

Sizey: Memory-Efficient Execution of Scientific Workflow Tasks

Jonathan Bader*, Fabian Skalski*, Fabian Lehmann[†], Dominik Scheinert*,
Jonathan Will*, Lauritz Thamsen[‡], and Odej Kao*

* {firstname.lastname}@tu-berlin.de, Technische Universität Berlin, Germany

[†] fabian.lehmann@informatik.hu-berlin.de, Humboldt-Universität zu Berlin, Germany

[‡] lauritz.thamsen@glasgow.ac.uk, University of Glasgow, United Kingdom

Abstract—As the amount of available data continues to grow in fields as diverse as bioinformatics, physics, and remote sensing, the importance of scientific workflows in the design and implementation of reproducible data analysis pipelines increases. When developing workflows, resource requirements must be defined for each type of task in the workflow. Typically, task types vary widely in their computational demands because they are simply wrappers for arbitrary black-box analysis tools. Furthermore, the resource consumption for the same task type can vary considerably as well due to different inputs. Since underestimating memory resources leads to bottlenecks and task failures, workflow developers tend to overestimate memory resources. However, overprovisioning of memory wastes resources and limits cluster throughput.

Addressing this problem, we propose *Sizey*, a novel online memory prediction method for workflow tasks. During workflow execution, *Sizey* simultaneously trains multiple machine learning models and then dynamically selects the best model for each workflow task. To evaluate the quality of the model, we introduce a novel resource allocation quality (RAQ) score based on memory prediction accuracy and efficiency. *Sizey*'s prediction models are retrained and re-evaluated online during workflow execution, continuously incorporating metrics from completed tasks.

Our evaluation with a prototype implementation of *Sizey* uses metrics from six real-world scientific workflows from the popular *nf-core* framework and shows a median reduction in memory waste over time of 24.68% compared to the respective best-performing state-of-the-art baseline.

Index Terms—Resource Management, Scientific Workflow, Memory Allocation, Memory Prediction, Machine Learning

I. INTRODUCTION

Scientific workflow management systems (SWMS) such as Nextflow [1], Pegasus [2], or Snakemake [3] enable the creation of reproducible data analysis workflows. SWMS help scientists from various domains, such as genomics, remote sensing, and material science [4]–[9], to abstract execution-specific details and to cope with huge amounts of data. Workflows are often defined as a directed acyclic graph (DAG), consisting of a set of black-box task types B and a set of directed edges E . Each task type b serves as a wrapper for a black-box analysis tool and as a template for physical task instances T with concrete inputs. Each individual edge e connects two tasks and expresses their dependency, dataflow, and execution order.

SWMS submit physical task instances to a cluster resource manager such as Kubernetes [10], Slurm [11], or

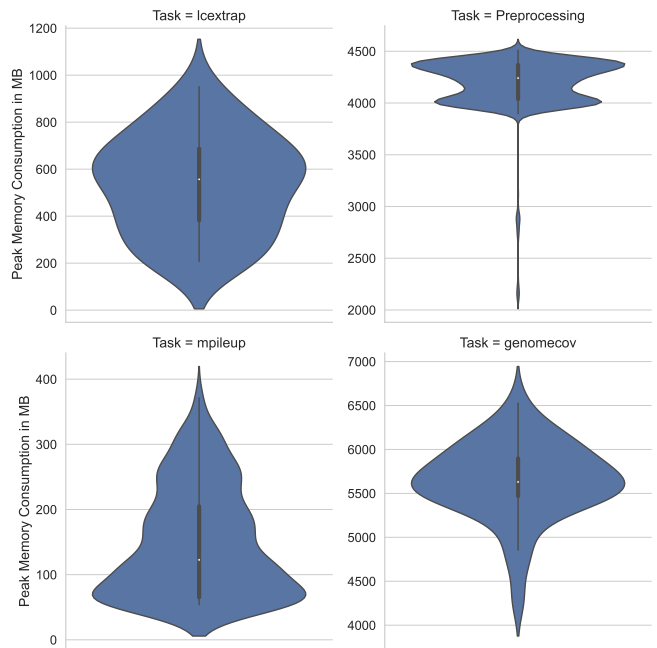


Fig. 1: Distribution of peak memory consumption of four task types, each executed repeatedly with varying input sizes.

HTCondor [12]. To assign each physical task to a suitable cluster node, resource managers rely on adequate resource requirements, such as the number of processor cores and main memory required. Typically, task type resource requirements are provided by users and workflow developers and are essentially estimates [13]–[15]. However, estimating the resource requirements of tasks is known to be difficult and error prone [16]–[19]. At the same time, accurate estimates of especially the required memory are crucial. The underestimation of memory can result in bottlenecks, costly offloading to disks, and even task failures. Therefore, users are incentivized to overestimate resource usage to avoid partial or complete restarts of workflows. However, requesting more memory than needed wastes precious cluster resources and effectively limits parallelism within and across workflow applications.

The memory consumption of executing different physical



Fig. 2: Memory consumption in relation to the input read of the physical task instances of two different task types, with a linear regression applied to these data points.

tasks that use the same task template but different inputs can vary substantially and often depends on the input data. Figure 1 shows the distribution of memory consumption for four different task types that were executed repeatedly with varying input sizes. It can be observed that the memory consumption for executions of the same task type differs significantly. Furthermore, the memory consumption also varies between the different task types. For example, the *lcextrap* task instances consume between 200MB and 1GB with a median of approximately 550MB. Consequently, the prediction of the peak memory usage of task instances should be automated, taking into account the differences between task types, the variances in memory consumption between their concrete physical task executions, as well as the impact of input data. Further, incorporating resource measurements online, i.e., during the execution of the workflow, allows refining prediction models during runtime.

The majority of existing methods for predicting memory requirements in scientific workflows employ machine learning to address the challenge of memory prediction [19]. However, these methods neglect that not each workflow task can be modeled well with the same type of machine learning model. For instance, while many workflow tasks yield a linear memory consumption behavior correlating with the input data size, this is not the case for all. Figure 2 shows how the memory consumption of tasks depends on input reads. It can be seen that there is a clear linear correlation for *MarkDuplicates*. However, using a linear model for the *BaseRecalibrator* task would lead to half of the task instances failing due to insufficient memory, and the other half would waste significant memory due to overestimates. Therefore, an adaptive prediction method is needed that selects an adequate machine learning model for all task types.

