

Do *you* precondition on the left or on the right?

Nicole Spillane, Pierre Matalon, and Daniel B. Szyld

Report 25-12-01
December 2025. Revised January 2026.

Department of Mathematics
Temple University
Philadelphia, PA 19122

This report is available in the World Wide Web at
<https://https://faculty.cst.temple.edu/~szyld/>
at the ArXiv at <https://arxiv.org/abs/2512.05160>
and at HAL at <https://hal.science/hal-05389479v1>

Do you precondition on the left or on the right?

Nicole Spillane^[0000-0001-5067-946X] and
Pierre Matalon^[0000-0002-4288-2289] and
Daniel B. Szyld^[0000-0001-8010-0391]

Abstract This work is a follow-up to a poster that was presented at the DD29 conference. Participants were asked the question: “Do you precondition on the left or on the right?”. Here we report on the results of this social experiment. We also provide context on left, right and split preconditioning, share our literature review on the topic, and analyze some of the finer points. Two examples illustrate that convergence bounds can sometimes lead to misleading conclusions.

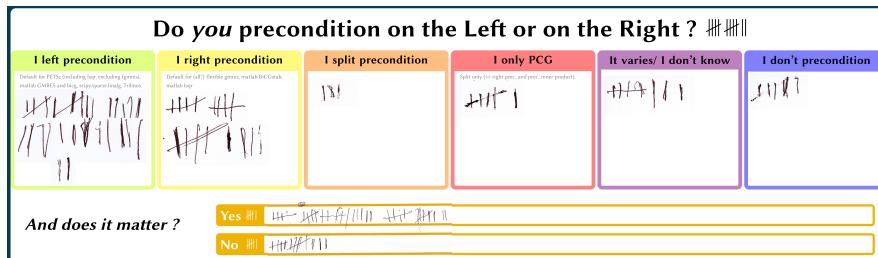


Fig. 1 Attendees of DD29 were asked to vote for their favourite way of preconditioning during the poster session. See Section 4 and Table 1 for more details.

Nicole Spillane

CNRS, CMAP, École polytechnique, Route de Saclay, 91128 Palaiseau cedex, France e-mail:
nicole.spillane@cnrs.fr

Pierre Matalon

CNRS, CMAP, École polytechnique, Route de Saclay, 91128 Palaiseau cedex, France e-mail:
pierre.matalon@polytechnique.edu

Daniel B. Szyld

Department of Mathematics, Temple University, Philadelphia, PA 19122, USA e-mail:
szyld@temple.edu

1 Introduction

This work considers Krylov subspace methods [11, Chapters 6–9] for solving a linear system $\mathbf{Ax} = \mathbf{b}$ given $\mathbf{A} \in \mathbb{C}^{n \times n}$ non-singular and $\mathbf{b} \in \mathbb{C}^n$. More precisely, we focus on the subset of Krylov subspace methods that minimize the residual. These methods start with an initial vector \mathbf{x}_0 , and thus the initial residual is $\mathbf{r}_0 := \mathbf{b} - \mathbf{Ax}_0$. Then, each iteration $k \geq 1$ is characterized by an approximate solution \mathbf{x}_k which is the unique element such that the residual $\mathbf{r}_k := \mathbf{b} - \mathbf{Ax}_k$ satisfies

$$\|\mathbf{r}_k\| = \min_{\mathbf{y} \in \mathbf{A}\mathcal{K}_k(\mathbf{A}, \mathbf{r}_0)} \|\mathbf{r}_0 - \mathbf{y}\|, \quad (1)$$

where the Krylov subspace is $\mathcal{K}_k(\mathbf{A}, \mathbf{r}_0) := \text{span}\{\mathbf{r}_0, \mathbf{Ar}_0, \dots, \mathbf{A}^{k-1}\mathbf{r}_0\}$.

A well-established idea is to accelerate convergence by preconditioning. Given two non-singular matrices \mathbf{H}_L and \mathbf{H}_R , the preconditioned system is

$$\mathbf{H}_L \mathbf{A} \mathbf{H}_R \mathbf{u} = \mathbf{H}_L \mathbf{b}; \text{ and } \mathbf{x} = \mathbf{H}_R \mathbf{u}. \quad (2)$$

We refer to this system as being split-preconditioned by $(\mathbf{H}_L, \mathbf{H}_R)$ in which \mathbf{H}_L is the left preconditioner and \mathbf{H}_R is the right preconditioner. If $\mathbf{H}_L = \mathbf{I}$, then the system is right-preconditioned by \mathbf{H}_R . If $\mathbf{H}_R = \mathbf{I}$, then the system is left-preconditioned by \mathbf{H}_L . Both systems $\mathbf{Ax} = \mathbf{b}$ and (2) are equivalent. The objective is to select the preconditioners in such a way that the Krylov subspace method is more efficient on (2) than on $\mathbf{Ax} = \mathbf{b}$. Applying the characterization (1) to a Krylov subspace method for solving (2) gives, with $\mathbf{r}_k = \mathbf{b} - \mathbf{Ax}_k$,

$$\|\mathbf{H}_L \mathbf{r}_k\| = \min_{\mathbf{y} \in \mathbf{A}\mathcal{K}_k(\mathbf{H}\mathbf{A}, \mathbf{H}\mathbf{r}_0)} \|\mathbf{H}_L(\mathbf{r}_0 - \mathbf{y})\|, \quad (3)$$

where $\mathcal{K}_k(\mathbf{H}\mathbf{A}, \mathbf{H}\mathbf{r}_0) := \text{span}\{\mathbf{H}\mathbf{r}_0, \mathbf{H}\mathbf{A}\mathbf{H}\mathbf{r}_0, \dots, (\mathbf{H}\mathbf{A})^{k-1}\mathbf{H}\mathbf{r}_0\}$ in which the combined preconditioner is defined by $\mathbf{H} := \mathbf{H}_R \mathbf{H}_L$.

