Tracking and Analyzing the Evolution of Provenance from Scripts

João Felipe Pimentel[1(✉)], Juliana Freire[2], Vanessa Braganholo[1], and Leonardo Murta[1]

[1] Universidade Federal Fluminense, Niterói, Brazil
{jpimentel,vanessa,leomurta}@ic.uff.br
[2] New York University, New York, USA
juliana.freire@nyu.edu

Abstract. Script languages are powerful tools for scientists. Scientists use them to process data, invoke programs, and link program outputs/inputs. During the life cycle of scientific experiments, scientists compose scripts, execute them, and perform analysis on the results. Depending on the results, they modify their script to get more data to confirm the original hypothesis or to test a new hypothesis, evolving the experiment. While some tools capture provenance from the execution of scripts, most approaches focus on a single execution, leaving out the possibility to analyze the provenance evolution of the experiment as a whole. This work enables tracking and analyzing the provenance evolution gathered from scripts. Tracking the provenance evolution also helps to reconstruct the environment of previous executions for reproduction. Provenance evolution analysis allows comparison of executions to understand what has changed and supports the decision of which execution provides better results.

1 Introduction

The life cycle of script-based experiments is usually composed of three main phases [12]: establishing hypotheses and coding scripts that enact the programs involved in the experiment; running the script over input data, which represent a specific context or population for the experiment; and analyzing the produced results through visualizations or queries to confirm the research hypotheses. However, the results of the latter phase may motivate the repetition of the cycle. For instance, when a trial (i.e., one execution of the experiment) is inconclusive, scientists repeat the cycle after adapting the script or changing the programs. When a hypothesis is confirmed for a restrict population, scientists repeat the experiment for a broader population by changing the input data. Similarly, when a hypothesis is refuted for a broad population, scientists restrict the population and repeat the cycle also by changing the input data. Moreover, some scientists design experiments considering multiple inputs or variable programs and the experiment execution entails many trials at once via parameter sweeping. Thus, script, programs, data, and the execution environment evolve over time as a natural consequence of the experimental process.

In the last decade some approaches emerged for capturing provenance from experiments encoded in scripts [2, 3, 11, 14, 19]. The captured provenance usually

includes the script structure with its functions and variables, all input data, intermediate data, output data, the required libraries, the environment characteristics (computer architecture, operating system, etc.), and the execution flow of the trial (function activations, variable assignments, etc.). However, these approaches either do not track the evolution of the experiment between trials or rely on external tools for such tracking. In both cases, the scientists are limited to intra-trial queries, not being able to contrast the provenance of two trials or to visualize the difference among trials' provenance.

Understanding and visualizing a single trial through intra-trial queries is not enough for the analysis of the whole experiment. To support this claim, we refer to a set of questions related to experiment evolution analysis, which were obtained and adapted from the first Provenance Challenge[1] and ProvBench workshops[2]: Q1 (see footnote 1): if a scientist has executed an experiment twice, but has replaced some procedures in the second trial, what are the trial differences? Q2[3]: comparing multiple executions according to their parameters, what are the differences on execution behavior? Q3[4]: how differences in the input data relate to differences in the output values? Q4 (see footnote 4): using historical provenance, which parts of the execution fail frequently? Q5[5]: which trials are related to a given trial? Q6 (see footnote 5): a given trial was derived from which trial? Q7[6]: what are the available trials, and what are their durations? Q8 (see footnote 6): how many trials are associated to a given source code? Q9 (see footnote 6): how many trials present failures?

To be able to answer these questions, in this work we propose a version model that supports tracking and analyzing the experiment provenance as a whole, considering its multiple trials. This model also allows us to restore any past trial, thus enabling reproducibility. Moreover, our version model supports comparison of different trials for analysis. As a proof of concept, we implemented our version model on top of noWorkflow [14, 16, 17]. noWorkflow is an approach that automatically collects provenance from Python scripts without requiring any modifications on the source code of the experiment. For every trial, noWorkflow generates an identifier and all provenance collected during the execution is stored in a database related to that identifier. Provenance collected by noWorkflow contains function activations (calls) with parameters, variable values, returned values, duration, and caller; imported modules with their versions; environment variables; and all the files accessed during the trial, including source files, module files, and input files, intermediate files, and output files.

This paper is organized in six sections, besides this introduction. Section 2 discusses related work. Section 3 presents our approach to track evolution, analyze provenance, and compare trials. Section 4 presents the implementation details on top of noWorkflow. Section 5 shows the evaluation of our work using the aforementioned

---

[1] http://twiki.ipaw.info/bin/view/Challenge/FirstProvenanceChallenge.

[2] https://sites.google.com/site/provbench/home/provbench-provenance-week-2014.

[3] https://github.com/provbench/Swift-PROV.

[4] https://github.com/provbench/CSIRO-PROV.

[5] https://github.com/provbench/VisTrails-PROV.

[6] https://github.com/provbench/Wf4Ever-PROV.

questions. Finally, Sect. 6 concludes the paper summarizing the contributions and discussing future work.

2   Related Work

Work related to our approach can be grouped into three main categories: (i) configuration management tools, (ii) script-based provenance tools, and (iii) workflow-based provenance tools. Many configuration management tools, such as Git, track the evolution of software through versioning [6]. These tools allow developers to inform the files they want to track and provide different mechanisms for querying the history, such as *bisect*, *blame*, and even some simple *lookups* on previous versions. Developers can also use external visualization tools to have a broader view of source code evolution [5]. Although generic and fast, these tools capture only prospective provenance [8] at coarse grain, when used to version experiment scripts. Thus, they do not track the inner structures of files, the evolution of computations that indicate which input files and parameters actually influence each output files, nor the multiple intermediate states of files. In other words, they do not capture fine-grained prospective provenance [8] nor retrospective provenance [8].