This paper presents Sizey, a novel online task memory prediction method for scientific workflows that employs the best out of multiple machine learning models for each task. Upon submission of a task, Sizey searches a provenance database for previous task executions of the same task type. If the task is of an unknown task type, it is submitted directly to the resource manager, resorting to the user-provided, usually

conservative memory estimate. On completion of the task, the monitoring data are stored in the provenance database. Based on the previous executions, Sizey trains a set of diverse machine learning models in parallel, with each model’s prediction yielding a memory estimate. Subsequently, the prediction performance of the models is assessed using a new resource allocation quality (RAQ) score, which measures both the observed accuracy and the estimated resource efficiency. Based on the RAQ score, Sizey selects the best-performing prediction model for each task and adjusts the prediction with a dynamic fault tolerance offset to minimize costly task failures caused by underestimates. Finally, as new data become available after the completion of each task, a lightweight - and thus fast - online learning step updates the existing models.

Contributions. This paper makes the following contributions:

- We present our novel task memory prediction method, Sizey, which employs a set of diverse machine learning models to adaptively predict a task’s memory in order to reduce memory resource wastage.
- We provide an open-source implementation of our method Sizey, as an easily extendable interface¹.
- We evaluate Sizey using six large-scale workflows from the nf-core repository [4] and compare it to four state-of-the-art methods. The experimental results demonstrate that Sizey outperforms all state-of-the-art baselines with a median memory wastage reduction of at least 24.68% compared to all baselines.

II. APPROACH

In this section, we first provide an overview of our proposed method. Second, we explain the granularity of the prediction model used by Sizey and the selection of the model classes. Third, we explain our resource allocation quality (RAQ) score that combines model accuracy and resource efficiency. Fourth, we describe how we weight and combine the prediction results into a final prediction. Lastly, we explain our initial sizing strategy and handling of task failures due to underprovisioning.

A. Overview

Sizey estimates the peak memory consumption of a submitted task instance in an online manner during the execution of a workflow. Figure 3 illustrates how this works.

a) Data Retrieval: In Phase ①, a task is submitted for execution, which causes Sizey to retrieve provenance details from the provenance database connected to the SWMS, such as the number of running tasks and information about historical task executions, including the task name and resource consumption.

b) Memory Prediction: In Phase ②, Sizey estimates the memory consumption of the task to be executed. For this, in Step 2.1, a set of diverse machine learning models is used, where each predictor produces a resource estimate. Next, our resource allocation quality (RAQ) score compares the accuracy and efficiency of the individual predictors. In Step 2.2, a gating

¹github.com/dos-group/sizey




Fig. 3: The figure provides an overview of how Sizey (green) is used in a scientific workflow execution environment. Sizey uses data from a provenance database and trains multiple machine learning models in parallel. The individual predictions are then combined into a single prediction that also includes an offset to ensure sufficient memory is allocated for tasks executing on the cluster. During workflow execution and after task completion, the model is retrained and reevaluated.

mechanism selects the models and, depending on the strategy, decides between aggregating and maximizing the individual predictions according to their RAQ scores to produce a single prediction. In Step 2.3, this single prediction is adjusted based on the selected offset strategy and how often the task execution was attempted before. Subsequently, the task instance and resource prediction are submitted to the cluster resource manager for executing the task with the determined memory.

c) *Online Learning*: In Phase ③, a task has been completed and new training data becomes available. Therefore, in this online learning phase, all models are updated. The update steps are efficient due to the lightweight inputs and models used.

d) *Assumptions*: We make the following assumptions for predicting the memory consumption of workflow tasks:

- A1: We assume that workflows are used on many input files, thus involving many data-parallel task executions and repeated executions of the same task type, allowing online learning during workflow execution.
- A2: We assume that scheduling, which involves ordering task instances and assigning them to cluster nodes, is the responsibility of the resource manager and is therefore outside the scope of this work.
- A3: We assume that the resource manager enforces strict resource limits on memory allocations, resulting in a failed task execution when exceeding these limits.

B. Model Granularity, Classes, and Features

To achieve the goal of memory wastage reduction, precise models for resource predictions are needed. Sizey uses multiple fine-granular models to predict the memory consumption. This technique is visualized in Figure 4, differentiat-

ing between a workflow-specific model, a task type specific model, a task/machine-specific model, and our proposed model granularity. Sizey employs the most fine-granular model by incorporating several models for each task-machine configuration. The rationale behind this approach is that tasks exhibit heterogeneous computational patterns that vary even more with different machine configurations. This implies that a model type that performs well on one task type may perform poorly on another.

Figure 5 shows the model classes incorporated by Sizey. As stated in [18], [20], [21], tasks frequently yield a linear relationship between input data size and memory consumption. We exploit this observation by integrating a linear regression model class into the pool of used model classes. However, as presented in the introduction, the relationship between input data size and memory consumption does not have to be linear. The k-nearest neighbors (KNN) regression model allows historical observations similar to the task currently estimated to influence the resource prediction. We include MLP regression to accurately model more complex, nonlinear relationships, such as memory usage that grows as the square of the amount of input data. Incorporating a random forest regression model makes our method more resistant to overfitting, especially when there are not many historical task executions.

Since Sizey runs online, there is no classical segmentation into training and test data during workflow execution. Instead, the prediction accuracy of individual models is permanently assessed. When new task resource measurements become available as tasks finish, each model performs an update step. Such an update step can also be a complete retraining of the models if 1) the training time is short, 2) the time until the next task submission is long enough for a complete retraining,




Fig. 4: Granularity levels for a given workflow with model A, task types 1, 2 on machines I and II; Sizy method in bold green box.




Fig. 5: The four model classes used by Sizy for modeling linear and non-linear relationships between input data sizes and peak memory consumption.

or 3) a complete retraining for better performance is desired.

C. Resource Allocation Quality (RAQ)

Sizy makes use of a set of machine learning models. This poses the question of how to select the best-fitting model or even a model combination for certain abstract task-machine configurations.

The RAQ score describes the presumed goodness of the task's resource estimate provided by one associated trained model. The composite score is derived from two subscores: the accuracy score and the efficiency score. All three scores, the RAQ score, the accuracy score, and the efficiency score, are normalized scalars between 0 (worst) and 1 (best).

a) *Accuracy Score*: The accuracy score expresses a model's average prediction performance. Let $T = \{t^{(i)} | i \in \mathbb{N}_0, \text{ and } 0 \leq i \leq S-1\}$ be the set of historical task instances, with S denoting the size of the set. The resource estimate for a historical task instance $t^{(j)}$ given by the i -th model is denoted as $\hat{y}_{i,t^{(j)}}$, while $y_{t^{(j)}}$ denotes the actual resource consumption of $t^{(j)}$. Then, the accuracy score AS_{i,t^*} of the i -th model associated with the combination of task type and machine type for the current task t^* can be calculated as follows:

$$AS_{i,t^*} = \frac{1}{S} \sum_{t^{(j)} \in T} \left(1 - \min\left(\left| \frac{\hat{y}_{i,t^{(j)}} - y_{t^{(j)}}}{y_{t^{(j)}}} \right|, 1\right) \right). \quad (1)$$

In Equation 1, the individual error terms are bounded at one to prohibit large estimation outliers from skewing the computed scores. The resulting accuracy score describes the mean prediction performance of a model across the set of task instances. A mean prediction error of zero indicates a perfectly accurate prediction and would result in an accuracy score of 1, the best possible value. Accuracy scores are updated over time, while models predict and learn from new task data.

