

Implementing Recycling Methods for Linear Systems in Python with an Application to Multiple Objective Optimization

Ainara Garcia

College of Engineering, Computing & Applied Sciences
Clemson University
Clemson, SC, USA
ainarag@clemson.edu

Sihong Xie

Thrust of AI
HKUST-GZ
Guangzhou, China
xiesihong1@gmail.com

Arielle Carr

Computer Science & Engineering
Lehigh University
Bethlehem, PA, USA
arg318@lehigh.edu

Abstract—Sequences of linear systems arise in the predictor-corrector method when computing the Pareto front for multi-objective optimization. Rather than discarding information generated when solving one system, it may be advantageous to recycle information for subsequent systems. To accomplish this, we seek to reduce the overall cost of computation when solving linear systems using common recycling methods. In this work, we assessed the performance of recycling minimum residual (RMINRES) method along with a map between coefficient matrices. For these methods to be fully integrated into the software used in Enouen et al. (2022), there must be working version of each in both Python and PyTorch. Herein, we discuss the challenges we encountered and solutions undertaken (and some ongoing) when computing efficient Python implementations of these recycling strategies. The goal of this project was to implement RMINRES in Python and PyTorch and add it to the established Pareto front code to reduce computational cost. Additionally, we wanted to implement the sparse approximate maps code in Python and PyTorch, so that it can be parallelized in future work.

Index Terms—MINRES, Sparse Approximate Maps, Python, PyTorch, Multiobjective Optimization, Recycling

I. INTRODUCTION

Issues of fairness often arise in graphical neural networks used for misinformation detection. Accurately reducing unfairness when satisfying more than one constraint can come at a significant computational cost. Multi-objective optimization (MOO) methods allow us to explore different trade-offs by generating a set of solutions, known as the Pareto front [1]. Traditional first-order MOO methods, such as multi-gradient descent (MGD) [2], are computationally expensive on large graphs due to the gradient computation at each iteration, and may require many iterations to generate a Pareto front. Using the predictor-corrector method introduced in [3], and iterative methods like MINRES [4] and Conjugate Gradient (CG) [5], the linear systems arising in the predictor step can be solved efficiently.

We specifically focus on the iterative method MINRES [4] and recycling MINRES (RMINRES) [6] for solving the linear systems arising in the predictor step. In [3], the authors found that the computational cost associated with solving this sequence of systems can be reduced by dramatically limiting the

maximum number of iterations when using MINRES, while still achieving an accurate Pareto front [3]. In the current work, we are interested in assessing the performance of recycling strategies across the sequence of linear systems to further reduce the total number of iterations. In this project, we also consider a mapping strategy between the coefficient matrices that may allow us to reuse previous solutions. Specifically, we explore the potential for a parallel version of the preconditioner update reported in [7] to further reduce computational costs associated with many preconditioners for long sequences of systems arising in the predictor step. In particular, we seek to implement both RMINRES and the preconditioner updates in Python and PyTorch [8], using the respective language and libraries when computing the Pareto front's predictor-corrector methods. We provide preliminary results related to our work, noting that the implementation is still in progress. Aiming to help other groups working on the same topic, we also detail the complications encountered.

II. PRELIMINARIES

A. Predictor Corrector Method

Given a point of a function, the predictor-corrector method allows us to approximate the value of that function at a nearby point. This method consists of two steps (see Figure 1). First, in the predictor step, we determine an approximate direction to a neighboring point, and then move along that direction based on a predetermined, fixed step size. Then, we refine this initial approximation in the corrector step.