It is interesting to note that:

- the minimization space depends only on the choice of the combined preconditioner $\mathbf{H} = \mathbf{H}_R \mathbf{H}_L$,
- the minimized norm depends only on the left preconditioner \mathbf{H}_L .

So far, we have not specified the norm $\|\cdot\|$. In fact, $\|\cdot\|$ can represent any weighted norm $\|\mathbf{y}\|_{\mathbf{W}} := \sqrt{\mathbf{y}^* \mathbf{W} \mathbf{y}}$ with \mathbf{W} Hermitian positive definite. Using such a norm yields the so-called *weighted GMRES* method. Nonetheless, we can choose the Euclidean norm without loss of generality. This is justified by the fact that *weighting is preconditioning by similarity transformations* [6]: unpreconditioned GMRES using the norm $\|\cdot\|_{\mathbf{W}}$ with $\mathbf{W} := \mathbf{L}\mathbf{L}^*$ is equivalent to GMRES using the Euclidean norm and split-preconditioned by $(\mathbf{L}^*, \mathbf{L}^{-*})$. Consequently, the paper can be formulated either from the perspective of preconditioning or of weighted norms. Since the preconditioning viewpoint is more familiar to most GMRES practitioners, we adopt this view, together with the Euclidean norm.

Before setting weighted GMRES aside, we want to emphasize that the triplet $(\mathbf{H}_L, \mathbf{H}_R, \mathbf{W})$ contains redundancies that lead to equivalent configurations. For example, (i) GMRES left-preconditioned by \mathbf{H} in norm $\|\cdot\|_{\mathbf{W}}$ is equivalent to GMRES right-preconditioned by \mathbf{H} in norm $\|\cdot\|_{\mathbf{H}^* \mathbf{W} \mathbf{H}}$; (ii) if $\mathbf{H} := \mathbf{L} \mathbf{L}^*$ is Hermitian positive definite, then GMRES right-preconditioned by \mathbf{H} in norm $\|\cdot\|_{\mathbf{H}}$ is equivalent to GMRES left-preconditioned by \mathbf{H} in norm $\|\cdot\|_{\mathbf{H}^{-1}}$, which is also equivalent to GMRES split-preconditioned by $(\mathbf{L}^*, \mathbf{L})$ in Euclidean norm. In practice, these equivalences offer significant flexibility. In particular, (ii) shows how split preconditioning by the Cholesky factors of a Hermitian positive definite preconditioner can be realized without explicitly forming those factors. The preconditioned conjugate gradient method exploits this technique to preserve symmetry [11, Section 9.2.1]. Weighted GMRES is also valuable for the theoretical analysis of Krylov methods. Firstly, it creates a unifying framework that includes the conjugate gradient (CG) method by remarking that, if \mathbf{A} is Hermitian positive definite, then CG is simply GMRES expressed in the \mathbf{A}^{-1} -norm. Furthermore, many analytical results are established in weighted norm (see, e.g., [3, 14, 15, 16]), leaving open whether they extend to standard Euclidean-norm GMRES. Addressing this latter question is one of the central motivations of the present work. This question has not been fully answered yet.

In the remainder of this work we compare left, right and split preconditioning for unweighted GMRES, which means that $\|\cdot\|$ is the Euclidean norm. We set out to determine which is preferable in terms of convergence and quality of the solution. In order to answer the question, Section 2 presents a literature review from which we extract the key points. Section 3 analyzes some of these in more detail. In a more unconventional fashion, Section 4 presents the results of a social experiment conducted at the DD29 conference on Domain Decomposition Methods where the participants were asked about their preferred way of preconditioning.

2 Literature Review

We have turned to the literature and found surprisingly few concrete answers. In what follows we summarize what we found in the following references:

- *Iterative methods for sparse linear systems* by Y. Saad [11]
- *Preconditioning* by A. J. Wathen [18]
- *Preconditioning and convergence in the right norm* by A. J. Wathen [17]
- *Preconditioning techniques for large linear systems: A survey* by M. Benzi [2]
- *Algorithms for sparse linear systems* by J. Scott and M. Tůma [13]
- *Any nonincreasing convergence curves are simultaneously possible for GMRES and weighted GMRES, as well as for left and right preconditioned GMRES* by P. Matalon and N. Spillane [9]

Iterative methods for sparse linear systems by Y. Saad [11]

(Page 271)¹ “A question arises on the differences between the right, left, and split preconditioning options. The fact that different versions of the residuals are available in each case may affect the stopping criterion and may cause the algorithm to stop either prematurely or with delay. This can be particularly damaging in case $[\mathbf{H}]$ is very ill-conditioned. The degree of symmetry, and therefore performance, can also be affected by the way in which the preconditioner is applied. For example, a split preconditioner may be much better if \mathbf{A} is nearly symmetric. Other than these two situations, there is little difference generally between the three options.”