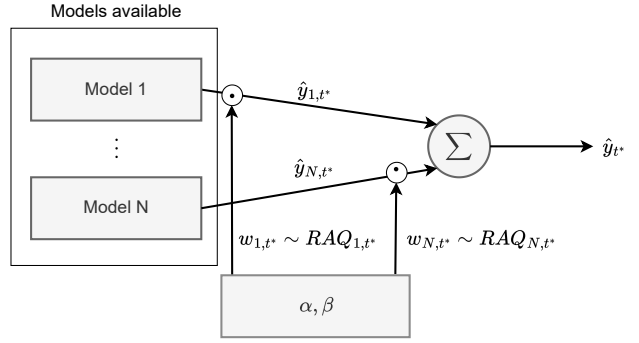


Fig. 6: Role of resource allocation quality scores in model selection and weighting

b) *Efficiency Score*: The efficiency score compares a model's individual memory predictions to the other model's memory predictions. With N being the number of models available for the desired combination of a task type and a machine configuration and \hat{y}_{i,t^*} the prediction of the i -th model for the current task t^* , its dynamic efficiency score ES_{i,t^*} is calculated as follows:

$$ES_{i,t^*} = 1 - \frac{\hat{y}_{i,t^*}}{\max_{j=1,\dots,N}(\hat{y}_{j,t^*})}. \quad (2)$$

Efficiency scores are model output-specific and punish outlying resource estimates that are huge in magnitude relative to other model outputs. These large estimation outliers may occur in various cases, for instance, when extrapolation occurs during the early training stages. By validating individual estimate magnitudes with those from other predictors, the risk of wasting anomalously large amounts of memory resources is mitigated. The efficiency score associated with the largest resource estimate is always 0.

c) *Computing the Resource Allocation Quality*: Our RAQ score includes one mandatory hyperparameter α . With $\alpha \in [0, 1]$ and AS_{i,t^*} and ES_{i,t^*} being the respective accuracy and efficiency scores, the resource allocation quality RAQ_{i,t^*} of the i -th model for the current task t^* can be calculated as:

$$RAQ_{i,t^*} = (1 - \alpha) \cdot AS_{i,t^*} + \alpha \cdot ES_{i,t^*}. \quad (3)$$

When α is close to zero, accurate models are favored, whereas values close to one punish large outlying estimates more strongly.

D. Model Selection Strategy

Figure 6 shows the structure of our proposed gating mechanism that is used to select the models that contribute to the final prediction. Similar to a gating network from the mixture of experts domain [22], [23], our gating mechanism assigns individual weights to the outputs of the individual predictors. These weights depend on our RAQ score that we introduced before.

Sizy can be used with two different gating strategies. The gating strategy determines the predictor's weights from Figure 6.

a) *Argmax Strategy*: The first strategy is called *Argmax* and weights all predictors with zero except the one corresponding to the highest RAQ score. In that case, the output of the associated predictor is chosen as the aggregate resource estimate \hat{y}_{t^*} .

b) *Interpolation Strategy*: Whenever the *interpolation* strategy is enabled, the hyperparameter $\beta \in [1, \infty)$ becomes mandatory and weights are calculated using softmax:

$$\hat{y}_{t^*} = \sum_{i=1}^N \left(\hat{y}_{i,t^*} \cdot w_{i,t^*} \right) = \sum_{i=1}^N \left(\hat{y}_{i,t^*} \cdot \frac{e^{\beta \cdot RAQ_{i,t^*}}}{\sum_{j=1}^N e^{\beta \cdot RAQ_{j,t^*}}} \right). \quad (4)$$

While the *Argmax* strategy is more opportunistic towards the presumably best available model, the *interpolation* strategy seeks to utilize a model output consensus.

E. Prediction Offset and Failure Handling

Since *Sizey* aims to predict memory as accurately as possible, small underpredictions can quickly lead to task failures. Therefore, we use a dynamic offset strategy that increases *Sizey*'s initial memory prediction for the task. As with the models, our dynamic offset strategy uses several possible offsets, the standard deviation, the standard deviation of underpredictions, the median prediction error, and the median prediction error of underpredictions. During online learning, *Sizey* selects the offset that would have caused the least wastage based on the tasks already executed.

If a task instance still fails due to underprediction, the maximum amount of task memory ever observed is allocated. For each subsequent attempt to size a previously failed task instance, the given resource estimate is continuously doubled until the machine resources are exhausted.

III. EVALUATION

Our evaluation section is divided into five subsections: experimental setup, baselines, experimental results, *Sizey* model insights, and the discussion of the results.

A. Experimental Setup

For our experiments, we executed and measured six real-world workflows on our cluster infrastructure, five workflows from the bioinformatics domain, and one workflow from the field of remote sensing:

- 1) For the **eager workflow** [24], we use input data published in a study in 2018 by de Barros Damgaard et al. on the population history of the Eurasian steppe [25].
- 2) For the **rnaseq workflow**², we use data from a bladder cancer cells study [26].
- 3) For the **mag workflow**³, we use input data from a study of the gut microbiome [27].
- 4) For the **chipseq workflow**⁴ we use data from a prostate cancer study [28].

- 5) The input data for the **methyseq workflow**⁵ are from a study of human pluripotent stem cells in 2020 by Thompson et al. [29].
- 6) The **iwd workflow** [30] uses more than 300 images within and around fire scars in Western Alaska and uses computer vision and graph analysis to examine the landscapes.

For data collection, we run these workflows on an eight-node Kubernetes cluster, each equipped with an AMD EPYC 7282 16-core 32-thread processor, 128GB of DDR4 memory, and two 960GB SATA III SSDs. The Kubernetes cluster uses a Ceph file system for data storage.