To determine the direction in which to move in the predictor step, we solve for the vector v in Equation 1,

$$H(x_0^*)v = \nabla f(x_0^*)\beta \quad (1)$$

where x_0^* is a starting Pareto optimal point, and $H(x_0^*)$ and $\nabla f(x_0^*)$ are the Hessian and Jacobian of our loss functions at x_0^* , respectively; and β is a weighting vector. Intuitively, this approach works by finding a suitable direction to move along the tangent plane of f at x_0^* as depicted in Figure 1. While more information on this method can be found in [3]; this report focuses on the reducing costs associated with




Fig. 1: The predictor step (yellow) and corrector step (green) given an initial point (red).

solving the sequence of systems of the form (1) that arise in the predictor step. Directly solving the linear systems of the form (1) requires $O(n^3)$ time, which is generally expensive, even for modest matrix dimensions, n . Furthermore, in this application, H is not explicitly known, so a matrix-free method, such as MINRES is required [3]. Thus, we consider solvers whereby the solution is approximated iteratively (i.e., iterative solvers). Since the coefficient matrix H is symmetric, we use a symmetric solver such as MINRES, which requires only a linear operator to compute a matrix-vector (the main computational cost in such iterative methods). Instead of starting the computation of MINRES from scratch for each new linear system, the recycling approach allows us to reuse information from previous solutions to accelerate convergence and reduce computational costs by reducing the total number of iterations to solve the sequence of linear systems. To accomplish this, we consider the RMINRES method [6].

III. ITERATIVE METHODS SOLVING LINEAR SYSTEMS

Iterative methods compute an approximate solution to a linear system by repeating a series of calculations to update the initial guess until convergence is achieved. Unlike direct methods, which compute an exact solution, iterative methods progressively refine an approximate solution until the accuracy is below the pre-determined tolerance. The processes includes taking an initial guess, which will be the starting vector for the rest of the iteration process.

The MINRES algorithm [4] is a Krylov iterative method used to solve systems of linear equations when the coefficient matrix is symmetric and possibly indefinite. MINRES was developed as an alternative to other well-known symmetric solvers, such as the CG algorithm [5], which requires the coefficient matrix to be symmetric, positive, and definite. A distinguishing features of MINRES is its ability to handle indefinite matrices. In contrast to positive definite matrices, which have all positive eigenvalues, indefinite matrices might have both positive and negative eigenvalues.

Central to MINRES is the emphasis on minimizing the residual vector norm, $\|r_k\|_2$, at each iteration, where the residual of an approximate solution x_k is defined as $r_k =$




Fig. 2: Normwise difference between each pair of matrices in a sequence of 100 systems extracted from [1]. Note that the y-axis shows $\log(\|A_k - A_{k-1}\|_F)$.

$Ax_k - b$. The Krylov subspace of dimension k is defined as $\mathcal{K}^k(A, r_0) = \text{span}\{r_0, Ar_0, \dots, A^{k-1}r_0\}$, where A is our coefficient matrix from the linear system $Ax = b$, $r_0 = b - Ax_0$ is the initial residual, and x_0 is the initial guess. $\mathcal{K}^k(A, r_0)$ is expanded by one vector with each iteration of MINRES. The k^{th} iteration k of MINRES computes $z_k \in \mathcal{K}^k(A, r_0)$ and updates x_k as $x_k = x_0 + z_k$ such that the residual norm $\|r_k\|_2 = \|Ax_k - b\|_2$ is minimized. This continues until $\|r_k\|_2$ is within a predefined tolerance. More information on MINRES and other Krylov methods can be found in [12]. This process ensures that convergence can be achieved in n iterations, but ideally an accurate-enough solution can be achieved in $m \ll n$ iterations. To further reduce the number of iterations even more, we can recycle some of the Krylov vectors generated for a previous system to start the next call to the solver for a subsequent system rather than starting with just one vector. This new set of vectors is stored in the columns of a matrix we refer to as U , called the recycle space [6].

For linear systems that are close to one another, recycling subspaces (i.e., U) can be advantageous to reduce the total number of iterations of the iterative solver across the entire sequence. We provide iteration comparisons in the next section, but refer to Figure 2 to demonstrate how closely related the systems are using 100 matrices extracted from [3]. We plot the norm-wise difference between each pair of matrices $\|A_k - A_{k-1}\|_2$, for $k = 2, 3, \dots, 100$. The dimension of these systems is $n = 1234$. The peaks generally occur at each new iteration of the predictor-corrector method; as noted in [3] and references within. Even in cases of these “larger” peaks, the differences between the matrices are quite small, motivating our conjecture that the subspaces generated by one call to MINRES will be a close approximation to the subspace for the next call (and thus will serve as a good recycle space for RMINRES).

