

Article

The Case for Adopting the Sequential Jacobi's Diagonalization Algorithm in Neutrino Oscillation Physics

Gustavo A. Valdivieso 

Instituto de Ciências e Tecnologia, Universidade Federal de Alfenas UNIFAL-MG,
Poços de Caldas 37715-400, MG, Brazil; gustavo.valdivieso@unifal-mg.edu.br

Abstract: Neutrino flavor oscillations and conversion in an interacting background (MSW effects) may reveal the charge-parity violation in the next generation of neutrino experiments. The usual approach for studying these effects is to numerically integrate the Schrödinger equation, recovering the neutrino mixing matrix and its parameters from the solution. This work suggests using the classical Jacobi's diagonalization in combination with a reordering procedure to produce a new algorithm, the Sequential Jacobi Diagonalization. This strategy separates linear algebra operations from numerical integration, allowing physicists to study how the oscillation parameters are affected by adiabatic MSW effects in a more efficient way. The mixing matrices at every point of a given parameter space can be stored for speeding up other calculations such as model fitting and Monte Carlo productions. This approach has two major computation advantages, namely, being trivially parallelizable, making it a suitable choice for concurrent computation, and allowing for quasi-model-independent solutions which simplify Beyond Standard Model searches.

Keywords: numerical methods; Jacobi diagonalization; neutrino oscillations; MSW effect; BSM searches



Citation: Valdivieso, G.A. The Case for Adopting the Sequential Jacobi's Diagonalization Algorithm in Neutrino Oscillation Physics.

Symmetry **2024**, *16*, 1617. <https://doi.org/10.3390/sym16121617>

Academic Editor: Stefania Ricciardi

Received: 24 July 2024

Revised: 12 August 2024

Accepted: 14 August 2024

Published: 6 December 2024



Copyright: © 2024 by the author. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

1. Introduction

Neutrino flavor oscillations in the presence of matter are described by a continuously varying, finite-dimension set of Schrödinger equations along a propagation path, in what is known as the Mikheyev–Smirnov–Wolfenstein (MSW) effect [1,2]. The varying nature of such backgrounds prevents any practical case from being solved analytically, with numerical methods being the only option. Although these computations by themselves are not intensive, the problem scales in complexity when dealing with any sort of model fitting or any situation requiring the system to be solved for a large number of configurations. Moreover, knowing the values of the oscillation parameters, mixing angles and mass differences as a function of a model's parameter space gives valuable insights into neutrino physics itself. In order to map the oscillation parameters, the solutions are used to find the eigenvalues and eigenvectors of the mixing matrix, which adds a second numerical task to the computational load. This second step has an inherent complication: by definition, linear algebra algorithms are agnostic to eigenvalues and eigenvector ordering. In fact, the output of a numerical diagonalization algorithm has unpredictable ordering, forming what is known as Newton's Fractal [3]. However, this ordering has implications for neutrino physics. The Pontecorvo–Maki–Nakagawa–Sakata (PMNS) parametrization assumes a mass/flavor ordering when defining the mixing matrix in vacuum, and this ordering has to be known when matter effects are present in order to properly recover the oscillation parameters. This work makes the case for the use of Jacobi's diagonalization algorithm [4–6], which finds both the set of eigenvalues and eigenvectors at the same time. This is accomplished by swapping the diagonalizing with the integration steps, finding the eigenvectors before solving the Schrödinger equation. This allows us to find the correct ordering of its eigenvalues by a simple comparison to neighbors, maintaining the PMNS

parametrization which connects the matter-affected oscillation parameters to their vacuum counterparts. By storing the eigenvectors in a look-up table (LUT), the strategies proposed here can trivially offload the computation of fitting and mapping tasks. This procedure is particularly advantageous for exploring exotic matter backgrounds and flavor-changing Beyond Standard Model (BSM) interactions, which are common avenues of study in the neutrino research field. Since future experiments such as the Deep Underground Neutrino Experiment (DUNE) [7] will be able to explore BSM physics, new algorithms with better computational complexity are required.

The modern version of the Jacobi method for the diagonalization of Hermitian matrices is reviewed, followed up by the addition of extra steps with the goal of preserving any pre-existing parametrization. This is the Sequential Diagonalization Strategy (SDS), which is accomplished by successively comparing the eigenvalues over an arbitrary smooth path in small, discrete steps, transporting the ordering of the eigenvector across the parameter space. This algorithm offers advantages to neutrino physicists, mainly when studying adiabatic evolution in active interacting media. See Refs. [8–10] for excellent reviews. Although analytical solutions do exist for the most common scenarios [11,12], these are Standard Model (SM)-dependent and cannot be easily expanded for more general models, such as the study of Non-Standard Neutrino Interactions (NSNIs) and sterile neutrinos. It has been shown that sequential diagonalizations can replace first- and second-order terms when perturbatively calculating the oscillation parameters [13], though this is limited to uniform backgrounds. The methods discussed here are not limited by the model being considered, nor by the medium's uniformity.

Although the methods described here are aimed at neutrino physicists, this paper is organized in such a way that the methodology can be appreciated by a more general reader. Section 2 defines the ordering problem and its implications; Section 3 reviews the original Jacobi algorithm for a Hermitian matrix; Section 4 outlines the SDS, which ensures parametrization over a continuous path and; Finally, Section 5 illustrates an example application in neutrino physics using an example which can also be compared with an analytical solution, followed by Section 6 with the conclusions. After this, two appendices showcase discussions and details that might not be of interest for the general reader: Appendix B performs a benchmark test by solving a random case and comparing it with its analytical solution and; finally, Appendix A contains an analysis of the convergence, precision and stability of the algorithm.