Figure 7 illustrates the memory, CPU, and I/O resource consumption of the experimental workflows. One can see that all workflows yield different resource usage patterns. For instance, the **methyseq workflow** yields many I/O-intensive tasks, while being less CPU-intensive. Table I illustrates the number of task types and the average number of task instances per task type on a workflow basis. The **rnaseq** and the **chipseq workflow** have the greatest number of task types, yet the lowest average number of task instances per task type. In contrast, the **iwd workflow** has the lowest number of task types, yet the second highest average number of task instances per task type. Using such a heterogeneous set of workflows and task types allows us to test the heterogeneous computing landscape of shared cluster environments.

To enable a comprehensive evaluation of memory prediction of large-scale workflow executions with multiple methods, we simulate an online environment where our measured real-world metrics from completed task executions can be incorporated into the learning process, similar to real-world systems. Simulation parameters can be used to define aspects such as the amount of historical data, failure strategies, model parameters, etc. An important parameter we will use in our experiments is the time-to-failure value, which expresses at what point in time a task will fail. A time-to-failure value of 1.0 means that we expect tasks to fail at the end of their execution. We filter out very short-running tasks and tasks with a single or few executions.

Our implementation is written in Python 3.11 and uses *scikit-learn*, *numpy*, and *scipy* as the main libraries and keeps the models in memory. For our simulation experiments, we used all four machine learning models: linear regression, knn regression, MLP regression, and random forest regression. We have implemented a version that always fully retrains the models and a version that provides incremental updates and caches the best hyperparameters over the workflow execution. In addition, we select an α value of 0.0 and use the *Interpolation* strategy for *Sizey*. A comparison of the full model training and the incremental updates, as well as an analysis of selecting the parameters, can be found in the model insights evaluation (Section III-D).

²github.com/nf-core/rnaseq

³github.com/nf-core/mag

⁴github.com/nf-core/chipseq

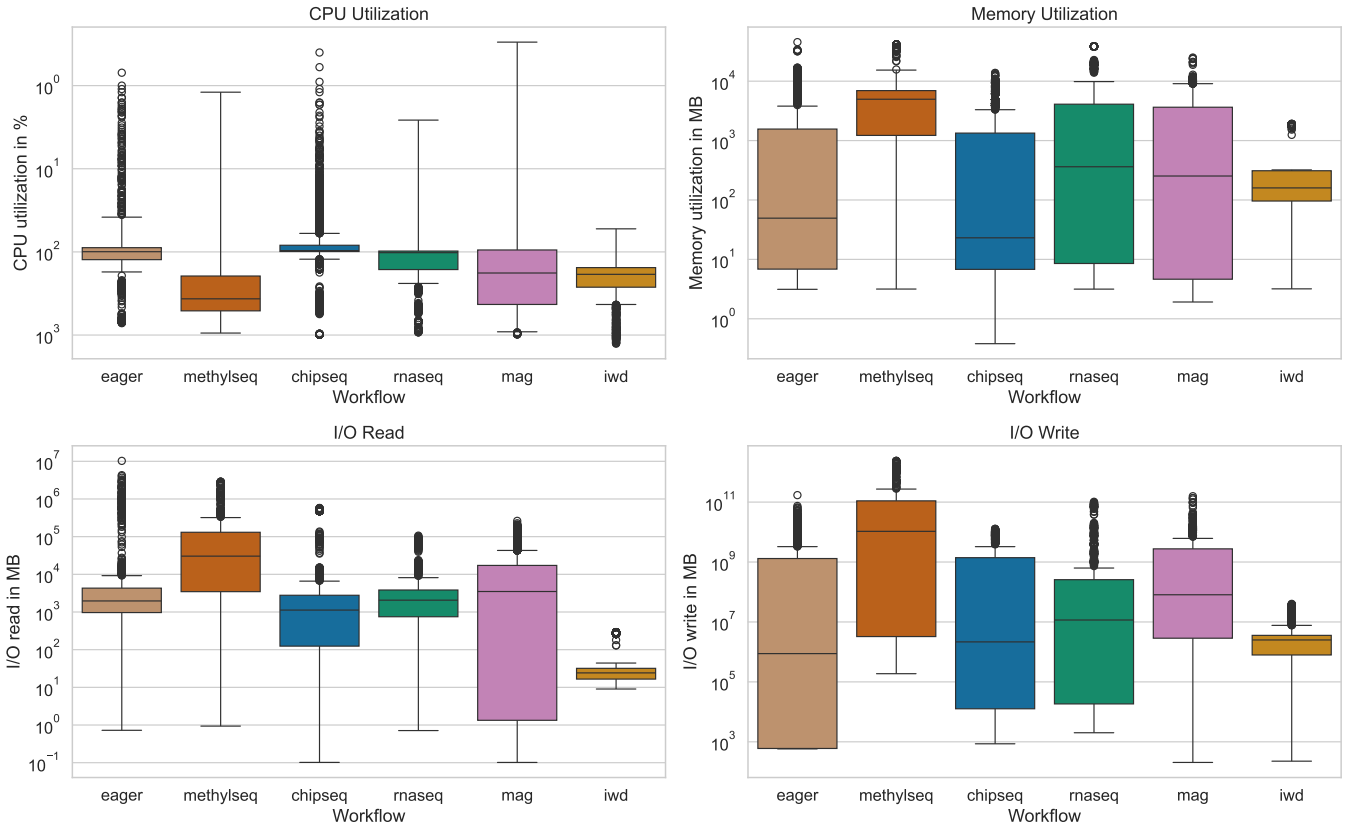


Fig. 7: Distribution of memory, CPU, and I/O utilization of the executed workflows on a logarithmic x-axis scale.

TABLE I: The table shows the average number of task instances per task type for each experimental workflow.

Workflow	# Task Types	AVG # Task Instances per Task Type
eager	13	121
methylseq	9	100
chipseq	30	82
rna-seq	30	39
mag	8	720
iwd	5	332

B. Baselines

Our method is evaluated against four state-of-the-art baselines along with the default workflow setups provided by the workflow developers, which serve as a sanity baseline.

Tovar et al. [31] use historical peak memory probabilities (**Tovar-PPM**) to determine the initial memory allocation for workflow tasks. Their strategy aims to reduce the overall probability that resource peaks will exceed the allocated memory. Should the initial allocation underestimate the required resource, resulting in task failure, Tovar et al. allocate a node’s maximum memory. We used the authors’ source code and integrated it into our setup.

Witt et al. [18] propose a low-wastage regression that optimizes the resource wastage instead of the prediction error

and is also applied in an online fashion. Their method is based on a linear model and doubles the predicted memory upon task failure. We use the source code provided on the author’s Github account and refer to the method as **Witt-Wastage**.