IV. RECYCLING STRATEGIES

The iterative nature of the MINRES algorithm enables it to handle large-scale linear systems that might be computationally expensive to solve using direct methods. By gradually refining the approximate solution with each iteration, MINRES offers an attractive compromise between accuracy and computational efficiency. Many algorithms also use preconditioners

as part their computation to further accelerate convergence when the matrix is ill-conditioned, that is $\|A\|\|A^{-1}\| \gg 1$. However, because of the computational cost associated with computing a good preconditioner, we can instead recycle preconditioners by updating a previous preconditioner to then be reuse for the next system.

Preconditioners are used in iterative methods to speed up the convergence of linear systems when the coefficient matrix is ill-conditioned. By transforming the original system into a better-conditioned one, preconditioners reduce the number of iterations required to reach an accurate solution. More details on preconditioning techniques can be found in [9] and references therein. Preconditioners can enhance the robustness of iterative solvers, making them more tolerant to ill-conditioned or difficult-to-solve problems. They can handle cases where the original system matrix is poorly conditioned or close to singular. However, they can be expensive to compute and so we consider preconditioner updates, specifically the sparse approximate map [7].

The authors in [7] present the central concept behind SAMs: the mapping of one matrix to another for which we have a good preconditioner (i.e., one that facilitates fast convergence of the iterative solver) [7]. For instance, consider a matrix A_k in a sequence of matrices, we want to compute a map N_k to another matrix A_0 with the constraint of an imposed sparsity pattern to limit the cost of computing and applying N_k . In particular, the equation $A_k N_k = A_0$ expresses the exact mapping, where A_k represents the current system matrix, A_0 is the target system matrix, and N_k is the map. The objective is to find an approximate N_k such that it approximates A_0 well enough. Due to the iterative nature of performing SAMs, we can reduce the computational cost of solving linear systems and increase efficiency of our implementation by exploring the potential of a parallel implementation. As shown in Figure 3, the condition numbers (i.e., $\|A\|\|A^{-1}\|$) of our systems is relatively low, so while the use of a preconditioner in this application was ultimately not necessary, we still consider a Python implementation for future use in cases where the matrices are, in fact, ill-conditioned. In section VI-B, we give the explicit mathematical expressions that define a SAM and discuss the corresponding implementation challenges and (proposed) solutions.




Fig. 3: Condition number (y-axis) of matrices from a sequence of 100 systems extracted from [1].

There is often a trade-off between accuracy and the cost of




Fig. 4: Iterations for MINRES and RMINRES for a sequence of 100 systems extracted from [1].

(preconditioned) iterative methods. Iterative methods typically allow control over the accuracy by specifying a convergence criterion (e.g., a tolerance level or maximum iterations allowed). Achieving higher accuracy may require more iterations, leading to increased computational cost. Choosing an appropriate iterative method, setting suitable convergence criteria, and considering the resources available are crucial in striking the right balance between accuracy and cost. Krylov methods, such as MINRES and RMINRES, are particularly useful in our case since these methods do not require the matrix to be explicitly defined, and are often referred to as ‘matrix-free methods’ [10]. In our case, H in 2 is not known explicitly, only its action on a vector.

The goal of this project is to implement RMINRES in Python and add it to the established Pareto front code used in [3] to reduce the computational cost. Additionally, the project aimed to implement the sparse approximate maps code in Python so that in future work, it can be parallelized. We ran initial experiments using a MATLAB implementation of MINRES and RMINRES on the same sequence of linear systems as in Figures 2 and 3. In Figure 4, we compare the iterations for the two methods. We see that RMINRES reduces total iterations from 828 to 796, and we observe the largest reduction in iterations for the linear systems arising earlier in the sequence. These results motivate our implementation of RMINRES in Python/PyTorch as the cost of the Hessian-vector products impose a significant computational computational bottleneck. As the size of the systems grow (i.e., as the number of constraints increase in the optimization), such reductions become critical.