2. Parametrized Hermitian Matrices

Consider a Hermitian matrix A of order n , with $n(n-1)/2$ independent elements $A_{jk} \in \mathbb{C}$. While developing a physical model, one may want to describe each element as a continuous function over a p -dimensional parameter space, i.e., $A_{jk} \equiv A_{jk}(\vec{q})$, with $\{\vec{q} = (q_1, \dots, q_p) \mid q_i \in \mathbb{R}\}$. In this situation, its real eigenvalues λ_k and corresponding eigenvectors V_k are also functions of $\vec{q}$. By the spectral theorem, all Hermitian matrices are normal matrices and, as such, can be written as $A = U D U^\dagger$, where $D = \text{diag}(\lambda_1, \lambda_2, \dots, \lambda_n)$ and

$$U = \left[\begin{pmatrix} \vdots \\ V_1 \\ \vdots \end{pmatrix} \begin{pmatrix} \vdots \\ V_2 \\ \vdots \end{pmatrix} \cdots \begin{pmatrix} \vdots \\ V_n \\ \vdots \end{pmatrix} \right] \quad (1)$$

where U is a unitary transformation, i.e., $U U^\dagger = 1$. The ordering of the λ_k in D might be arbitrary but is assumed to have physical meaning, and so it has to be preserved. Note that, by choosing to represent the eigenvalues as the elements of D and the eigenvector as a column of U , their ordering is preserved in these matrices by definition, and the relation $A V_k = \lambda_k V_k$ becomes equivalent to $A U = U D$. For the sake of brevity, the pair $\{D, U\}$ will be referred to as the eigensystem of A .

We are interested in describing the eigensystem of Hermitian matrices as a function of the original parametrization $\vec{q}$, i.e., since $A \equiv A(\vec{q})$, so it must be that $D \equiv D(\vec{q})$ and

$U \equiv U(\vec{q})$. The Jacobi diagonalization method is well suited for this goal since it results in the complete eigensystem, already in the form $\{D, U\}$. However, the resulting ordering of their columns is never guaranteed since this information is arbitrary and not related to the method's input, A . In fact, the resulting ordering is unpredictable, and its dependence on the elements A_k is fractal-like, presenting self-similarity and recursive characteristics [3]. In practice, this means that by defining the first element of $D(\vec{q})$ as $\lambda_1(\vec{q})$, one should not assume that the first element of $D(\vec{q} + \Delta\vec{q})$ is still $\lambda_1(\vec{q} + \Delta\vec{q})$, even after an arbitrarily small step $\Delta\vec{q}$. This implication prevents us from reconstructing the functions $\lambda_k(\vec{q})$ and $V_k(\vec{q})$, unless we transport the ordering information along, with every step. This strategy will be addressed in Section 4, after the following review of the Jacobi method.

3. Jacobi's Algorithm

The original algorithm proposed by Carl G. J. Jacobi in 1845 [4] established a numerical procedure to calculate eigenvalues of a real, symmetric matrix. Since then, several variations have been developed in the literature, including an extension to a general complex matrix [5,6]. The focus of this work is the diagonalization of Hermitian matrices, which is reviewed in this section for completion's sake, using notation and steps mainly based on ref. [14]. For a more comprehensive review, see ref. [15].

Let us start by defining a way to measure the magnitude of a matrix's off-diagonal elements, d^2 given by

$$d^2 = \frac{2}{n(n-1)} \sum_{i>j}^n A_{ij} A_{ji}, \quad (2)$$

Jacobi proved that for all Hermitian matrices A , there is an infinite sequence $A_0, A_1, \dots, A_k, A_{k+1}, \dots$, with off-diagonal magnitudes $d_{k+1}^2 < d_k^2$ for any k , meaning that A_k converges to a diagonal matrix as $k \rightarrow \infty$. The sequence in Equation (2) has a general term given by

$$\begin{aligned} A_{k+1} &= S_k^\dagger A_k S_k, \quad \text{or} \\ &= S_k^\dagger S_{k-1}^\dagger \dots S_0^\dagger A_0 S_0 \dots S_{k-1} S_k, \end{aligned} \quad (3)$$

with each matrix S_k being a unitary transformation that has to be constructed. The strategies for constructing S_k will be discussed in a moment. From the definitions in Section 2, A can be factored as a diagonal matrix D and a unitary transformation U , as $A = U D U^\dagger$. This relation can be inverted in order to express D as a function of A and U ,

$$D = U^\dagger A U \quad (4)$$

which is recognizable as the limit of Equation (3) when $A = A_0$, with $A_k \rightarrow D$ and $S_0 \dots S_k \rightarrow U$. From the perspective of a numerical approximation, one may stop the sequence $\{A_k\}$ when the condition $d_k^2 \leq \varepsilon^2$ is met, for an arbitrary precision ε . In this case, Equation (3) can be read as

$$D \approx A_k = S^\dagger A_0 S, \text{ for large enough } k, \quad (5)$$

and

$$U \approx S, \text{ with } S \equiv S_0 S_1 \dots S_k, \quad (6)$$

with a global truncation error $E \leq \varepsilon$.

Several strategies are available for constructing the sequence of rotations that satisfies these definitions. The total computational complexity depends on the number of steps in the sequence and which decisions are considered between each one. In particular, the S_k can be organized in groups called sweeps where all the off-diagonal elements are systematically

rotated, one by one, in what is known as the Cyclic Jacobi Method (CJM) [6,16,17]. This strategy requires no decision making regarding the elements themselves, thus employing the least amount of time per step. Another known strategy is to eliminate the largest remaining off-diagonal element with each rotation, which is Jacobi's original strategy [4]. This strategy is proven to have quadratic convergence [15], at the cost of a search for the largest element, between rotations. This is the implementation chosen for this work, which was confirmed to achieve quadratic convergence, with the test and its results presented in the Appendix A. For a modern review and variations on the implementation presented here, please refer to refs. [14,15] and references therein.

A Jacobi rotation S_k represents a single step in the process of diagonalizing the target matrix A , and according to the chosen strategy, it is applied to the largest off-diagonal element, A_{rc} , with $r \neq c$. Each S_k can be decomposed into two consecutive rotations, $S_k = K^{[rc]} G^{[rc]}$, which was first introduced by W. Givens [14], with K and G known as Givens rotations [14]. Each of these independent rotations, K and G , is responsible for rotating away one of the two degrees of freedom of this element, since A_{rc} is a complex number. In other words, the first rotation $A' = K^{\dagger[rc]} A K^{[rc]}$ makes the resulting element A'_{rc} real, while the second one $A'' = G^{\dagger[rc]} A' G^{[rc]}$ is responsible for vanishing with A''_{rc} . Under these requirements, $K^{[rc]}$ may be written as

$$K^{[rc]} = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 & & & \\ & e^{i\theta_1} & \cdots & e^{-i\theta_1} \\ & \vdots & 1 & \vdots \\ & -e^{i\theta_1} & \cdots & e^{-i\theta_1} \\ & & & 1 \end{bmatrix}, \quad (7)$$