Witt et al. [32] use two different methods using an online learning mode. The percentile predictor predicts the percentile peak memory usage of all historical tasks. The authors propose a conservative estimate, using the 95th percentile to avoid task failures. As a second method, they propose a linear regression (LR), using the input size as a feature and adding an offset on the prediction. We refer to these methods as **Witt-Percentile** and **Witt-LR**. Given the absence of accessible source code, our implementation of the method was based on the descriptions provided in the paper.

C. Experimental Results

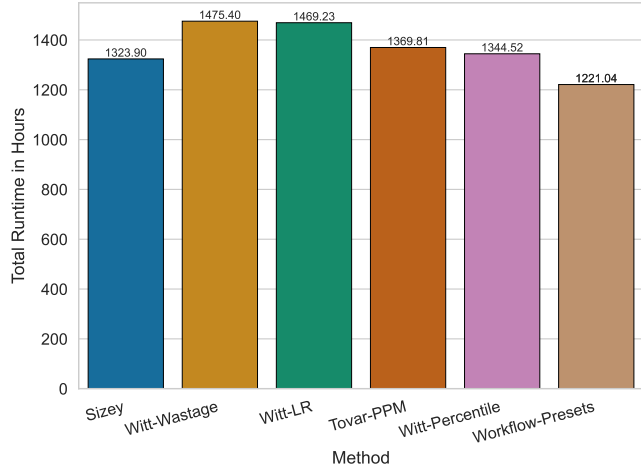

Figure 8 shows the aggregated results of the memory prediction for six different workflows, including the memory wastage over time for two different time-to-failure values, the number of task failures, and the aggregated task runtimes.

Figure 8a shows the aggregated wastage over time with a time-to-failure value of 1.0 for all methods in gigabyte-hours. Among all methods, the Workflow-Presets show the highest wastage, approximately 17 times higher than Sizey and 6 times higher than the best-performing baseline. Sizey exhibits the lowest wastage over time for all methods and achieves

⁵github.com/nf-core/methylseq



(a) Wastage over time in gigabyte-hours (GBh) aggregated over all workflows with a time-to-failure of 1.0 on a logarithmic y-axis scale. (b) Wastage over time in gigabyte-hours (GBh) aggregated over all workflows with a time to failure of 0.5 on a logarithmic y-axis scale.



(c) Distribution of the methods' task failures, aggregated by task type, with outliers hidden.

(d) Aggregated workflow task runtimes for each method.

Fig. 8: Experimental results of predicting memory consumption for tasks of six workflows.

a 64.58% lower wastage compared to the best-performing baseline, Witt-Wastage. Also, all other baselines are able to outperform the Workflow-Presets but still show significantly higher wastage over time compared to Sizey.

Figure 8b shows the aggregated wastage over time for all methods in gigabyte-hours. However, we now configure a time-to-failure value of 0.5, which means that tasks will fail earlier during execution. Here we can see that Sizey also shows the lowest wastage, a decrease of 60.60% compared to the Witt-Wastage baseline. In contrast to Figure 8a, the relative difference in wastage between Sizey and the Witt-Wastage method has slightly decreased. Again, the Workflow-Presets show the highest wastage and also the same wastage as before as they are not impacted by the time-to-failure value. It is also observed that all state-of-the-art methods benefit from a lower time-to-failure and that Witt-Percentile outperforms

Tovar-PPM with tasks failing faster.

Figure 8c depicts the distribution of task failures aggregated by task type. Witt-Wastage exhibits the highest median number of task failures followed by Witt-LR and Sizey. The methods' whiskers and the difference between the first and the third quartile indicate a high variability in the number of failures. Witt-Percentile and Tovar-PPM exhibit a comparable low number of failures, which is anticipated given that the Witt-Percentile method employs a highly conservative 95th percentile predictor and Tovar-PPM uses a very conservative failure handling. Because workflow defaults are user-defined estimates set to prevent task failures, task failures do not occur.

Figure 8d exhibits the aggregated task runtimes without considering the memory wastage. Here, the Workflow-Presets achieve the lowest aggregate. This is expected as we did not observe any task failures, resulting in no mandatory task

TABLE II: Aggregated memory wastage over time for all workflows and methods evaluated.

Method	methylseq	chipseq	eager	rnaseq	mag	iwd
Sizey	631.62	79.38	678.19	43.62	251.05	0.36
Witt-Wastage	3565.11	214.60	491.16	176.39	323.62	0.55
Witt-LR	988.90	136.33	3585.19	57.91	301.00	2.94
Tovar-PPM	4080.60	211.02	624.14	195.26	309.36	16.70
Witt-Percentile	4372.19	94.70	860.16	128.90	309.81	1.44
Workflow-Presets	22596.14	260.61	2304.53	1238.62	1955.01	15.86




Fig. 9: Time required to train Sizey for full retraining and incremental retraining, including model hyperparameter optimization. y-axis uses a logarithmic scale.

restarts. Sizey achieves the second-lowest aggregated task runtime, slightly followed by Witt-Percentile. The Witt-Wastage method shows the highest runtime, which can be explained by the high number of average task failures.

Table II shows the wastage on a workflow basis for all methods and experimental workflows. We can observe that in five out of six workflows, Sizey achieves the lowest wastage over time and a considerable improvement to the best-performing competitor. The highest relative difference compared to the best-performing baseline can be observed for the methylseq workflow, showing a reduction of memory wastage of 36.13%, followed by the iwd workflow with a reduction of 34.55%. Notably, the performance of the baselines fluctuates, and no baseline is consistently second best after Sizey. Four out of five times, the Workflow-Presets result in the highest wastage over time; for the chipseq workflow, Workflow-Presets outperform Witt-LR.

D. Sizey Model Insights

Doing a complete retraining using Sizey to predict a task’s memory consumption, including hyperparameter optimization, we observed a median training time of 1.09 seconds. Using an incremental approach, we achieved a reduction of 98.39% with a median training time of 17.5 milliseconds. Figure 9 illustrates these findings and demonstrates that the training time for full retraining is comparable across all workflows.




Fig. 10: Impact of the parameter alpha on wastage over time in GBh for two rnaseq tasks.




Fig. 11: Proportion of model classes selected by Sizey for the rnaseq workflow.

A comparable behavior can be observed in the context of incremental training. Further, we observed that using Sizey with incremental training increased its median wastage over time by 6.1%.

Figure 10 exemplifies the impact of alpha when Sizey is executed with an alpha value ranging from 0 to 1.0 for two rnaseq workflow tasks. We can observe that the FastQC task shows a tendency for wastage to be lower for a lower alpha parameter. In contrast, the MarkDuplicates task shows the opposite pattern with decreasing wastage as alpha increases. We believe this observation is due to the definition of our efficiency score, which favors models with lower predictions and penalizes possible overprediction outliers to achieve more accurate predictions. However, this may also lead to more task

failures as it favors lower predictions. Also, in general, the additional data we have analyzed do not show a clear tendency to choose a higher or lower alpha.