(Page 271) “When comparing the left, right, and split preconditioning options, a first observation to make is that the spectra of the three associated operators $[\mathbf{H}\mathbf{A}$, $\mathbf{A}\mathbf{H}$, and $\mathbf{H}_L\mathbf{A}\mathbf{H}_R]$ are identical. Therefore, in principle one should expect convergence to be similar, although, as is known, eigenvalues do not always govern convergence.”

(Page 272) “In most practical situations, the difference in the convergence behavior of the two approaches is not significant. The only exception is when $[\mathbf{H}]$ is ill-conditioned which could lead to substantial differences.”

Preconditioning by A. J. Wathen [18]

(Pages 356-357) “One aspect which arises in the non-symmetric but not the symmetric case is the possibility to left-precondition [...], to right-precondition [...], or, if $[\mathbf{H}]$ is available in split form $[\mathbf{H}] = [\mathbf{H}_R\mathbf{H}_L]$, to use [split-preconditioning]. In every case the coefficient matrix is generally again non-symmetric, and obvious similarity transformations involving $[\mathbf{H}]$ or $(\mathbf{H}_L, \mathbf{H}_R)$ show that all three coefficient matrices are mathematically similar and so have the same eigenvalues. There is little evidence, however, that any one of these is better than the others, in terms of its effect on the convergence rate of a Krylov subspace method, even though, theoretically, this could be the case. Traditionally, right-preconditioning is favoured since then GMRES, for example, minimizes the residual $\mathbf{r}_k = \mathbf{b} - \mathbf{A}\mathbf{x}_k$ for the original linear system; left-preconditioning would lead to minimization of the preconditioned residual $[\mathbf{H}\mathbf{r}_k]$. It is certainly not always clear which is preferable: a good preconditioner fortunately seems to give fast convergence for all three forms in practice.”

Preconditioning and convergence in the right norm by A. J. Wathen [17]

(Page 5) “[W]e show how preconditioning affects this balance between approximation and iteration error. In particular for widely used iterative methods which minimise the residual [...] we show that care is needed to avoid selection of preconditioners which apparently give rapid convergence, but which in fact merely distort the relevant norm so that poor solutions are achieved for all but extremely small convergence tolerances.”

Preconditioning techniques for large linear systems: A survey by M. Benzi [2]

(page 420) “Which type of preconditioning to use depends on the choice of the iterative method, problem characteristics, and so forth. For example, with residual

¹ The selected quotes also appear in the first edition of the book published in 1996 with a different page numbering. In both cases the quote can be found in Section 9.3.

minimizing methods, like GMRES, right preconditioning is often used. In exact arithmetic, the residuals for the right-preconditioned system are identical to the true residuals $\mathbf{r}_k = \mathbf{b} - \mathbf{A}\mathbf{x}_k$.

Notice that the matrices $[\mathbf{H}\mathbf{A}$, $\mathbf{A}\mathbf{H}$, and $\mathbf{H}_L\mathbf{A}\mathbf{H}_R]$ are all similar and therefore have the same eigenvalues. If \mathbf{A} and $[\mathbf{H}]$ are spd, the convergence of the CG method will be the same (except possibly for round-off effects) in all cases. On the other hand, in the nonnormal case, solvers like GMRES can behave very differently depending on whether a given preconditioner is applied on the left or on the right (see the discussion in [11, page 272]²; see also [8, page 66] for a striking example)."

Algorithms for sparse linear systems by J. Scott and M. Tůma [13] (page 167)
 "The following result states that it is not possible to determine a priori which variant is the best.

Theorem ([Mendelsohn [10]]: Let δ and Δ be positive numbers. Then, for any $n \geq 3$, there exist non-singular $n \times n$ matrices \mathbf{A} and \mathbf{H} such that all the entries of $(\mathbf{H}\mathbf{A} - \mathbf{I})$ have absolute value less than δ and all the entries of $(\mathbf{A}\mathbf{H} - \mathbf{I})$ have absolute values greater than Δ .

Nevertheless, the choice between left and right preconditioning is still important and may be based on the properties of the coupling of the preconditioner with the iterative method or on the distribution of the eigenvalues of \mathbf{A} . The computed quantities that are readily available during a preconditioned iterative method depend on how the preconditioner is applied and this may influence the choice. These quantities may be used, for example, to decide when to terminate the iterations. An obvious advantage of right preconditioning is that in exact arithmetic, the residuals for the right preconditioned system are identical to the true residuals, enabling convergence to be monitored accurately. In some cases, the numerical properties of an implementation and/or the computer architecture may also play a part."

Any nonincreasing convergence curves are simultaneously possible for [...] left and right preconditioned GMRES by P. Matalon and N. Spillane [9]

(page 3) " Consider two prescribed convergence curves of matching length:

- $\mathbf{r}_0 > \mathbf{r}_1 > \mathbf{r}_2 > \dots$ for right preconditioned GMRES,
- $\tilde{\mathbf{r}}_0 > \tilde{\mathbf{r}}_1 > \tilde{\mathbf{r}}_2 > \dots$ for left preconditioned GMRES.