V. IMPLEMENTATION

It is important to have iterative solver methods in an industry-standard computer science language. For the code implemented in this project, the Python language and several libraries were used. NumPy, KryPy, and SciPy [11] packages played a significant role in these implementations. PyTorch is an integral tool in machine learning and parallel computing. It allows for scalable distributed training and performance optimization in research. However, like all tools, it has limitations. Initially, when implementing RMINRES in PyTorch, an error within the implementation occurred. However we discovered

that the sparse tensor data structure and operations required for this project, are not developed enough for the scale of our research. The PyTorch linear algebra solver, `torch.linalg.solve`, often used to find the solution of systems of linear equations, is compatible with float, double, cfloat, and cdouble dtypes [8]. This limitation in the PyTorch operations caused errors in the implementation when solving the incomplete Cholesky factorization of coefficient matrix. This implementation in PyTorch presented many challenges when trying to compute matrix-vector multiplication. Therefore, the alternative choice was to implement the code in Python initially, then integrate it into PyTorch from Python. Python has many packages capable of handling large-scale computing, matrices, linear algebra, and iterative methods.

KryPy, a Krylov subspace methods package [11], includes many of the functions called into our implementations. Among these packages is `krypy.recycling`, which allows for the computation of solutions of linear systems using recycling. Calling `krypy.recycling` allows us to build the U subspace and return it to be implemented in the next call to RMINRES. Many combinations of KryPy functions were attempted before coming to the conclusion that KryPy's recycling functions are limited to the number of iterations allowed, and many times does not converge within those iterations for the computations done in this project. The SciPy library has a package that allows for the computation of sparse linear systems, for example the compressed sparse column format used in this project, as well as functions used to process matrices, like returning a sparse matrices from diagonals and performing linear algebra operations. NumPy is commonly used for scientific computation in Python. While these libraries were important in the development of the code for recycling strategies, their data structures are sometimes not compatible with PyTorch, and therefore require significant modification.

VI. WORKFLOW AND STEPS

As the goal for the project is to integrate these commonly used methods to improve performance of the predictor-corrector method, we began with sequential implementations of our methods. When referencing MINRES, it is important to again emphasize that it begins with an initial vector, while RMINRES begins with many vectors, or a subspace generally defined by vectors from the Krylov space built for the previous linear system in the sequence. This often leads to a reduction in solver iterations compared with standard MINRES (as shown in Figure 4). The reader can refer to [10] for the detailed algorithm for RMINRES.

A. Recycling Minimum Residual (RMINRES)

In our case, several functions are required to implement RMINRES. We also need to calculate the k eigenvalue-eigenvector pairs, referred to as Ritz pairs (see e.g., [12]), that correspond to the smallest eigenvalues of the coefficient matrix. Then, we generate a sparse matrix A and right-hand-side b from test matrices. Algorithm 1 gives the conjugate residual method [10], an equivalent version of MINRES; we

Algorithm 1 Conjugate Residual Method for $H\mathbf{v} = \mathbf{b}$

```

1:  $\mathbf{r}^{(0)} = \mathbf{b} - H\mathbf{v}^{(0)}$ .
2:  $\mathbf{p}^{(0)} = \mathbf{r}^{(0)}$ .
3:  $i = 0$ .
4: while  $i < \text{maxIter}$  and  $\|\mathbf{r}^{(i)}\| > \text{tol}$  do
5:    $\alpha_i = \frac{(H\mathbf{r}^{(i)})^\top \mathbf{r}^{(i)}}{\|H\mathbf{p}^{(i)}\|^2}$ .
6:    $\mathbf{x}^{(i+1)} = \mathbf{x}^{(i)} + \alpha_i \mathbf{p}^{(i)}$ .
7:    $\mathbf{r}^{(i+1)} = \mathbf{r}^{(i)} - \alpha_i H\mathbf{p}^{(i)}$ .
8:    $\beta_i = \frac{(H\mathbf{r}^{(i+1)})^\top \mathbf{r}^{(i+1)}}{(H\mathbf{r}^{(i)})^\top \mathbf{r}^{(i)}}$ .
9:    $\mathbf{p}^{(i+1)} = \mathbf{r}^{(i+1)} + \beta_i \mathbf{p}^{(i)}$ .
10:   $i = i + 1$ .
11: end while
```