with the main elements given by $K_{rr}^{[rc]} = e^{+i\theta_1}/\sqrt{2} = K_{cc}^{\star[rc]}$, $K_{rc}^{[rc]} = e^{i\theta_1}/\sqrt{2} = -K_{cr}^{\star[rc]}$, and $K_{jj}^{[rc]} = 1$ for the diagonal elements, except at rr and cc . All other elements are zero. By imposing that $\text{Im}\{A'_{rc}\} = 0$, the rotation angle θ_1 becomes

$$\tan 2\theta_1 = \frac{\text{Im}\{A_{rc}\}}{\text{Re}\{A_{rc}\}}. \quad (8)$$

One can verify that, in the case where the target matrix A is already real, $\theta_1 = 0$ and $K^{[rc]}$ becomes the identity 1. The second transformation $G^{[rc]}$ must be a real rotation,

$$G^{[rc]} = \begin{bmatrix} 1 & & & \\ & \cos \theta_2 & \cdots & \sin \theta_2 \\ & \vdots & 1 & \vdots \\ & -\sin \theta_2 & \cdots & \cos \theta_2 \\ & & & 1 \end{bmatrix}, \quad (9)$$

with notation analogous to the one used in Equation (7), $G_{rr}^{[rc]} = \cos \theta_2 = K_{cc}^{[rc]}$, $K_{rc}^{[rc]} = \sin \theta_2 = -K_{cr}^{[rc]}$, with $K_{jj}^{[rc]} = 1$ for the diagonal elements, except at rr and cc , with all others being zero. The rotation angle θ_2 responsible for vanishing with the element in position rc is given by

$$\tan 2\theta_2 = \frac{2 A'_{rc}}{A'_{cc} - A'_{rr}}. \quad (10)$$

Extra care should be taken when calculating the angles θ_1 and θ_2 since Equations (8) and (10) are prone to overflow when numerically evaluating $\tan^{-1}$. The final implementation was tested with random 3×3 matrices, so that the numerical results could be compared to the analytical ones (see Appendix B). In the next section, the discussion returns to how

to preserve the eigenvalues ordering, making the Jacobi's algorithm suitable for studying adiabatic matter effects in neutrino physics.

4. Sequential Diagonalization Strategy

As previously discussed in Section 2, once the parametrization of a Hermitian matrix is defined, the objective is to obtain its eigensystem as a function of a given model's parameters, correcting for the randomness in the eigensystem ordering. The solution proposed here was loosely inspired by the parallel transport of the tangent vector, in Riemannian geometry. Since the change in both eigenvalues and eigenvectors is continuous, it should be possible to detect any unwanted reordering by a simple comparison of neighboring results. This requires connecting the point of study in the parameter space with another where the ordering is known, diagonalizing and correcting the ordering along the way, transporting the eigensystem from one point to the other. This method will be referred to as the Sequential Diagonalization Strategy (SDS) which can be summarized as follows:

- Starting from a point in the parameter space where the eigensystem is known (including ordering), a small step is taken to a new position;
- The eigensystem is obtained in this new position, and a comparison is drawn between the original set of eigenvalues and the new ones;
- Assuming that the step is short enough, it is always possible to arrive at a one-to-one match between the two sets, which allows the post-step eigensystem to be *reordered* following their pre-step counterparts;
- This is now regarded as a new reference value, and the process is repeated over a predefined path, transporting the known ordering along it.

Figure 1 illustrates this strategy with a graphical example. In what follows, a formal definition of SDS is presented.




Figure 1. Graphical summary of the Sequential Diagonalization Strategy (SDS). Each block shows the resulting diagonalizing transformation U , where each column represents an eigenvector, with its corresponding eigenvalue marked below it. The background arrow shows the sequence of diagonalizations, with the top row following the Jacobi diagonalization while the bottom row shows the reordered matrices. (A) is the starting point, representing the ordering to be preserved. It is stored as the first reference, shown in (B); after a small step $d\vec{q}$, a new eigensystem is obtained with the Jacobi method, shown as (C). This is compared with the previous reference, using Equation (15), in order to know the correct way to reorder (C) to match (B); (C) is stored as a new reference, and the cycle repeats for the next step (E). In summary, while (A,C,E) result from the diagonalization, with random ordering, (B,D,F) are the values actually stored.

From Section 2, the relation $A = U D U^\dagger$ defines the eigensystem of A as $\{D, U\}$, where its eigenvalues are represented as the matrix $D(\vec{q}) \equiv \text{diag}(\lambda_1 \cdots \lambda_n)$, with $\lambda_k(\vec{q})$, while its collection of eigenvectors $V_k(\vec{q})$ is organized as the columns of $U(\vec{q})$ as defined in Equation (1). In order to obtain $\{D, U\}$ at a given point $\vec{q}_* \neq \vec{q}_0$, a path between the two points is drawn, $\vec{q}(t)$, as a function of a single parameter t . This new parametrization has no physical meaning, and it is unrelated to how A is parametrized over $\vec{q}$. This relation is introduced to reduce the number of degrees of freedom (d.o.f.) from p parameters to a single one, t . The path is then divided into smaller steps $\delta t \ll 1$ in a total of $N = 1/\delta t$. Without any loss of generality, one might consider a straight line,

$$\vec{q}(t) = \vec{q}_0 + \left(\frac{t - t_i}{t_f - t_i} \right) (\vec{q}_* - \vec{q}_0) \quad (11)$$

as long as $\vec{q}(t)$ never leaves the domain of $A(\vec{q})$. When this is not feasible, Equation (11) can be generalized by a series of line segments or a curve of any kind. Nevertheless, the reader should keep in mind that the resulting $\{D, U\}$ are independent of the taken path, so the curve $\vec{q}(t)$ should be as simple as possible. From this point forward, consider that Equation (11) is enough to define $\vec{q}(t)$. In this case, a short step δt leads to a step $\delta \vec{q}$, with

