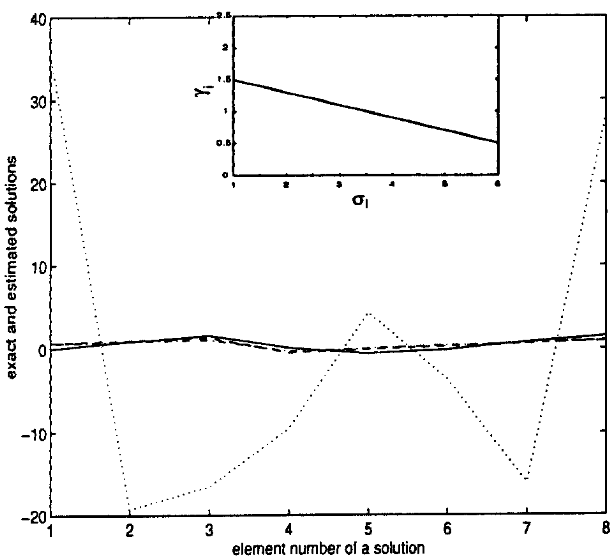

Figure 1: Noise free solution (solid line), solution with no regularization (dotted line), solution using Tikhonov regularization (dashdot line), and TSVD solution (dashed line) for the case when γ is a increasing function of σ_i. The input vector $\mathbf{b}$ satisfies $\mathbf{U}'\mathbf{b} = [1.5 \quad 0.3344 \quad 0.2022 \quad 0.0821 \quad 0.0063 \quad 0.0005]$ and $\gamma = [1.5 \quad 1.3 \quad 1.1 \quad 0.9 \quad 0.7 \quad 0.5]$. The values of γ are plotted in the insert window.

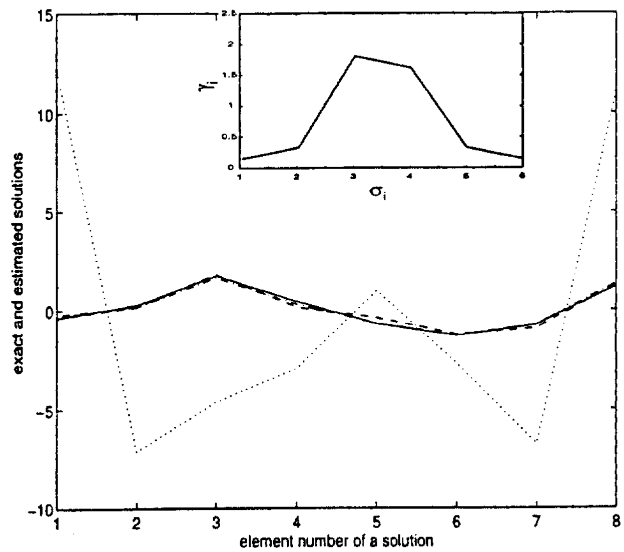

Figure 2: Noise free solution (solid line), solution with no regularization (dotted line), solution using Tikhonov regularization (dashdot line), and TSVD solution (dashed line) for the case when the largest γ values are in the mid-range singular values. The input vector $\mathbf{b}$ satisfies $\mathbf{U}'\mathbf{b} = [0.2 \quad 0.1029 \quad 0.3677 \quad 0.1642 \quad 0.0036 \quad 0.0002]$ and $\gamma = [0.2 \quad 0.4 \quad 2 \quad 1.8 \quad 0.4 \quad 0.2]$. The values of γ are plotted in the insert window.

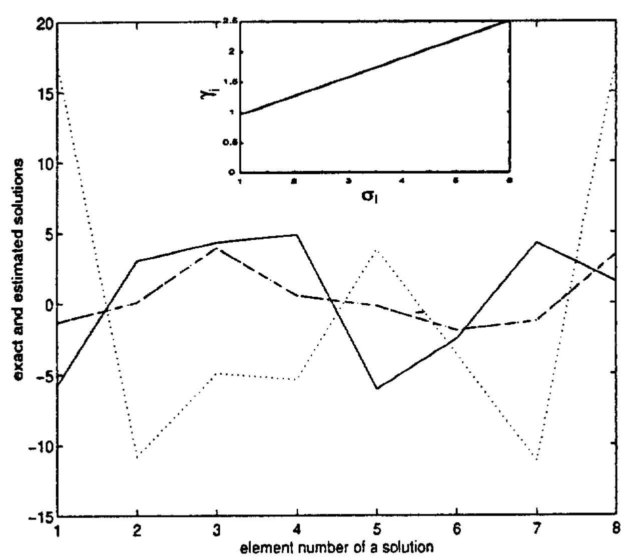

Figure 3: Noise free solution (solid line), solution with no regularization (dotted line), solution using Tikhonov regularization (dashdot line), and TSVD solution (dashed line) for the case when γ is a decreasing function of σ_i. The input vector $\mathbf{b}$ satisfies $\mathbf{U}'\mathbf{b} = [1 \quad 0.4615 \quad 0.5353 \quad 0.4021 \quad 0.0572 \quad 0.008]$ and $\gamma = [1 \quad 1.3 \quad 1.6 \quad 1.9 \quad 2.2 \quad 2.5]$. The values of γ are plotted in the insert window.