use this algorithm for simplicity, but refer the reader to [4] for an in depth discussion of the MINRES algorithm.

RMINRES function requires several initial inputs: A (from $Ax = b$), b (again from $Ax = b$), an initial guess (x_0), a convergence tolerance, the maximum number of iterations, a preconditioner (if available/desired), and the recycle space U . The code checks if a preconditioner and a recycle space were provided, then computes the initial residual (corresponding to line 1 of Algorithm 1), $r = b - Ax_0$. It then computes the norm of the initial residual and checks if the initial guess x_0 is a converged solution. If yes, it returns the result (but generally this is not the case). If a recycle space U is provided, it orthogonalizes it against the current approximation to the solution (current approximation computed on line 6 of Algorithm 1), updates the initial guess and residual, and checks for convergence again. The symmetric tridiagonal matrix T generated by the Lanczos method [12] is not given in Algorithm 1, but we note that this is the procedure used by MINRES and RMINRES to build an orthogonal basis for the Krylov space using successive matrix-vector products. These are fairly easy to implement in Python, however we must compute the QR factorization of T to compute the update to the approximate solution.

The code continues this iterative loop to refine the solution while monitoring convergence. Within the loop, several key things happen. (1) It applies (right) preconditioning, if necessary, (2) It computes a matrix-vector product to expand the Krylov space, and (3) It orthogonalizes the new vector against the recycle space as well as the Krylov space. Finally, (4) it computes the elements of the tridiagonal matrix, T , arising in the Lanczos iteration. After computing the QR factorization of T to compute an update to the approximate solution x_k at the k th iteration, we must compute the eigenvalues of T , and update the solution. We have pointed out key steps in the algorithm that posed particular challenges when implementing the RMINRES algorithm. There are other steps that are fairly straightforward and that we omit for brevity. For instance, the code updates various variables for the next iteration, including values related to the Lanczos method. Upon convergence at iteration m , it returns the approximate solution x_m , as well as a number of other outputs, if requested by the user (e.g., a vector *relres* containing the residuals at each iteration). It

also provides the new recycle space to then use for the next linear system as the set of vectors we have referred to as U .

B. Sparse Approximate Maps Implementation

We seek to eventually integrate SAMs as part of the larger predictor-corrector method code to explore their use as a mapping between matrices to recycle preconditioners when they are necessary. We again note that they were not required for the linear systems discussed in Figures 2 - 4, but in cases when the systems may become ill-conditioned, the role of SAMs will become vital to reduce costs associated with preconditioning many linear systems in a long sequence; see [7] for experiments demonstrating how SAMs reduced such costs. The update scheme in [7] solves

$$N_k = \operatorname{argmin}_{N \in \mathcal{S}} \|A_k N - A_0\|_F \quad (2)$$

and defines the updated preconditioner as $P_k = N_k P_0$. Here, $\mathcal{S}$ represents the subspace defined by a selected sparsity pattern [7]. We assume the matrices A_0 and A_k are part of a given sequence and are relatively close to each other.