Figure 11 shows the distribution of selected model classes across all experimental workflows when Sizey is run with the Argmax strategy. It can be observed that the MLP regressor is employed for 42.7% of the task memory predictions, followed by KNN with 29.1%, Random Forest with 19.4%, and Linear Regression with 8.8%. The linear model is a good fit when only a few training data points are available, for instance, at the beginning of a workflow execution. Once more data become available, Sizey can switch to more complex models.

E. Discussion

In our experimental results, we observed that Sizey is able to significantly reduce memory wastage compared to all state-of-the-art baselines. While the largest reduction in memory wastage compared to the best-performing baseline is observed for the methylseq workflow with 36.13%, the aggregated reduction compared to the best-performing baseline is about 60.60%. This is because the performance of the baselines varies greatly for each workflow, while Sizey achieves low memory wastage for all workflows.

Our experimental results show that a reduction in memory wastage is accompanied by an increase in task execution time, due to the occurrence of failed tasks. Although the additional time is relatively small, especially when compared to the memory savings, it is a factor to be considered. One approach to reducing task failures and therefore time, while still reducing memory, is to choose a more conservative offset.

Our experiments show that Sizey’s lightweight models allow a quick training time. Using Sizey with an incremental model reduces the training time further while still significantly outperforming the baselines. The analysis of the alpha parameter showed no clear trend. For some workflow tasks, a lower alpha results in less wastage, while for others it will be the other way around. We assume that this is due to the definition of our efficiency score, which is weighted with 1 when an alpha of 1 is chosen, leading to lower memory predictions. In some cases, this can lead to very accurate predictions, while certain more error-prone tasks, such as those with a high variance in memory usage, may fail due to underprediction. Switching between alpha parameters adaptively during workflow execution, as we do with the models, could address this problem and is an idea for future work.

IV. RELATED WORK

First, we discuss related methods focusing on workflow task memory prediction. Second, we discuss related work that is similar to our proposed method but targets resource managers in general without a specific focus on workflows.

A. Workflow Task Memory Prediction

Substantial research has been done on predicting the memory requirements of workflow tasks. We describe seminal work on this problem and compare each with our own method.

Da Silva et al. [33] present a method that uses an online feedback loop to predict resource consumption, e.g., peak memory, of scientific workflow tasks. Its prediction process uses a regression tree, built offline from historical monitoring data, that first classifies the workflow, application, and resource to be predicted. Then, a data subset with high correlation is identified. Their proposed method, Online-M, predicts the ratio between the I/O read and the resource metrics for correlated data and the median for uncorrelated data. The extended method, Online-P [34], checks for a normal and gamma probability distribution in case there is no correlation identified it draws a value from these distributions. Their online method includes updating task estimates upon completion of the workflow task. Our method is also applied online and is able to update the models and thus the prediction results during the runtime of the workflow. However, we use multiple models that are also able to capture non-linear memory consumption patterns and use a more selective failure handling and offsetting strategy to further reduce failed task executions caused by underestimation.

Witt et al. [32] propose a feedback-based resource allocation system for the scheduling of scientific workflows. Their approach uses an online feedback loop to improve resource usage predictions and thus sizes resources for the tasks of a scientific workflow. The authors examine two approaches to estimate the task’s peak memory usage. The first approach, the percentile predictor, predicts the task’s peak memory usage by using certain percentiles, e.g., P50 for the median or P99 for the 99th percentile. The second method is based on a linear regression model and estimates the task’s peak memory consumption based on the relationship between input data size and memory usage. The predictions of the linear model are then offset by the expected difference between the actual and the predicted peak memory usage. Our work addresses a similar problem. However, instead of relying on a single model class that is applied to all tasks, our method dynamically switches between model classes to find the most suitable for each task.

The peak memory allocation problem is also addressed in [18]. Here, again, the authors use the input data size to predict peak memory usage and train a prediction model that minimizes resource wastage instead of the prediction error. They test quantile regression lines and select the parameters of the one with the least wastage. Further, the authors examine different failure-handling strategies. Again, we do not rely only on a single regression model class. Sizey uses multiple regression models and includes dynamic switching on a per-task basis, allowing us to find the best model for heterogeneous tasks.

Tovar et al. [31] present a job sizing strategy for tasks in scientific workflows that either minimizes wastage or maximizes throughput. The authors use the probabilities of peak resource values from the respective historical task traces to initialize a task. Their goal is to minimize the sum of probabilities of resource peaks where the resource peak is greater than the allocated resource value. The scope of our paper is similar since we also address the sizing of memory for tasks. However,

we also address task heterogeneity by including multiple parallel predictors while dynamically selecting the one with the best prediction results for each task. Additionally, we propose a more selective retry strategy to avoid high resource wastage once a task fails.

Tovar et al. [17] propose another domain-specific approach that, instead of sizing the memory for tasks, aims to split the tasks into subtasks, matching a specific memory requirement. Without historical task executions, the highest observed memory consumption is assigned to the first five tasks. Upon task failure due to insufficient memory, the task is split into smaller tasks and resubmitted. The initial size of a task is determined by a linear correlation between chunk size and task resource usage. Contrary to our and the other presented approaches, the author’s method does not directly predict the memory but rather the input size a task needs to consume a certain amount of memory. The viability of such an approach depends on the structure and divisibility of the task input data and is thus not generally applicable to black-box tasks used in scientific workflows.

Bader et al. [35] proposed two reinforcement learning methods based on gradient bandits and Q-learning. The object of their reinforcement learning methods is the minimization between allocated and used memory while avoiding task failure. The authors use no offsetting technique, as the reward functions implicitly discourage the agents from underestimating the resources. One drawback of the proposed reinforcement learning methods is that they do not incorporate the dependency between task input size and resource usage, leading to higher wastage for tasks with fluctuating memory usage.

Bader et al. [16] also presented a method to predict the memory consumption of workflow tasks over time. To do this, they used time series monitoring data and predicted the expected task runtime. The expected runtime is then divided into segments. For each segment, they then trained a linear regression model that predicts peak memory, resulting in a step function that models a task’s memory consumption over time. Both the internal runtime prediction model and the segment-by-segment peak memory prediction models incorporate an offset strategy to avoid overall task failure. In contrast to our method, which is a combination of multiple predictors, the resultant prediction function is a piecewise linear regression. The method also requires time-series monitoring data, which may not be available in many workflow systems, and relies on the resource manager’s ability to adjust memory allocations over time, a feature that is widely unavailable.