$$\delta \vec{q} = (\vec{q}_* - \vec{q}_0) \delta t \quad (12)$$

with

$$\vec{q}_{i+1} = \vec{q}_i + \delta \vec{q} \quad (13)$$

being the discrete representation of the chosen path. At any given step i , the next step will lead to $A_{i+1} = A(\vec{q}_{i+1}) = A(\vec{q}_i + \delta \vec{q})$. When diagonalized, the resulting eigenvalue set $\{\lambda_k^{(i+1)}\}$ should have the form

$$\lambda_k^{(i+1)} = \lambda_k^{(i)} + \delta \lambda_k. \quad (14)$$

This relation can be used as a tool to correct for the ordering of $\{\lambda_k^{(i+1)}\}$ by defining the quantity

$$\Delta = \sum_{k=0}^n \left| \lambda_k^{(i+1)} - \lambda_k^{(i)} \right| \quad (15)$$

it is possible to search for the correct ordering of $\lambda^{(i+1)}$ among all possible permutations. Given that δt is small enough, the relation $\Delta < (n \max\{\lambda_k\})$ can only be true if the ordering of $\{\lambda_k^{(i+1)}\}$ matches the previous one for $\lambda_k^{(i)}$. This is performed by simple inspection, placing all permutations of $\{\lambda_k^{(i+1)}\}$ in the definition of Δ and choosing the smallest one. Once the correct permutation is known, both $\{\lambda_k^{(i+1)}\}$ and $\{V_k^{(i+1)}\}$ can be reordered and stored in $\{D_{i+1}, U_{i+1}\}$. This procedure is repeated until the endpoint is reached, leading to the desired $\{D_*, U_*\}$.

The combination of SDS and Jacobi's diagonalization will be referred to as Sequential Jacobi Diagonalization (SJD). Although the SDS can be used with any diagonalization method, it is worth noting that Jacobi's is the most suitable one for neutrino physics since it offers the possibility of evaluating both D and U at the same time with precision ε , predefined only by the stopping condition. In fact, since no other information is kept from one point to the next, besides the ordering, there are no cumulative numerical errors involved. In other words, the only errors affecting $\{D_*, U_*\}$ are those coming from the last diagonalization, at the point $\vec{q}_*$ (see Appendix A for an in-depth discussion about precision).

A few remarks are in order. A major advantage of the SDS is that all the diagonalizations over a path can be computed concurrently, with the reordering performed afterwards,

in a serialized fashion. Also, if the intention is to map the eigensystem over a volume of parameter space, finding a way to run over such space in a continuous manner becomes a trivial task. Yet, this method is not without its limitations. The SDS relies on the premise that there is a reference order. As a consequence, the eigenvalues have to be non-degenerate to begin with. Not only that, they also have to be different enough so Equation (15) is applicable. It could be the case, however, that some particular parametrization causes two or more eigenvalues to cross each other, becoming degenerate at that point. This kind of ambiguity can be solved by adopting a higher order discriminant, such as comparing λ_k^{i+1} with λ_k^{i-1} , which is equivalent to comparing the derivatives of $d\lambda_k/dt$.

5. Neutrino Physics Application

This section offers an example application of SJD in neutrino physics. The goal is to obtain the mixing (oscillation) parameters, defined by the PMNS parametrization, as a function of the matter background in the MSW effect. (see refs. [10,18] for a modern review). In it, the presence of an interacting medium shifts the energy levels of the Hamiltonian, which in turn leads to a new set of effective values for the neutrino mixing parameters, either enhancing or suppressing the oscillation pattern, depending on the matter profile along the propagation path. In what follows, the MSW effect is briefly reviewed, and the usage of the SJD is illustrated for a 3-neutrino, SM case.

As the mass-flavor mixing model states [18], a three-neutrino system can be represented by a free Hamiltonian which is diagonal when expressed in the mass basis, $H_m = (\Delta m_{21}^2/2p) \times \text{diag}(0, 1, \alpha)$, with $\alpha = \Delta m_{31}^2/\Delta m_{21}^2$ and the Δm_{ij}^2 as the squared-mass differences between the neutrino mass-states. The unitary mixing matrix U takes the diagonal H_m to the flavor basis via a similarity transformation,

$$H_f = U H_m U^\dagger, \quad (16)$$

with U being the result of three real rotations (with Euler angles θ_{12} , θ_{23} and θ_{13}) and at least one complex phase δ_{CP} . In the presence of an interacting background, represented by a potential matrix V , the total Hamiltonian of the system becomes

$$\tilde{H}_f = U H_m U^\dagger + V, \quad (17)$$

where the $\sim$ sign represents non-vacuum values, with V being a general real matrix, encoding how each neutrino flavor interacts with the medium. The physical observables are those related to the neutrino oscillation pattern, namely the oscillation length and amplitudes, given by the eigenvalues and eigenvectors of $\tilde{H}_f$, respectively. Let $\tilde{U}$ be the diagonalizing transformation that realizes the following,

$$\begin{aligned} \tilde{H}_m &= \tilde{U}^\dagger \tilde{H}_f \tilde{U} \\ &= \tilde{U}^\dagger (U H_m U^\dagger + V) \tilde{U} \end{aligned} \quad (18)$$

where $\tilde{H}_m \rightarrow H_m$, and $\tilde{U} \rightarrow U$, when $V \rightarrow 0$. Equation (18) is equivalent to Equation (4), with $A = \tilde{H}_f$ and $D = \tilde{H}_m$. This single realization evokes the motivation behind this study, since while Equation (4) is just the starting point of a diagonalization tool, Equation (18) has actual meaning in neutrino physics.

Consider $\tilde{H}_m = (\Delta m_{21}^2/2p) \times \text{diag}(\lambda_1, \lambda_2, \lambda_3)$ as the diagonal form of $\tilde{H}_f$, with λ_k as the relevant factors of its eigenvalues. From the elements of the diagonalizing transformation of U (or $\tilde{U}$), it is possible to define three Mixing Amplitudes,

$$\sin^2 2\theta_{12} = 4 \frac{|U_{e1}|^2 |U_{e2}|^2}{(1 - |U_{e3}|^2)^2}, \quad \sin^2 2\theta_{23} = 4 \frac{|U_{\mu 1}|^2 |U_{\tau 2}|^2}{(1 - |U_{e3}|^2)^2},$$

and

$$\sin^2 2\theta_{13} = 4|U_{e3}|^2 (|U_{\mu 3}|^2 + |U_{\tau 3}|^2). \quad (19)$$

representing how the mass-eigenstates are mixed into the flavor states. This notation corresponds to the PMNS parametrization [10]. It is also possible to isolate the effects of CP-violation, represented by the Jarlskog invariant,

$$J_{\text{CP}} = \text{Im} \left\{ U_{\mu 3} U_{\mu 2}^* U_{e2} U_{e3}^* \right\} \quad (20)$$

where J_{CP} represents the difference between neutrino and antineutrino oscillation.