The first result is that there exists a system $\mathbf{A}\mathbf{x} = \mathbf{b}$ and a preconditioner \mathbf{H} such that both convergence curves are realized. Additionally, the eigenvalues of $\mathbf{A}\mathbf{H}$ can be prescribed. [...] Besides highlighting that GMRES convergence does not solely depend on the eigenvalues, these results show that for some cases, the decision to apply preconditioning on the left or on the right, or to apply weighted GMRES may lead to significant differences in convergence speed. "

From these quotes we draw the following conclusions:

1. The core difference between left and right preconditioning is the residual that gets minimized. This can be seen in the characterization (3) (same \mathbf{H} but different \mathbf{H}_L).

² One of the quotations above by Y. Saad with page number edited for correct edition of book

2. As a consequence of the previous observation, with right preconditioning, the residual norm can be monitored for free throughout the iterations. With left preconditioning, monitoring the preconditioned residual norm can lead to premature or delayed stopping of the algorithm in terms of the unpreconditioned residual norm. This important question, central to [17], is also addressed, *e.g.* in [1, 12].
3. The left, right and split preconditioned matrices have the same spectrum so long as $\mathbf{H}_R \mathbf{H}_L$ is unchanged.
4. It is expected that convergence be faster for the matrix that is close to normality, or to symmetry. Ease of implementation could play a role too.
5. In practice, the behaviour is often not modified by switching the preconditioner between left and right. In fact there are few illustrations of opposite behaviours.
6. There is no easy answer. Indeed, any two convergence curves can be simultaneously prescribed for left and right preconditioned GMRES. And it is possible to find \mathbf{A} and \mathbf{H} such that the entries in $\mathbf{AH} - \mathbf{I}$ are arbitrarily large and those of $\mathbf{HA} - \mathbf{I}$ are arbitrarily small.

3 Analysis

Our own findings are in agreement with the literature review. It is theoretically possible for the behaviours of left and right preconditioned GMRES to be very different but usually nothing much happens. Below we detail three points. Before we go into those details we point out the relation between the eigenvectors in the three cases considered. Let \mathbf{v} be an eigenvector of \mathbf{HA} , then both \mathbf{Av} and $\mathbf{H}^{-1}\mathbf{v}$ are eigenvectors of \mathbf{AH} while both $\mathbf{H}_L\mathbf{Av}$ and $\mathbf{H}_R^{-1}\mathbf{v}$ are eigenvectors of $\mathbf{H}_L\mathbf{AH}_R$ (all corresponding eigenvalues being equal).

3.1 Left and right preconditioned GMRES can only differ if \mathbf{H} is ill-conditioned

If right preconditioned GMRES produces residuals \mathbf{r}_k and left preconditioned GMRES produces residuals $\tilde{\mathbf{r}}_k$, they satisfy

$$\|\mathbf{r}_k\| \leq \|\tilde{\mathbf{r}}_k\| \leq \kappa(\mathbf{H})\|\mathbf{r}_k\|,$$

where $\kappa(\mathbf{H}) = \sigma_{\max}(\mathbf{H})/\sigma_{\min}(\mathbf{H})$ is the condition number of \mathbf{H} . In other words, the residuals produced by left and right preconditioned GMRES, when measured in the Euclidean norm, cannot differ by a factor more than $\kappa(\mathbf{H})$. The proof is simply that

$$\sigma_{\min}(\mathbf{H})\|\mathbf{r}_k\| \leq \sigma_{\min}(\mathbf{H})\|\tilde{\mathbf{r}}_k\| \leq \|\mathbf{H}\tilde{\mathbf{r}}_k\| \leq \|\mathbf{H}\mathbf{r}_k\| \leq \sigma_{\max}(\mathbf{H})\|\mathbf{r}_k\|,$$

where the first and third inequalities are the optimality condition (3) for right and left preconditioned GMRES, while the second and fourth inequalities are spectral estimates.

We can generalize this analysis to two split preconditioners $(\mathbf{H}_L, \mathbf{H}_R)$ and $(\tilde{\mathbf{H}}_L, \tilde{\mathbf{H}}_R)$ that satisfy $\mathbf{H}_R \mathbf{H}_L = \tilde{\mathbf{H}}_R \tilde{\mathbf{H}}_L$. The two minimization spaces are identical. Their corresponding residuals \mathbf{r}_k and $\tilde{\mathbf{r}}_k$ satisfy

$$\sigma_{\min}(\mathbf{G}) \|\mathbf{H}_L \mathbf{r}_k\| \leq \sigma_{\min}(\mathbf{G}) \|\mathbf{H}_L \tilde{\mathbf{r}}_k\| \leq \|\tilde{\mathbf{H}}_L \tilde{\mathbf{r}}_k\| \leq \|\tilde{\mathbf{H}}_L \mathbf{r}_k\| \leq \sigma_{\max}(\mathbf{G}) \|\mathbf{H}_L \mathbf{r}_k\|,$$

with $\mathbf{G} = \tilde{\mathbf{H}}_L \mathbf{H}_L^{-1}$. Again, the first and third inequalities are the optimality conditions, while the second and fourth inequalities are spectral estimates. It follows that:

$$\|\mathbf{H}_L \mathbf{r}_k\| \leq \|\mathbf{H}_L \tilde{\mathbf{r}}_k\| \leq \kappa(\tilde{\mathbf{H}}_L \mathbf{H}_L^{-1}) \|\mathbf{H}_L \mathbf{r}_k\|.$$