B. Ensemble Memory Prediction for Resource Managers

There is related work that has focused more generally on memory prediction for cluster workloads. Similar to our workflow memory prediction method, closely related methods employ the best of multiple models.

Tanash et al. [36] proposed a Slurm extension that predicts a job’s memory usage using machine learning models. Their method intercepts the submission process by analyzing the job,

predicting the expected memory usage, and adjusting the job limits. While the prediction is done online, the actual model training is done offline and the model is not updated using completed jobs. Tanash et al. decided to use the decision tree regression as it achieved the best prediction results. Instead of selecting a single model to apply globally, Sizey dynamically selects the best performing model for each task. In addition, Sizey works in a full online fashion, incorporating and updating the internal model at runtime to improve prediction accuracy by either updating the current model or even switching to a different model if it performs better with more data.

Rodrigues et al. [37] present a memory prediction extension for the popular LSF HPC resource manager. Their predictive method trains multiple machine learning models based on online and offline data collection and uses a sliding window to focus on the most recent submitted and finished jobs. The machine learning models included are Support Vector Machines, Random Forests, Multilayer Perceptrons, and k-Nearest Neighbors. Instead of addressing it as a regression problem, they discretize the memory predictions and turn it into a classification problem with memory sizes of 512MiB. We approach this as a regression problem, which allows for more fine-grained predictions that can result in less wastage. Further, our method allows switching between a weighting mechanism as presented by Rodrigues et al. and an Argmax strategy where the predictor with the highest resource allocation quality score is used for the prediction. Our results show that such an Argmax strategy is beneficial in scenarios where the predictors are divergent, and weighting would bias the result. Instead of relying solely on the accuracy of the validation data, we introduced an efficiency score that measures the deviations of the models from each other.

V. CONCLUSION

This paper presented Sizey, a novel online workflow task memory prediction method that employs multiple machine learning models. The method dynamically selects the best-fitting model for each task based on assessing the model’s accuracy and memory prediction efficiency. To this end, Sizey incorporates monitoring information during the workflow execution and updates its prediction models while continuously re-evaluating the models’ performance based on gathered data.

We implemented a prototype of Sizey and compared it to four state-of-the-art methods from the literature, using metrics from six real-world nf-core Nextflow workflows. The results demonstrate that Sizey is able to significantly outperform the baselines, with a reduction in memory wastage of 24.68% compared to the best state-of-the-art baseline. Moreover, our experiments offer insights into the parameter selection of our method, thereby facilitating its deployment in other systems.

We plan to integrate our method into state-of-the-art online scheduling by considering the predicted required memory already during scheduling.

ACKNOWLEDGMENTS

Funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) as FONDA (Project 414984028, SFB 1404).