The code first allows the user to set a sparsity pattern, $\mathcal{S}$, that is imposed on the SAM. The code computes a map to transform one system matrix into another matrix in the sequence using equation (2), which defines the full least squares problem (e.g., $n \times n$). However, with a judiciously chosen sparsity pattern (i.e., one that balances a small number of nonzeros with the accuracy as a map), the least squares equation given in (2) results n independent (and thus parallel), very small least squares problems. Specifically, it computes the least squares solution n_j for the equation $G n_j = A s_0$, where G and $A s_0$ are submatrices of A_k and A_0 , respectively, that respect the chosen sparsity pattern. Each solution, n_j , for the n least squares problems defines the j^{th} column of the approximate map N_k . The algorithms for SAMs and an analysis of the choice in sparsity pattern for several applications can be found in [7] (see Algorithms 4.1 and 4.2).

VII. CHALLENGES AND PROPOSED SOLUTIONS

We first discuss challenges we encountered while implementing the RMINRES algorithm in Python and PyTorch. Python requires a series of preprocessing steps to initialize appropriate data structures, variables, and parameters before beginning the algorithm directly. One example is what happens after loading the files supplied (i.e., csv files storing the matrix elements), where the data must then be converted to arrays to perform calculations. Another key difference is often the data must be reshaped to fit the variety of the SciPy sparse library and to suit user/application necessity (e.g., compressed sparse column or compressed sparse row).

We use compressed sparse column, and so the `scipy.sparse` function is used. The section of code seen in Figure 5 includes the first attempts at using `torch.linalg.solve` function for solving the systems of matrices. However, as seen in Figure 6, the function does not execute. The error provided in Figure 6 is longer than shown below, however the most significant comments are provided.

```
if (existM1):
    r = torch.linalg.solve(M1, r)
if (existM1):
    b = torch.tensor(b)
    normb = torch.linalg.norm(torch.linalg.solve(M1, b))
else:
    normb = torch.norm(b)
```

Fig. 5: Code from RMINRES that fails due to underdevelopment of sparse tensors and `torch.linalg.solve` function.

```
NotImplementedError                                Traceback (most recent call last)
Cell In[28], line 33
    31 A1 = torch.from_numpy(A1) #sparse tensor
    32 b1 = torch.from_numpy(b1) #tensor
--> 33 x, flag, rres, iter, resvec = rminres(r, A1, b1, x0, tol, 1000, 1, 1, 1, 0, cycle, recycle)
    36 NRBE = torch.norm(b1 - A1 @ x) / (torch.norm(b1 - A1 @ x0) + torch.norm(A1, float('inf')) * torch.norm(x))
    38 plot = np.log10(resvec / resvec[0])

Cell In[27], line 165, in rminres(r, A, b, x0, tol, maxiter, M1, M2, U, dim, k)
    164 if (existM1):
--> 165     r = torch.linalg.solve(r, M1) # r= M1\ r
    166 if (existM1):
    167     normb = torch.linalg.matrix_norm(torch.linalg.solve(M1, b))

NotImplementedError: Could not run 'aten::linalg_solve_ex' with arguments from the 'SparseCPU' backend. This could be because the operator doesn't exist for this backend, or was omitted during the selective/custom build process (if using custom build). If you are a Facebook employee using PyTorch on mobile, please visit https://fburl.com/ptmfixes for possible resolutions. 'aten::linalg_solve_ex' is only available for these backends: [CPU, CUDA, HIP, MPS, iPU, XPU, HPU, VE, Meta, MTIA, PrivateUse1, PrivateUse2, PrivateUse3, FPGA, ORT, Vulkan, Metal, QuantizedCPU, QuantizedCUDA, QuantizedHIP, QuantizedMPS, QuantizedIPU, QuantizedXPU, QuantizedHPU, QuantizedVE, QuantizedMeta, QuantizedMTIA, QuantizedPrivateUse1, QuantizedPrivateUse2, QuantizedPrivateUse3, CustomRNGKeyid, MklDnnCPU, SparseCsrCPU, SparseCsrCUDA, BackendSelect, Python, FuncTorchDynamicallyLayerBackMode, Functionalize, Named, Conjugate, Negative, ZeroTensor, ADInplaceOrView, AutogradOther, AutogradCPU, AutogradCUDA, AutogradHIP, AutogradXLA, AutogradMPS, AutogradIPU, AutogradXPU, AutogradHPU, AutogradVE, AutogradLazy, AutogradMeta, AutogradMTIA, AutogradPrivateUse1, AutogradPrivateUse2, AutogradPrivateUse3, AutogradNestedTensor, Tracer, AutocastCPU, AutocastCUDA, FuncTorchBatched, FuncTorchVMmapMode, Batched, VmapMode, FuncTorchGradWrapper, PythonTLSSnapshot, FuncTorchDynamicallyLayerFromTMode, PythonDispatcher].
```