Applied to right preconditioning by \mathbf{H} (by setting $(\mathbf{H}_L, \mathbf{H}_R) = (\mathbf{I}, \mathbf{H})$) versus split preconditioning by $(\tilde{\mathbf{H}}_L, \mathbf{H}\tilde{\mathbf{H}}_L^{-1})$ we obtain

$$\|\mathbf{r}_k\| \leq \|\tilde{\mathbf{r}}_k\| \leq \kappa(\tilde{\mathbf{H}}_L) \|\mathbf{r}_k\|,$$

so that right and split preconditioning can differ in Euclidean residual norm only by a factor $\kappa(\tilde{\mathbf{H}}_L)$.

3.2 Different convergence bounds, same behaviour (1/2)

The striking example cited by M. Benzi in [2] is from the PhD dissertation of S. Lee [8, page 66], and is adapted from the work of A. S. Householder [7, page 96]. It consists in a matrix \mathbf{A} and a preconditioner \mathbf{H} such that \mathbf{H} is a poor right inverse and a good left inverse in two senses: entrywise and field of values.

The example is as follows. Let \mathbf{A} be a real matrix with eigenvalues in a real positive interval $[\mu, \lambda]$ with $0 < \mu < \lambda$. Let \mathbf{u} (respectively \mathbf{v}) be a unit right (respectively left) eigenvector associated with λ (respectively μ), *i.e.*,

$$\mathbf{A}\mathbf{u} = \lambda\mathbf{u}; \quad \mathbf{v}^\top \mathbf{A} = \mu\mathbf{v}^\top; \quad \|\mathbf{u}\| = \|\mathbf{v}\| = 1; \text{ and } \mu \neq \lambda.$$

The proposed preconditioner is then a rank-one perturbation of \mathbf{A}^{-1} defined by

$$\mathbf{H} := \mathbf{A}^{-1} + \mathbf{u}\mathbf{v}^\top.$$

It is easy to see that the preconditioned operators are $\mathbf{H}\mathbf{A} = (\mathbf{A}^{-1} + \mathbf{u}\mathbf{v}^\top)\mathbf{A} = \mathbf{I} + \mu\mathbf{u}\mathbf{v}^\top$ and $\mathbf{A}\mathbf{H} = \mathbf{A}(\mathbf{A}^{-1} + \mathbf{u}\mathbf{v}^\top) = \mathbf{I} + \lambda\mathbf{u}\mathbf{v}^\top$. Moreover,

$$\mathbf{I} - \mathbf{H}\mathbf{A} = -\mu\mathbf{u}\mathbf{v}^\top \text{ and } \mathbf{I} - \mathbf{A}\mathbf{H} = -\lambda\mathbf{u}\mathbf{v}^\top.$$

so the entries of $\mathbf{I} - \mathbf{H}\mathbf{A}$ and $\mathbf{I} - \mathbf{A}\mathbf{H}$ are in the ratio λ/μ which can be very large. This illustrates [13, Theorem 9.2] (originally from [10]) cited in Section 2.

Moreover, the field of values of $\mathbf{H}\mathbf{A}$ is the circle centered at 1 of radius $\mu/2$ and the field of values of $\mathbf{A}\mathbf{H}$ is the circle centered at 1 of radius $\lambda/2$ (see below for a proof). A well-known convergence bound for GMRES, the Elman bound [4],[5, §5.3] relates the rate of convergence of GMRES to the field of values of the operator. In particular, in the case where the field of values is a small disk around 1, convergence is provably fast. When 0 is included in the field of values, Elman's bound does no longer apply. For this reason, if $\lambda > 1 \gg \mu$ it is tempting to conclude that \mathbf{H} is a poor right-hand side inverse and a good left-hand side inverse in terms of the field of values.

However, when it comes to convergence of GMRES, the reality is quite different and this can be seen by observing that

$$(\mathbf{H}\mathbf{A})^2 = (\mathbf{I} + \mu\mathbf{u}\mathbf{v}^\top)^2 = \mathbf{I} + 2\mu\mathbf{u}\mathbf{v}^\top + \mu^2\mathbf{u}\mathbf{v}^\top\mathbf{u}\mathbf{v}^\top = \mathbf{I} + 2\mu\mathbf{u}\mathbf{v}^\top = 2\mathbf{H}\mathbf{A} - \mathbf{I},$$

where the third equality is the result of $\mathbf{v} \perp \mathbf{u}$ since \mathbf{u} and \mathbf{v} are a pair of right and left eigenvectors associated with two different eigenvalues. Consequently, the grade of any vector with respect to $\mathbf{H}\mathbf{A}$ is 2 and, both left and right preconditioned GMRES converge in at most 2 iterations (for any right hand side and any starting vector).