REFERENCES

- [1] P. Di Tommaso, M. Chatzou, E. W. Floden, P. P. Barja, E. Palumbo, and C. Notredame, “Nextflow enables reproducible computational workflows,” *Nature biotechnology*, vol. 35, no. 4, 2017.
- [2] E. Deelman, K. Vahi, G. Juve, M. Rynge, S. Callaghan, P. J. Maechling, R. Mayani, W. Chen, R. F. Da Silva, M. Livny *et al.*, “Pegasus, a workflow management system for science automation,” *Future Generation Computer Systems*, vol. 46, 2015.
- [3] J. Köster and S. Rahmann, “Snakemake—a scalable bioinformatics workflow engine,” *Bioinformatics*, vol. 28, no. 19, pp. 2520–2522, 2012.
- [4] P. A. Ewels, A. Peltzer, S. Fillinger, H. Patel, J. Alneberg, A. Wilm, M. U. Garcia, P. Di Tommaso, and S. Nahnsen, “The nf-core framework for community-curated bioinformatics pipelines,” *Nature Biotechnology*, 2020.
- [5] F. Lehmann, D. Frantz, S. Becker, U. Leser, and P. Hostert, “FORCE on Nextflow: Scalable Analysis of Earth Observation data on Commodity Clusters,” in *Proceedings of the CIKM 2021 Workshops*, Online, 2021.
- [6] J. Schaarschmidt, J. Yuan, T. Strunk, I. Kondov, S. P. Huber, G. Pizzi, L. Kahle, F. T. Bölle, I. E. Castelli, T. Vegge *et al.*, “Workflow engineering in materials design within the battery 2030+ project,” *Advanced Energy Materials*, p. 2102638, 2021.
- [7] C. S. Liew, M. P. Atkinson, M. Galea, T. F. Ang, P. Martin, and J. I. V. Hemert, “Scientific workflows: moving across paradigms,” *ACM Computing Surveys (CSUR)*, vol. 49, no. 4, pp. 1–39, 2016.
- [8] T. Coleman, H. Casanova, L. Pottier, M. Kaushik, E. Deelman, and R. F. da Silva, “Wfcommons: A framework for enabling scientific workflow research and development,” *Future Generation Computer Systems*, vol. 128, pp. 16–27, 2022.
- [9] J. Bader, K. Styp-Rekowski, L. Doehler, S. Becker, and O. Kao, “Macaw: The machine learning magnetometer calibration workflow,” in *2022 IEEE International Conference on Data Mining Workshops (ICDMW)*, 2022, pp. 1095–1101.
- [10] B. Burns, B. Grant, D. Oppenheimer, E. Brewer, and J. Wilkes, “Borg, omega, and kubernetes,” *Queue*, vol. 14, no. 1, 2016.
- [11] A. B. Yoo, M. A. Jette, and M. Grondona, “Slurm: Simple linux utility for resource management,” in *Workshop on Job Scheduling Strategies for Parallel Processing*. Springer, 2003.
- [12] D. Thain, T. Tannenbaum, and M. Livny, “Distributed computing in practice: the condor experience,” *Concurrency and computation: practice and experience*, vol. 17, no. 2-4, 2005.
- [13] A. Ilyushkin and D. Epema, “The impact of task runtime estimate accuracy on scheduling workloads of workflows,” in *2018 18th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CC-GRID)*. IEEE, 2018, pp. 331–341.
- [14] D. G. Feitelson, *Workload modeling for computer systems performance evaluation*. Cambridge University Press, 2015.
- [15] A. Hirales-Carbajal, A. Tchernykh, R. Yahyapour, J. L. González-García, T. Röblitz, and J. M. Ramírez-Alcaraz, “Multiple workflow scheduling strategies with user run time estimates on a grid,” *Journal of Grid Computing*, vol. 10, no. 2, pp. 325–346, 2012.
- [16] J. Bader, N. Diedrich, L. Thamsen, and O. Kao, “Predicting dynamic memory requirements for scientific workflow tasks,” in *2023 IEEE International Conference on Big Data (Big Data)*, 2023. [Online]. Available: <https://arxiv.org/pdf/2311.08185.pdf>
- [17] B. Tovar, B. Lyons, K. Mohrman, B. Sly-Delgado, K. Lannon, and D. Thain, “Dynamic task shaping for high throughput data analysis applications in high energy physics,” in *2022 IEEE International Parallel and Distributed Processing Symposium (IPDPS)*. IEEE, 2022, pp. 346–356.
- [18] C. Witt, J. van Santen, and U. Leser, “Learning low-wastage memory allocations for scientific workflows at icecube,” in *2019 International Conference on High Performance Computing & Simulation (HPCS)*. IEEE, 2019, pp. 233–240.
- [19] C. Witt, M. Bux, W. Gusew, and U. Leser, “Predictive performance modeling for distributed batch processing using black box monitoring and machine learning,” *Information Systems*, vol. 82, pp. 33–52, 2019.
- [20] M. Bux, J. Brandt, C. Witt, J. Dowling, and U. Leser, “Hi-way: Execution of scientific workflows on hadoop yarn,” in *20th International Conference on Extending Database Technology, EDBT 2017, 21 March 2017 through 24 March 2017*. OpenProceedings.org, 2017, pp. 668–679.
- [21] J. Bader, F. Lehmann, L. Thamsen, U. Leser, and O. Kao, “Lotaru: Locally predicting workflow task runtimes for resource management on heterogeneous infrastructures,” *Future Generation Computer Systems*, vol. 150, pp. 171–185, 2024.
- [22] S. Masoudnia and R. Ebrahimpour, “Mixture of experts: a literature survey,” *Artificial Intelligence Review*, vol. 42, no. 2, pp. 275–293, 2014.
- [23] N. Shazeer, A. Mirhoseini, K. Maziarz, A. Davis, Q. Le, G. Hinton, and J. Dean, “Outrageously large neural networks: The sparsely-gated mixture-of-experts layer,” *arXiv preprint arXiv:1701.06538*, 2017.
- [24] J. A. F. Yates, T. C. Lamnidis, M. Borry, A. A. Valtueña, Z. Fagernäs, S. Clayton, M. U. Garcia, J. Neukamm, and A. Peltzer, “Reproducible, portable, and efficient ancient genome reconstruction with nf-core/eager,” *PeerJ*, vol. 9, p. e10947, 2021.
- [25] P. d. B. Damgaard, N. Marchi, S. Rasmussen, M. Peyrot, G. Renaud, T. Korneliussen, J. V. Moreno-Mayar, M. W. Pedersen, A. Goldberg, E. Usmanova *et al.*, “137 ancient human genomes from across the eurAsian steppes,” *Nature*, vol. 557, no. 7705, 2018.
- [26] J. L. Green, R. E. Osterhout, A. L. Klova, C. Merkwirth, S. R. McDonnell, R. B. Zavareh, B. C. Fuchs, A. Kamal, and J. S. Jakobsen, “Molecular characterization of type 1 ifn-induced cytotoxicity in bladder cancer cells reveals biomarkers of resistance,” *Molecular Therapy-Oncolytics*, vol. 23, pp. 547–559, 2021.
- [27] F. H. Karlsson, V. Tremaroli, I. Nookaew, G. Bergström, C. J. Behre, B. Fagerberg, J. Nielsen, and F. Bäckhed, “Gut metagenome in european women with normal, impaired and diabetic glucose control,” *Nature*, vol. 498, no. 7452, pp. 99–103, 2013.
- [28] S. J. Baumgart, E. Nevedomskaya, R. Lesche, R. Newman, D. Mumberg, and B. Haendler, “Darolutamide antagonizes androgen signaling by blocking enhancer and super-enhancer activation,” *Molecular oncology*, vol. 14, no. 9, pp. 2022–2039, 2020.
- [29] O. Thompson, F. von Meyenn, Z. Hewitt, J. Alexander, A. Wood, R. Weightman, S. Gregory, F. Krueger, S. Andrews, I. Barbaric *et al.*, “Low rates of mutation in clinical grade human pluripotent stem cells under different culture conditions,” *Nature communications*, vol. 11, no. 1, p. 1528, 2020.
- [30] T. Rettelbach, M. Langer, I. Nitze, B. M. Jones, V. Helm, J.-C. Freytag, and G. Grosse, “From images to hydrologic networks—understanding the arctic landscape with graphs,” in *Proceedings of the 34th International Conference on Scientific and Statistical Database Management*, 2022, pp. 1–10.
- [31] B. Tovar, R. F. da Silva, G. Juve, E. Deelman, W. Allcock, D. Thain, and M. Livny, “A job sizing strategy for high-throughput scientific workflows,” *IEEE Transactions on Parallel and Distributed Systems*, vol. 29, no. 2, pp. 240–253, 2017.
- [32] C. Witt, D. Wagner, and U. Leser, “Feedback-based resource allocation for batch scheduling of scientific workflows,” in *2019 HPCS*. IEEE, 2019.
- [33] R. F. Da Silva, G. Juve, E. Deelman, T. Glatard, F. Desprez, D. Thain, B. Tovar, and M. Livny, “Toward fine-grained online task characteristics estimation in scientific workflows,” in *Proceedings of the 8th Workshop on Workflows in Support of Large-Scale Science*, 2013, pp. 58–67.
- [34] R. F. Da Silva, G. Juve, M. Rynge, E. Deelman, and M. Livny, “Online task resource consumption prediction for scientific workflows,” *Parallel Processing Letters*, vol. 25, no. 03, 2015.
- [35] J. Bader, N. Zunker, S. Becker, and O. Kao, “Leveraging reinforcement learning for task resource allocation in scientific workflows,” in *2022 IEEE International Conference on Big Data (Big Data)*. IEEE, 2022.
- [36] M. Tanash, B. Dunn, D. Andresen, W. Hsu, H. Yang, and A. Okanlawon, “Improving hpc system performance by predicting job resources via supervised machine learning,” in *Proceedings of the Practice and Experience in Advanced Research Computing on Rise of the Machines (learning)*, 2019, pp. 1–8.
- [37] E. R. Rodrigues, R. L. Cunha, M. A. Netto, and M. Spriggs, “Helping hpc users specify job memory requirements via machine learning,” in *2016 Third International Workshop on HPC User Support Tools (HUST)*. IEEE, 2016, pp. 6–13.