Fig. 6: `NotImplementedError` after attempting to run `torch.linalg.solve` function in RMINRES implementation

The use of additional operations is fundamental to this predictor-corrector method, as the matrices within this code must be in a format that allows for the solution of systems of linear equations and vector multiplication. In our implementations, it became clear that more development of the sparse tensor data structure is needed in order to successfully incorporate RMINRES into PyTorch.

Additionally, creating functions within PyTorch with diverse data type compatibility will allow for solving a larger range of calculations within the PyTorch language and environment. This also includes different ways to compute matrix-vector multiplication for the predictor-corrector code, as Hv is defined as a function and H is not explicitly known. In the RMINRES code, we require not just matrix-vector multiplication, but matrix-matrix products (i.e., when orthogonalizing the subspace). More development is needed in this area, and part of ongoing work. While the predictor-corrector method and RMINRES solvers are implemented in Python [13], the development of PyTorch sparse tensor data structure and compatible functions to complete the remaining portions of the RMINRES implementation is part of the ongoing work.

Within the SAMs code, the issues we have encountered are related to mismatched data types, allocation of the appropriate data structure, and errors within called functions. As matrices are imported and converted to a sparse data structure, the updated matrices and data collected from computation needs

to be converted and reshaped for the function at each step of the call. As with RMINRES, changes to the shape of the data are also required. Many of these issues have been resolved, however issues within the NumPy linalg.solve function called to solve linear systems are ongoing. An example of this is resolving data type matching errors for the implementation of the solver. Transitions between many data types occurs during the SAMs implementation that play a large role in the output. The segment of code shown in Figure 7 is part of the preprocessing of data for the functionality within the SAMs implementation.

```
data = np.ones(np.shape(rows))
rows = np.array(rows)
cols = np.array(cols)
msk = scipy.sparse.csc_matrix((data, (cols, rows)))
rows = rows.reshape(-1, 1)
cols = cols.reshape(-1, 1)
j, i = As.nonzero()
findA = scipy.sparse.csc_matrix((np.ones(np.shape(i)), (i, j)))
PP1 = findA != 0
PP = PP1.multiply(msk)
PP2 = PP.copy()
```

Fig. 7: Examples of the data type agreement in Python.

Lastly, to parallelize the code, we can utilize the Python multiprocessing library. This library allows us to create separate processes for individual tasks, which can run simultaneously and speed up the computation. We will divide the loop into separate tasks that can run concurrently. The code must be rewritten to reflect the parallel structure, particularly in the loop that computes the n small least squares problems. Parallelizing the existing code is also part of ongoing work.

VIII. DISCUSSION

In comparing the computational costs expected from our ongoing implementation, an important note is assessing the efficiency gains offered by RMINRES in contrast to the conventional MINRES method. RMINRES has the potential to significantly reduce computational cost when solving sequences of linear systems by reducing the total number of iterations to converge to an accurate-enough approximate solution. This method updates iteratively to refine solutions, capitalizing on the shared data between consecutive linear systems. However, it is still important to note that the computational cost and time savings will depend on factors such as the size and sparsity of the linear systems, the nature of the coefficient matrices, and the specific problem at hand. SAMs can be useful when dealing with ill-conditioned matrices by reducing costs associated with computing a preconditioner from scratch for every linear system in the sequence.