Proof of field of value formulae. The field of values of $\mathbf{A}\mathbf{H}$ is the set of all complex numbers of the form

$$\frac{\mathbf{z}^*\mathbf{A}\mathbf{H}\mathbf{z}}{\mathbf{z}^*\mathbf{z}} = 1 + \lambda \frac{\mathbf{z}^*\mathbf{u}\mathbf{v}^\top\mathbf{z}}{\mathbf{z}^*\mathbf{z}} = 1 + \lambda \frac{\mathbf{z}^*\mathbf{u} \cdot \mathbf{v}^\top \mathbf{z}}{\mathbf{z}^*\mathbf{z}}; \mathbf{z} \in \mathbb{C}^n,$$

with \mathbf{z}^* the conjugate transpose of \mathbf{z} . Applying again that $\mathbf{v} \perp \mathbf{u}$, any $\mathbf{z} \in \mathbb{C}^n$ can be split into $\mathbf{z} = \xi_1\mathbf{u} + \xi_2\mathbf{v} + \mathbf{w}$ with $\mathbf{w} \perp \mathbf{u}$, $\mathbf{w} \perp \mathbf{v}$, and $\xi_1, \xi_2 \in \mathbb{C}$. It follows that

$$\frac{\mathbf{z}^*\mathbf{A}\mathbf{H}\mathbf{z}}{\mathbf{z}^*\mathbf{z}} = 1 + \lambda \frac{\overline{\xi_1}\xi_2}{|\xi_1|^2 + |\xi_2|^2 + \|\mathbf{w}\|^2}.$$

The inclusion of the circle in the field of values can be seen by setting $\mathbf{w} = 0$, $\xi_1 = e^{-i\theta}$ and $\xi_2 = e^{i\theta}$ for all values of $\theta \in [0, \pi]$:

$$\frac{\mathbf{z}^*\mathbf{A}\mathbf{H}\mathbf{z}}{\mathbf{z}^*\mathbf{z}} = 1 + \frac{\lambda}{2}e^{2i\theta}.$$

The inclusion of the field of values in the circle is the result of

$$2|\overline{\xi_1}\xi_2| \leq |\xi_1|^2 + |\xi_2|^2 \leq |\xi_1|^2 + |\xi_2|^2 + \|\mathbf{w}\|^2 = \|\mathbf{z}\|^2.$$

The proof for the field of values of $\mathbf{H}\mathbf{A}$ is almost identical.

3.3 Different convergence bounds, same behaviour (2/2)

The convergence bound [11, Prop. 6.32] involving the eigenvalues and the conditioning of the eigenvectors is a well-known convergence bound for GMRES. In this section, we build a family of matrices and preconditioners such that this bound is

small when the preconditioner is applied on the right, large when it is applied on the left, and yet lead to the same convergence speed in practice.

Assume a non-singular, diagonalizable matrix $\mathbf{A} := \mathbf{P}\Lambda\mathbf{P}^{-1}$, where Λ is a diagonal matrix holding the eigenvalues of \mathbf{A} , and \mathbf{P} is a matrix whose columns hold the eigenvectors of \mathbf{A} . The bound [11, Prop. 6.32] for the GMRES method applied to the system $\mathbf{Ax} = \mathbf{b}$ reads

$$\frac{\|\mathbf{r}_k\|}{\|\mathbf{r}_0\|} \leq \epsilon^{(k)} \kappa(\mathbf{P}), \quad (4)$$

where

$$\epsilon^{(k)} := \min_{q \in \mathbb{P}^k, q(0)=1} \max_{\lambda \in \sigma(\mathbf{A})} |q(\lambda)|.$$

In $\epsilon^{(k)}$, \mathbb{P}^k denotes the space of polynomial of degree at most k . This term accounts for the spectrum of \mathbf{A} , while the second term of the bound, $\kappa(\mathbf{P})$, accounts for the eigenvectors of \mathbf{A} . Note that $\kappa(\mathbf{P})$ is often seen as a measure of \mathbf{A} 's departure from normality, in the sense that if \mathbf{A} is normal, then \mathbf{P} is unitary and $\kappa(\mathbf{P}) = 1$. On the other hand, if $\kappa(\mathbf{P})$ is large, it means that some eigenvectors are close to being linearly dependent, thus making $\mathbf{P}^*\mathbf{P}$ far from identity, and therefore $\mathbf{A}^*\mathbf{A}$ far from \mathbf{AA}^* .

Let us build a test case where left- and right-preconditioning lead to significantly different bounds. Let Λ be an invertible, well-conditioned, diagonal matrix. Let \mathbf{X} be an ill-conditioned invertible matrix. Let \mathbf{Y} be a unitary matrix. Now, define the matrix $\mathbf{A} := \mathbf{X}\Lambda\mathbf{Y}^{-1}$ and the preconditioner $\mathbf{H} := \mathbf{Y}\mathbf{X}^{-1}$. We then have the left- and right-preconditioned matrices $\mathbf{HA} = \mathbf{Y}\Lambda\mathbf{Y}^{-1}$ and $\mathbf{AH} = \mathbf{X}\Lambda\mathbf{X}^{-1}$.