The angles in Equation (19) and the Jar Equation (20) will lead to different values, depending on the elements of the potential V . As an example, in non-standard interaction searches, the elements of V can be either independent of each other or given by an underlying model. In the case of an ordinary-matter background, however, the potential can be as simple as $V = (\Delta m_{21}^2)/2p \times \text{diag}(a, 0, 0)$, with $a = 2pV_{cc}/\Delta m_{21}^2$ (more on V_{cc} in a moment). Regardless of the model, be it SM or BSM, SJD can be used to obtain the behavior of $\tilde{H}_f$'s eigensystem as a function of a specific model parameter or even the complete set of V elements, which would be model-independent. In the case of a constant and uniform background, these definitions are enough to completely define the system. When this is not the case, it becomes necessary to also know how $\tilde{U}$ varies along the neutrino's trajectory x , i.e.,

$$\frac{d\tilde{U}_{ij}}{dx} = \sum_{k\ell} \frac{dU_{ij}}{dV_{k\ell}} \frac{dV_{k\ell}}{dx}. \quad (21)$$

This is the scenario where the SJD can provide a sizable improvement over other methods. The values of $d\tilde{U}_{ij}/dx$ can be obtained prior to a full model analysis since they should be recalculated less frequently, if ever conducted twice in a single study. Monte Carlo productions, as well as model fitting, can make use of a LUT instead of performing thousands of diagonalizations at every step. The more demanding simulations for the next generation of neutrino detectors, such as DUNE [7], should benefit from this approach.


To take a concrete example, we can appreciate an application using only Standard Model physics, for which there are analytical solutions. In this case, the relevant matter potential is $V = \text{diag}(V_{cc}, 0, 0)$, with V_{cc} being the charged current potential between electrons in the medium and the electron-(anti)neutrino. For neutral baryonic matter, $V_{cc} = \sqrt{2}G_F n_e$, where n_e is the background's electron density, and G_F is Fermi's constant. Since global phases do not influence the final oscillation probabilities, we can place $\Delta m_{21}^2/2p$ in evidence, writing $V = \text{diag}(a, 0, 0)$, with $a = 2pV_{cc}/\Delta m_{21}^2$. The parameter a encodes all the background description such as density, interaction strength and uniformity. By using SJD, it is possible to obtain all the relevant observables and still be agnostic with respect to the background properties, which can be added at a later point of the computation.

The total Hamiltonian to be diagonalized is $\tilde{H}_m(a)$, where $a > 0$ means that both the neutrinos and the background are of the same nature, i.e., either both matter or both antimatter, while $a < 0$ represents the matter/antimatter combination. Using the vacuum values on Table 1, it is possible to obtain two distinct values for α , $\alpha_{\text{NH}} = 32.4$ and $\alpha_{\text{IH}} = -31.9$, corresponding to Normal Hierarchy (NH) and Inverted Hierarchy (IH), respectively. All results that follow will show four distinct cases: $\pm a$ and NH/IH.


Table 1. Best fit values, 1σ and 3σ ranges for the global fit of all relevant neutrino oscillation data [19]. Here, the notation used by the original reference is changed in favor of one that best suits this work.

Parameter	Best Fit	1σ Range	3σ Range
$\Delta m_{21}^2/10^{-5}$	7.37	7.21–7.54	6.93–7.97
$\sin^2 \theta_{12}/10^{-1}$	2.97	2.81–3.14	2.50–3.54
Normal Hierarchy (NH)			
$+\Delta m_{31}^2/10^{-3}$	2.39	2.35–2.43	2.27–2.51
$\sin^2 \theta_{13}/10^{-2}$	2.14	2.05–2.25	1.85–2.46
$\sin^2 \theta_{23}/10^{-1}$	4.37	4.17–4.70	3.79–6.16
δ_{CP}/π	1.35	1.13–1.64	0–2
Inverted Hierarchy (IH)			
$-\Delta m_{31}^2/10^{-3}$	2.35	2.31–2.40	2.23–2.48
$\sin^2 \theta_{13}/10^{-2}$	2.18	2.06–2.27	1.86–2.48
$\sin^2 \theta_{23}/10^{-1}$	5.69	4.28–4.91	3.83–6.37
δ_{CP}/π	1.32	1.07–1.67	0–2

The analytical solutions for the Mixing Amplitudes [11,12] are compared to the SJD, being in agreement up to the chosen precision ($\varepsilon = 10^{-14}$). Figure 2a shows the mixing amplitudes $\sin^2 2\theta_{12}$ and $\sin^2 2\theta_{13}$ as a function of $|a|$ ($\sin^2 2\theta_{23}$ is not shown since it is indistinguishable from 1 in this scale). It is possible to observe the resonant MSW effect, related to the two mass-scales. The lower resonance a_r affects $\sin^2 2\theta_{12}$, while the higher one a_R affects $\sin^2 2\theta_{13}$ and $\sin^2 2\theta_{23}$. The latter is not shown on the plots since it would be indistinguishable from the unit due to its large vacuum value.



(a) Mixing Amplitudes, as a function of $|a|$.



(b) Jarlskog Invariant, as a function of $|a|$.

Figure 2. (a) The matter-enhanced values of $\sin^2 2\theta_{12}$ and $\sin^2 2\theta_{13}$, with red and blue for $a > 0$ and magenta and cyan for $a < 0$, respectively, with NH at the top plot and IH at the bottom. (b) The matter enhanced values of J_{CP} , as defined on Equation (20). At the top is the NH scenario, with blue for $a > 0$ and red for $a < 0$, while at the bottom is the IH case, with magenta for $a > 0$ and orange for $a < 0$.

In Figure 3, we observe the eigenvalues of H_m for NH and IH. The vacuum eigenvalues are $\lambda_1 = 0$, $\lambda_2 = 1$ and $\lambda_3 = \alpha$, and it is possible to see that resonances a_r and a_R represent the points where the eigenvalues change asymptotes.