IX. CONCLUSIONS AND FUTURE WORK

We seek to reduce the overall cost of computing the linear systems arising in the predictor step of the predictor-corrector method utilized in [3] using recycling methods for both the iterative solver as well as for preconditioners. In this work, we considered the linear solver RMINRES and a map between

coefficient matrices, the sparse approximate map. SAMs allow us to map closely related matrices, and may allow us to recycle previously computed preconditioners. However, these methods do not yet exist in PyTorch, and thus have not yet been implemented in the Pareto front code described in [3]. While the predictor-corrector method and RMINRES solvers are implemented in Python, the development of PyTorch sparse tensor data structure and compatible functions is part of the ongoing and future work. With development in PyTorch linear algebra solver and sparse tensor data structure, we expect that the RMINRES and SAMs implementations will help to reduce computational cost in the overall predictor-corrector method for computing an optimal Pareto front [3].

REFERENCES

- [1] Ma, P., Du, T., & Matusik, W. (2020, November). Efficient continuous pareto exploration in multi-task learning. In International Conference on Machine Learning (pp. 6522-6531). PMLR.
- [2] Liu, S., & Vicente, L. N. (2019). The stochastic multi-gradient algorithm for multi-objective optimization and its application to supervised machine learning.
- [3] Enouen, Eric & Mathesius, Katja & Wang, Sean & Carr, Arielle & Xie, Sihong. (2022). Efficient first-order predictor-corrector multiple objective optimization for fair misinformation detection. 10.48550/arXiv.2209.07245.
- [4] Paige, Christopher C and Saunders, Michael A. Solution of sparse indefinite systems of linear equations. SIAM Journal on Numerical Analysis, Volume 12, 617–629, (1975).
- [5] Hestenes, M.R., & Stiefel, E. (1952). Methods of conjugate gradients for solving linear systems. Journal of research of the National Bureau of Standards, 49, 409-435.
- [6] Wang, Shun and de Sturler, Eric and Paulino, Glaucio H. Large-scale topology optimization using preconditioned Krylov subspace methods with recycling. International Journal for Numerical Methods in Engineering, Volume 69, Number 12, 2441–2468, (2007).
- [7] Carr, Arielle, & de Sturler, Eric, & Gugercin, Serkan Gugercin. Preconditioning Parameterized Linear Systems. SIAM Journal on Scientific Computing, Volume 43, Number 3, A2242-A2267, (2021). <https://doi.org/10.1137/20M11331123>.
- [8] PyTorch Development Team. (2023). torch.linalg.solve — PyTorch documentation. PyTorch. <https://pytorch.org/docs/stable/generated/torch.linalg.solve.html>
- [9] Benzi, Michele. Preconditioning techniques for large linear systems: a survey. Journal of Computational Physics, Volume 182, 418–477, (2002).
- [10] Saad, Yousef (2003). Iterative methods for sparse linear systems (2nd ed.). SIAM. ISBN 0-89871-534-2. OCLC 51266114.
- [11] Narchy, A. (2020). KryPy: A Python Library for Krylov Subspace Methods. Version 1.2. GitHub. <https://github.com/andrenarchy/krypy>
- [12] Yousef Saad. Iterative Methods for Sparse Linear Systems, 2nd Ed. SIAM, 2003.
- [13] Krypy. (n.d.). RecyclingMinres. In Krypy Documentation. From <https://krypy.readthedocs.io/en/latest/krypy.recycling.RecyclingMinres>
- [14] Gargiani, M., Zanelli, A., Diehl, M., & Hutter, F. (2020). On the promise of the stochastic generalized Gauss-Newton method for training DNNs. arXiv preprint arXiv:2006.02409.

X. ACKNOWLEDGMENTS

We thank the National Science Foundation for funding the Lehigh University Intelligent and Scalable Systems REU. Ainara Garcia is supported by the National Science Foundation under Grants NSF CNS-2051037. Any opinions, findings, conclusions, or recommendations expressed in this document are those of the author(s) and should not be interpreted as the views of the National Science Foundation.

We also thank Sean Wang and Rishad Islam for their helpful consultations on this project.