Solving $\mathbf{Ax} = \mathbf{b}$ with left-preconditioned GMRES applies GMRES to the system $\mathbf{H}\mathbf{Ax} = \mathbf{H}\mathbf{b}$. Denoting by $\tilde{\mathbf{r}}_k := \mathbf{H}(\mathbf{b} - \mathbf{Ax}_k)$ the k^{th} residual produced, (4) writes

$$\frac{\|\tilde{\mathbf{r}}_k\|}{\|\tilde{\mathbf{r}}_0\|} \leq \epsilon^{(k)}.$$

Indeed, since \mathbf{Y} is unitary, $\kappa(\mathbf{Y}) = 1$. We expect fast convergence in this case due to the sole dependency of the eigenvalues, which we have selected to be clustered. Similarly, right-preconditioned GMRES yields the bound

$$\frac{\|\mathbf{r}_k\|}{\|\mathbf{r}_0\|} \leq \epsilon^{(k)} \kappa(\mathbf{X}).$$

Since \mathbf{X} is ill-conditioned, this bound is large: we (naively) expect slow convergence. However, we present below an example for which both cases converge almost at the same speed, with only small discrepancies in the successive residual norms produced.

Experimental setup: \mathbf{Y} is obtained as the orthogonal matrix produced by the QR factorization of a randomly generated square matrix. \mathbf{X} is, first, initialized the same way. Then, denoting by \mathbf{x}_i the i^{th} column of \mathbf{X} , we replace \mathbf{x}_1 by $\frac{1}{K}\mathbf{x}_1 + \mathbf{x}_n$, where K is a large real number. This way, \mathbf{x}_1 is close to collinear to \mathbf{x}_n , while preserving its strict linear independence from all other columns. The coefficient K drives the condition number of \mathbf{X} . \mathbf{A} and \mathbf{H} are defined accordingly, and \mathbf{b} is randomly generated. Figure 2 shows such an example with a system of order 100 and $K := 10^8$.

($\approx \kappa(\mathbf{X})$). For right-preconditioning, we plot the residual norms $\|\mathbf{r}_k\|/\|\mathbf{r}_0\|$, and for left-preconditioning, we plot the preconditioned residual norms $\|\tilde{\mathbf{r}}_k\|/\|\tilde{\mathbf{r}}_0\|$. Note that the unpreconditioned residuals in this latter case sensibly produce the same convergence curve as the preconditioned residuals.

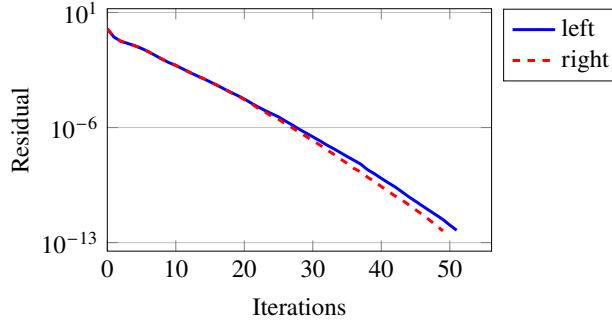


Fig. 2 Example of left-right preconditioning leading to similar convergence behaviour while presenting significantly different convergence bounds.

4 A social experiment at DD29

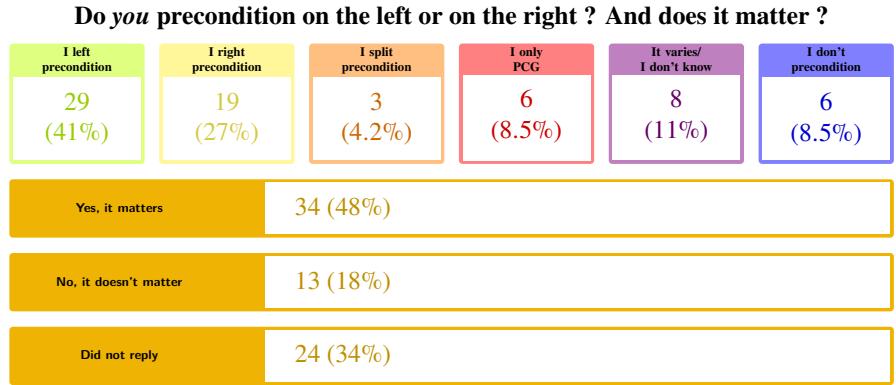


Table 1 Results of the poll. There are 71 replies to the first question. All percentages are with respect to these 71 participants. Figure 1 shows how the questions were displayed on the poster.

During the poster session at the DD29 conference on Domain Decomposition methods, the authors presented a poster which included a poll. The participants at

the conference were asked to vote in favour of their favourite way of preconditioning. The possible answers and results of the poll are shown in Table 1. There were 299 attendees at the conference according to the organizers. Of these, 71 people replied to the first question in the poll which is “Do you precondition on the left or on the right”. The participation rate is 24%. Significantly fewer people (47) replied to the second question which is “And does it matter ?”. In the presentation of the results in Table 1 we included a “Did not reply” field to account for this.

Preconditioned conjugate gradients (PCG) is given as a separate possible response to the first question since there is one overwhelmingly natural way of preconditioning the conjugate gradient method that is neither left- nor right-preconditioning [11, Section 9.2.1] but preserves the short recurrence property.

It turns out that 41% of respondents apply left preconditioning. One possible explanation is that left preconditioning is the default choice in several libraries including PETSc (with the exception of fgmres but including lsqr), MATLAB’s gmres and bicg, Python’s scipy.sparse.linalg and Trilinos. Note that right preconditioning is the default for MATLAB solvers bicgstab and lsqr, and the necessary choice for all implementations of flexible GMRES. In second place, 27% of respondents apply right preconditioning. The argument in favour of right preconditioning is to minimize and monitor the residual in the Euclidean norm³. Finally, there are 2.6 times more people who consider the choice between left and right preconditioning to be important than those who don’t.