Finally, Figure 2b shows how the Jarlskog invariant is affected by the background. Regardless of the hierarchy case, a_r represents a resonant minimum for $a > 0$, and $a > a_R$

will always lead to $J_{CP} = 0$, meaning that neutrinos and antineutrinos would behave the same. It is worth noting that the existence of Charge-Parity (CP)-violation in neutrino oscillations is not confirmed, and the results shown here only consider the best fit values for δ_{CP} , which are still compatible with zero. No matter what the true δ_{CP} is, it affects J_{CP} with all the previous observables remaining unchanged.




Figure 3. Eigenvalues of $\tilde{H}_m$, in Equation (18), for NH (a) and IH (b). Each individual eigenvalue λ_i is drawn as a function of the dimensionless background descriptor a , representing the MSW effect in ordinary matter. The differences $\Delta\lambda_{ij}$ are also shown, since they correlate to the oscillation length. All vacuum values are taken from Table 1.

6. Conclusions

The Sequential Jacobi Diagonalization, or SJD, proposed in this work combines a heuristic procedure with a well-established numerical method in order to satisfy the computation requirements for neutrino physics application. In this field, computational resources become a bottleneck whenever BSM hypotheses are being tested. In more general terms, given the description of a Hermitian system, modeled over a particular set of parameters, this method allows for the study of how the eigenvalues and eigenvectors are related to these parameters.

Funding: This research was funded by *Conselho Nacional de Desenvolvimento Científico e Tecnológico*, CNPq, grant No. 477588/2013-1, and *Fundação de Amparo a Pesquisa do Estado de Minas Gerais*, FAPEMIG, with grants No. APQ-01439-14 and APQ-01249-24, which also funded the APC with grant No. APQ-00544-23. This work was partially developed during a fellowship period granted by *Coordenação de Aperfeiçoamento de Pessoal de Nível Superior*, CAPES, grant No. 88881.121149/2016-01.

Data Availability Statement: Data is contained within the article.

Acknowledgments: The author wishes to thank Alberto M. Gago (ORCID 0000-0002-0019-9692) for his insightful and encouraging comments on this work.

Conflicts of Interest: The author declares no conflict of interest.

Abbreviations

The following abbreviations are used in this text:

BSM	Beyond Standard Model
CJM	Cyclic Jacobi Method
CP	Charge-Parity
d.o.f.	degrees of freedom
DUNE	Deep Underground Neutrino Experiment
LUT	look-up table
IH	Inverted Hierarchy
MSW	Mikheyev–Smirnov–Wolfenstein
NH	Normal Hierarchy
NSNI	Non-Standard Neutrino Interactions
PMNS	Pontecorvo–Maki–Nakagawa–Sakata
SDS	Sequential Diagonalization Strategy
SJD	Sequential Jacobi Diagonalization
SM	Standard Model

Appendix A. Precision and Efficiency

In order to evaluate stability, convergence and precision, I applied the Jacobi diagonalization to a sample of 10^6 random Hermitian matrices, with the real and imaginary parts of each element constrained to $[-1, +1]$. This sample is representative since any matrix can be normalized by its largest element in order to fit in this range. After reaching the stop condition $d^2 \leq \varepsilon^2$, as defined on Equation (2), the resulting diagonal form is rotated back with the obtained $U = S^\dagger$ and compared component-wise with the original matrix. The largest difference is found to be always smaller than the target precision ε , meaning that $-\log_{10} \varepsilon$ is a good indicator of the number of significant figures achieved in the solutions. For each sample, ε is varied from 10^{-2} to 10^{-16} , where the average number of rotations S_k is recorded. This virtual experiment is repeated from $n = 3$ to $n = 10$ matrices, and the results are shown in Figure A1.




Figure A1. Average number of sweeps before converging to a diagonal, with given precision ε (shown as $-\log \varepsilon$). One sweep is defined as the application of one full complex rotation for each off-diagonal element, $n(n-1)/2$. Each line represents a dimensionality from $n = 3$ to 10. This average is obtained from a random sample of 10^6 Hermitian matrices with real and imaginary parts limited to the interval $[-1, +1]$. It is possible to obtain diagonalizations with the largest off-diagonal element bounded as $<10^{-14}$ by employing an average of 3 to 4 sweeps, for matrices up to 10×10 .

Since the method targets the largest elements, not all in sequence (contrary to its cyclic variant), the average “sweep” is defined as the ratio between the number of complex

rotations (two real rotations from Equations (7) and (9) and the number of off-diagonal elements $n(n-1)/2$. This ratio is strictly larger than 1, regardless of the dimensionality, since a general matrix requires at least one complex rotation for each off-diagonal element. The actual number of rotations goes with $\mathcal{O}(n^2)$. By establishing the $d^2 \leq \varepsilon^2$ limit from Equation (2) as stopping criterion, there is a possibility that some elements might underflow if the required ε is too close to machine precision. Indeed, this is observed when requiring $\varepsilon \leq 10^{-15}$, using 64-bit floating-point variables (which can represent a maximum of 15 significant figures). Stability and convergence are observed with $\varepsilon \geq 10^{-14}$, which is the largest precision shown in Figure A1. In this limit, numerical diagonalization is achieved with an average of between 3 and 4 sweeps. This average holds even for matrices as large as 10×10 . Most physical applications would realistically require far less precision than the 10^{-14} tested, which translates to a less demanding process. Table A1 shows the average number of sweeps, the standard deviation and how many sweeps are needed to diagonalize 99% of each sample. Even in the most demanding case, with $n = 30$, five significant figures' precision can be obtained with a maximum of 3.2 sweeps. Also, the standard deviation around this average becomes narrower as n increases. Both Figure A1 and Table A1 show evidence of a possible limit, or at least a log-like growth, in the number of sweeps as a function of n . This cannot be verified by employing only numerical analysis, so no further statements will be made on this observation. It can be said, however, that the expected number of real rotations is $3n^2$, as a thumb rule. As a final remark, quadratic convergence (precision = sweeps²) was observed for all tested dimensionalities, as suggested by the literature [6].

Table A1. Average number of sweeps required for convergence a diagonal, with precision $\varepsilon = 10^{-5}$, from the a 10^6 random matrices sample. The first column (Avg. Sweeps) corresponds to a cut from Figure A1 at $-\log \varepsilon = 5$. The second one (Std. dev.) shows the standard deviation from each sample. The rightmost column shows how many sweeps were needed for 99% of each sample to reach the stopping condition. The two bottom rows show extra information not present in Figure A1, for $n = 20$ and $n = 30$. Even with such large matrices, less than 1% of the matrices reacquired more than 3.3 sweeps.

n	Avg. Sweeps	Std. Dev.	99% Less Than
3	2.30	0.19	2.7
4	2.51	0.19	3.0
5	2.66	0.13	3.1
6	2.74	0.10	3.1
7	2.81	0.10	3.1
8	2.85	0.09	3.2
9	2.88	0.08	3.2
10	2.92	0.08	3.2
20	3.07	0.05	3.2
30	3.15	0.04	3.3

Appendix B. Numerical vs. Analytical

In this section, a random example with an analytic solution is analyzed. The goal is to validate the numerical methods proposed in this work. A particular $n = 3$ case with a known analytic solution is used to exemplify the validity of the method. Starting with two Hermitian matrices, A and B , given by

$$A = \begin{bmatrix} 3 & i & 0 \\ -i & -2 & i \\ 0 & -i & 1 \end{bmatrix}, \quad (\text{A1})$$

and

$$B = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 2 & 0 \\ 0 & 0 & 3 \end{bmatrix}, \quad (\text{A2})$$

a linear parametrization is defined as

$$H(x) = A x + B. \quad (\text{A3})$$

These were chosen among several tests for no particular reason other than to provide a good example. The eigenvalues of B are not only explicit, since B is diagonal by definition, but also, their ordering is well determined. The eigenvalues of A are obtained by solving its order-3 characteristic polynomial, leading to

$$\lambda_1^A = \frac{1}{3} [2 - \operatorname{Re}\{\sigma (1 - i\sqrt{3})\}], \quad (\text{A4})$$

$$\lambda_2^A = \frac{1}{3} [2 + \operatorname{Re}\{\sigma\}], \quad (\text{A5})$$

and

$$\lambda_3^A = \frac{1}{3} [2 - \operatorname{Re}\{\sigma (1 + i\sqrt{3})\}], \quad (\text{A6})$$

where $\sigma = (-64 + 3i\sqrt{1281})^{1/3}$ (the first complex root). Their numbering is reflecting their relative positioning on the number line, $\lambda_1^A \geq \lambda_2^A \geq \lambda_3^A$, not parametrization ordering. For the sake of this analysis, all numerical values are quoted with 10^{-5} precision even when using analytical formulas. The eigenvalues of A are $\lambda_1^A = -2.47090$, $\lambda_2^A = 1.26071$ and $\lambda_3^A = 3.21018$, as defined by Equations (A4)–(A6).

One wishes to study the parametrized eigensystem of H , represented by $\{D, U\}$, as a function of x . By the definition in Equation (A3), $D(x=0) = B$; therefore, $U(x=0) = 1$. In other words, at $x = 0$, the eigenvalues of H are not only the same as those of B , but they follow the same order. It is also possible to infer just from Equation (A3) the behavior of H when $x \rightarrow \pm\infty$, since $A x$ becomes the dominant term, and the eigenvalues of H assume the form of $x \lambda_k^A$. This means that the eigensystem of H has an asymptotic behavior, and, for instance, one might be tempted to write $\lambda_1^H(x) = x \lambda_1^A + \lambda_1^B$ in order to describe $\lambda_1^H(x)$ asymptote. However, there is no explicit information stating which $x \lambda_k^A$ corresponds to which λ_j^B . Unless H is diagonalized, the true correspondence between the eigenvalues near zero and its value elsewhere is not clear yet, being implicitly determined by the parametrization. Additionally, the same λ^H can have different asymptotes for each limit.

By applying the Sequential Jacobi Diagonalization, described in Sections 3 and 4, a numerical representation of $D(x)$ and $U(x)$ can be calculated for a range of x around the origin. Figure A2a shows the functions $\lambda_j^H(x)$ for a $|x| \geq 3$, which contain all of this system's features. Their behavior is analogous to that of trains changing tracks. There are three asymptotes, of the form

$$f_k = x \lambda_k^A + a_k, \quad k = 1, 2, 3, \quad (\text{A7})$$

where k indicates a particular eigenvalue of A , as defined in Equations (A4)–(A6), and $a_1 = 2.04297$, $a_2 = 2.89648$ and $a_3 = 1.05859$ are constants, numerically obtained by the method. Each $\lambda_j^H(x)$ follows these asymptotes, changing allegiance every time they intersect. There are also three intersection points, in increasing order of $x_{12}^* \geq x_{13}^* \geq x_{23}^*$, with numerical values shown in Table A2. The intersections are obtained by considering the λ_k curves as hyperbolas, where their point of closest approach is where their derivatives are equal, which can be seen on Figure A2b. Finally, it is possible to examine the three eigenvectors by taking their real spherical representation, i.e.,

$$\theta_k = \cos^{-1}(U_{3k}^* U_{3k}) \quad \text{and} \quad \phi_k = \tan^{-1}(U_{2k}^* U_{2k}), \quad (\text{A8})$$

where it is implicit that $\sum_{jk} U_{jk}^* U_{jk} = 1$. Although this projection is a limited representation, where $V_k = (\theta_k, \phi_k)$, with $\theta/(\pi/2) \in [0, 1]$ and $\phi/(\pi/4) \in [0, 1]$, it is enough to observe their limiting behavior, as shown on Figure A3.




Figure A2. (a) Eigenvalues of $H(x)$, as defined by Equations (A1)–(A3). The continuous lines shows λ_1^H (yellow), λ_2^H (green) and λ_3^H (blue), as a function of the free parameter x . The dashed lines indicate the asymptotes Equation (A7) and their intersections (Table A2). (b) Derivative of the Eigenvalues of $H(x)$ with respect to x , as defined by Equations (A1)–(A3). The continuous lines shows $d\lambda_1^H/dx$ (yellow), $d\lambda_2^H/dx$ (green) and $d\lambda_3^H/dx$ (blue), as a function of the free parameter x . The dashed lines indicate the asymptotes (Equation (A7)), which correspond to the eigenvalues of A listed in Equations (A4)–(A6). Their intersections correspond to those of the asymptotes defined by Equation (A7) and their numerical values displayed on Table A2.