Both authors that were present at the conference responded that they left-precondition and that yes, it does matter. Our reasoning is that if \mathbf{H} is a good preconditioner for \mathbf{A} then the left preconditioned residual is a reasonable measure of the error. Since we are aware that the preconditioned residual norm may be much larger than the unpreconditioned residual norm, we do also compute the unpreconditioned residual (but not at every iteration) so that the solver does not stop prematurely.

5 Conclusion

The purpose of this work is to question whether there is a best way to position a given preconditioner. We have done our best to reflect existing considerations from the literature as well as to offer some enlightenments. Our examples show that convergence bounds that predict possible opposite behaviours between left and right preconditioned GMRES are not always correct predictions. If there is a norm that is particularly useful or meaningful for a given problem then it should be the norm that is minimized by GMRES. Otherwise, either choice is valid but GMRES users should be aware of what they are doing and perhaps compute several different residual norms for the output solution.

³ Even though the poll was not conducted in any other scientific meetings, the conversations that the first author has had with many scientists led her to believe that right preconditioning would have been the winner at a numerical linear algebra conference.

References

1. Ashby, S.F., Holst, M.J., Manteuffel, T.A., Saylor, P.E.: The role of the inner product in stopping criteria for conjugate gradient iterations. *BIT* **41**(1), 26–52 (2001). DOI 10.1023/A:1021961616621
2. Benzi, M.: Preconditioning techniques for large linear systems: A survey. *J. Comput. Phys.* **182**(2), 418–477 (2002). DOI 10.1006/jcph.2002.7176
3. Cai, X.C., Widlund, O.B.: Domain decomposition algorithms for indefinite elliptic problems. *SIAM J. Sci. Stat. Comput.* **13**(1), 243–258 (1992). DOI 10.1137/0913013. URL semantic-scholar.org/paper/106a429bba2ed3c761256c704edacd40770b8e1b
4. Eisenstat, S.C., Elman, H.C., Schultz, M.H.: Variational iterative methods for nonsymmetric systems of linear equations. *SIAM J. Numer. Anal.* **20**, 345–367 (1983). DOI <https://doi.org/10.1137/0720023>
5. Elman, H.C.: Iterative methods for large, sparse, nonsymmetric systems of linear equations. Ph.D. thesis, Department of Computer Science, Yale University (1982). URL <http://ftp.cs.yale.edu/publications/techreports/tr229.pdf>. Research Report #229
6. Gutknecht, M.H., Loher, D.: Preconditioning by similarity transformations: Another valid option? abstract, GAMM Workshop on Numerical Linear Algebra (2001)
7. Householder, A.S.: The theory of matrices in numerical analysis. Blaisdell, Waltham, Mass. (1964). Reprinted by Dover, New York, 1975
8. Lee, S.L.: Krylov methods for the numerical solution of initial-value problems in differential-algebraic equations. Tech. rep., University of Illinois at Urbana-Champaign (1993)
9. Matalon, P., Spillane, N.: Any nonincreasing convergence curves are simultaneously possible for GMRES and weighted GMRES, as well as for left and right preconditioned GMRES. Preprint, arXiv:2506.17193 [math.NA] (2025). URL <https://arxiv.org/abs/2506.17193>
10. Mendelsohn, N.S.: Some properties of approximate inverses of matrices. *Proc. Trans. R. Soc. Can. III* (3) **50**, 53–59 (1956)
11. Saad, Y.: Iterative methods for sparse linear systems., 2nd ed. Philadelphia, PA: SIAM Society for Industrial and Applied Mathematics (2003)
12. Sarkis, M., Szyld, D.B.: Optimal left and right additive Schwarz preconditioning for minimal residual methods with Euclidean and energy norms. *Comput. Methods Appl. Mech. Eng.* **196**(8), 1612–1621 (2007). DOI 10.1016/j.cma.2006.03.027
13. Scott, J., Tůma, M.: Algorithms for sparse linear systems. Nečas Cent. Ser. Cham: Springer (2023). DOI 10.1007/978-3-031-25820-6
14. Spillane, N.: Hermitian preconditioning for a class of non-Hermitian linear systems. *SIAM J. Sci. Comput.* **46**(3), a1903–a1922 (2024). DOI 10.1137/23M1559026
15. Spillane, N., Szyld, D.B.: New convergence analysis of GMRES with weighted norms, preconditioning, and deflation, leading to a new deflation space. *SIAM J. Matrix Anal. Appl.* **45**(4), 1721–1745 (2024). DOI 10.1137/23M1622398
16. Spillane, N., Szyld, D.B.: Improved Polynomial Bounds and Acceleration of GMRES by Solving a min-max Problem on Rectangles, and by Deflating. Preprint, arXiv:2504.05723 [math.NA] (2025). URL <https://arxiv.org/abs/2504.05723>
17. Wathen, A.: Preconditioning and convergence in the right norm. *Int. J. Comput. Math.* **84**(8), 1199–1209 (2007). DOI 10.1080/00207160701355961. URL ora.ox.ac.uk/objects/uuid:355787b7-ef50-4765-8f9a-970989034ffa
18. Wathen, A.J.: Preconditioning. *Acta Numer.* **24**, 329–376 (2015). DOI 10.1017/S0962492915000021