Table A2. Intersections between asymptotes. These are the values of x where each eigenvalue changes allegiance to one of the system's asymptotes.

Intersection Between	x_{jk}^*	y_{jk}^*
f_1 and f_2	−0.22873	2.60814
f_1 and f_3	0.17327	1.61484
f_2 and f_3	0.94276	4.08503

The initial position of each vector is $V_1 = (0, 1)$, $V_2 = (1, 1)$ and $V_3 = (0, 0)$, when $x = 0$, revolving around the unit sphere for other values of x . When we compare each eigenvector with its corresponding eigenvalue on Figure A2a, it is possible to correlate their behavior, for instance, how $V_1 \rightarrow V_3$ as $x \rightarrow \pm\infty$ or how although V_2 crosses the other eigenvectors several times, the eigenvalues are never degenerate.

In conclusion, all the values and functions obtained in this example match their analytical counterparts up to 10^{-14} , which is the precision set for the method's precision ϵ , while the expected precision for evaluating the analytical solutions is 10^{-15} (using 64-bit floats).




Figure A3. Spherical representation of $H(x)$'s eigenvectors. The eigenvectors V_1 (yellow), V_2 (green) and V_3 (blue). The starting point for each vector is $V_1 = (0,1)$, $V_2 = (1,1)$ and $V_3 = (0,0)$. The dots represent equal steps in x and help demonstrate the asymptotic behavior, when the density of points increases (meaning $x \rightarrow \infty$ (denoted by the three stars). The point $A = (0.25721, 0.94092)$ corresponds to either V_1 when $x \rightarrow -\infty$ or V_3 when $x \rightarrow +\infty$. Similarly, $B = (0.96496, 0.82418)$ is the limit of V_1 when $x \rightarrow +\infty$ or V_3 when $x \rightarrow -\infty$. And finally, $C = (0.31152, 0.18597)$ is the convergence of V_2 for both $x \rightarrow \pm\infty$.

References

- Wolfenstein, L. Neutrino oscillations in matter. *Phys. Rev. D* **1978**, *17*, 2369–2374. [\[CrossRef\]](#)
- Mikheev, S.P.; Smirnov, A.Y. Resonance Amplification of Oscillations in Matter and Spectroscopy of Solar Neutrinos. *Sov. J. Nucl. Phys.* **1985**, *42*, 913–917.
- Hubbard, J.; Schleicher, D.; Sutherland, S. How to find all roots of complex polynomials by Newton's method. *Invent. Math.* **2001**, *146*, 1–33. [\[CrossRef\]](#)
- Jacobi, C.G.J. Über ein leichtes Verfahren, die in der Theorie der Säkularstörungen vorkommenden Gleichungen numerisch aufzulösen. *J. Pure Appl. Math.* **1846**, 51–94. [\[CrossRef\]](#)
- Eberlein, P.J. A Jacobi-Like Method for the Automatic Computation of Eigenvalues and Eigenvectors of an Arbitrary Matrix. *J. Soc. Ind. Appl. Math.* **1962**, *10*, 74–88. [\[CrossRef\]](#)
- Hansen, E.R. On Cyclic Jacobi Methods. *J. Soc. Ind. Appl. Math.* **1963**, *11*, 448–459. [\[CrossRef\]](#)
- Abi, B.; Acciarri, R.; Acero, M.A.; Adamov, G.; Adams, D.; Adinolfi, M.; Ahmad, Z.; Ahmed, J.; Alion, T.; Monsalve, S.A.; et al. Long-baseline neutrino oscillation physics potential of the DUNE experiment. *Eur. Phys. J. C* **2020**, *80*, 1–34. [\[CrossRef\]](#)
- Blennow, M.; Smirnov, A.Y. Neutrino Propagation in Matter. *Adv. High Energy Phys.* **2013**, *2013*, 33. [\[CrossRef\]](#)
- Maltoni, M.; Smirnov, A.Y. Solar neutrinos and neutrino physics. *Eur. Phys. J. A* **2016**, *52*, 87. [\[CrossRef\]](#)
- Smirnov, A.Y. Solar neutrinos: Oscillations or No-oscillations? *arXiv* **2016**, arXiv:1609.02386.
- Zaglauer, H.W.; Schwarzer, K.H. The mixing angles in matter for three generations of neutrinos and the MSW mechanism. *Z. Phys. C Part. Fields* **1988**, *40*, 273–282. [\[CrossRef\]](#)
- Kimura, K.; Takamura, A.; Yokomakura, H. Exact formulas and simple CP dependence of neutrino oscillation probabilities in matter with constant density. *Phys. Rev. D* **2002**, *66*, 073005. [\[CrossRef\]](#)
- Denton, P.B.; Parke, J.; Zhang, X. Rotations versus perturbative expansions for calculating neutrino oscillation probabilities in matter. *Phys. Rev. D* **2018**, *98*, 033001. [\[CrossRef\]](#)
- Press, W.H.; Teukolsky, S.A.; Vetterling, W.T.; Flannery, B.P. *Numerical Recipes in C*, 2nd ed.; Cambridge University Press: Cambridge, MA, USA, 1992.
- Golub, G.H.; van der Vorst, H.A. Eigenvalue computation in the 20th century. *J. Comput. Appl. Math.* **2000**, *123*, 35–65. [\[CrossRef\]](#)
- Forsythe, G.E.; Henrici, P. The cyclic Jacobi method for computing the principal values of a complex matrix. *Trans. Am. Math. Soc.* **1960**, *94*, 1–23. [\[CrossRef\]](#)
- Nazareth, L. On the convergence of the cyclic Jacobi method. *Linear Algebra Appl.* **1975**, *12*, 151–164. [\[CrossRef\]](#)

18. Bilenky, S.M. Neutrino oscillations: Brief history and present status. *arXiv* **2014**, arXiv:1408.2864.
19. Capozzi, F.; Lisi, E.; Marrone, A.; Montanino, D.; Palazzo, A. Neutrino masses and mixings: Status of known and unknown 3ν parameters. *Nucl. Phys. B* **2016**, *908*, 218–234. [[CrossRef](#)]

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.