Some approaches can be used to capture provenance from scripts. YesWorkflow [13] captures prospective provenance from scripts through annotations. StarFlow [2] and RDataTracker [11] collect provenance from scripts through dynamic analysis and annotations. Bochner et al. [3] collect provenance from scripts using a library to connect to a remote server and send annotated provenance data. Tariq et al. [19] collect provenance from code compiled with a LLVM compiler. Stamatogiannakis et al. [18] perform dynamic taint analysis on binary files to capture provenance. noWorkflow [14] collects provenance from scripts without requiring any modifications on the script. Most of these approaches capture execution provenance (i.e., retrospective provenance) [14] with intermediate data, and support querying and visualizing provenance during analysis. However, they do not provide mechanisms to compare and contrast different trials. An outstanding exception in this category is Sumatra [7]. It stores each trial in a configuration management tool (either Git or Mercurial) and allows users to tag them and to compare the collected information. However, it does not record the intermediate states of files during execution and is subject to the problems of using configuration management tools for tracking the evolution of experiments.

Finally, workflow-based provenance tools [1, 4] track provenance from scientific experiments. Some tools, such as Vistrails [4] and Kepler [1], not only track the provenance, but also track the workflow evolution and offers all the data for users to analyze it. They also allow users to restore past versions of workflows and re-execute them. Although attractive in terms of features, these tools require converting script-based experiments into workflows, which is not an option for many scientists for different reasons. This motivates the creation of a version model for script-based experiments, detailed in the next section.

3 Script-Based Provenance Evolution

Supporting evolution analysis of experiments requires the provenance-capturing tool to be evolution-aware. This can be achieved through versioning. Versioning enables tracking the evolution of the experiment and also navigating on the evolution history, allowing the user to restore previous versions, if needed. Additionally, such an evolution-aware provenance capture system should provide a way to compare different trials on the history. In Sect. 3.1, we propose a version model for provenance collected from scripts. In Sect. 3.2, we propose techniques to compare provenance from different trials.

3.1 Version Model

Conradi and Westfechtel [6] state that a version model should define the organization of the *version space* (i.e., how a product is versioned) and the interrelation of the *product space* (i.e., how a product is structured) and the version space. We define our *product space* as an experiment, containing its scripts, data, execution traces, etc. The entry point of our product space is the main script of the experiment. From this script, we recursively capture imported modules, accessed files during execution, and the execution provenance. Thus, we have scripts (including imports), input files, intermediate files, and output files as *file objects*. We identify file objects solely by their path within the experiment directory.

File objects describe the structure of the experiment: that is, all files needed by the experiment, which includes the script itself (definition provenance [14]), imported modules (deployment provenance [14]), and accessed (read/write) files (execution provenance [14]). On the other hand, we also have logical provenance information that is not stored in files: functions called during execution, parameters values, variable values, etc. In our product space, we have a special object called *logical object* that contains all the aforementioned logical provenance information. This way, we can say that our product space is composed of multiple file objects and one logical object.

Our *version space* [6] has two levels of versioning: *trial version* (i.e., the trial *id*) and *file object version*. Trial versions represent the state of the experiment in terms of file object versions read or written within each trial, together with the logical object version produced by the trial. On the other hand, file object versions represent the state of file objects at each file access during the whole experiment execution (throughout all trials). File object versions may contain extra attributes (metadata) besides the state of file objects: modules may have their semantic versions declared by developers (e.g., 3.5.1), files may have their moment of opening and opening mode (read/write), etc.

We apply this distinction between trial versions and file object versions because scripts can write to some file objects more than once, generating more than one version of the file object within a single trial. Due to this distinction, our version space supports restoring trial versions as a whole, with all input file objects, or specific file object versions (e.g., an intermediate version of a file object). However, to restore a specific file object version, users should inform which object they want to restore individually

and in which moment (i.e., by indicating a timestamp, the file content hash code, or its access position in a sequential list by timestamp).

While we associate file objects to both version concepts (trial version and file object version), we associate logical objects only to trial versions, because they are unique for each trial and already contain all execution steps (i.e., each function activations, each variable state, etc.) within a trial. Nonetheless, restoring a trial version does not restore the logical object of that trial, as it is not a tangible object, even though it is still useful for auditing or reproducing a trial.

Figure 1 presents an example of this version model with two trial versions for an experiment, where the user only edited "experiment.py" and added "converter.py" before executing the second trial. Circles represent object versions and dotted squares represent trial versions. Note that the file "warp.warp" has four file object versions in Trial 1, and those versions were written four times, and read four times. Note also that Trial 1 does not have file object versions for "converter.py", "atlas-x.ppm", and "atlas-x.jpg" because file object versions refer to the state of files at their access and Trial 1 did not access these files. Equivalently, there is no file object version for "atlas-x.gif" at Trial 2, since Trial 2 did not access it. Moreover, we can observe that both trials accessed the same file object version of "external.py" and "anatomy1.img" and that the user edited "experiment.py" after Trial 1. The logical object, on the other hand, has a single and unique version on each trial, since it contains runtime data such as function activations, start and finish times, variable values, etc. This kind of data is already time-sensitive, not demanding an extra layer of versioning.
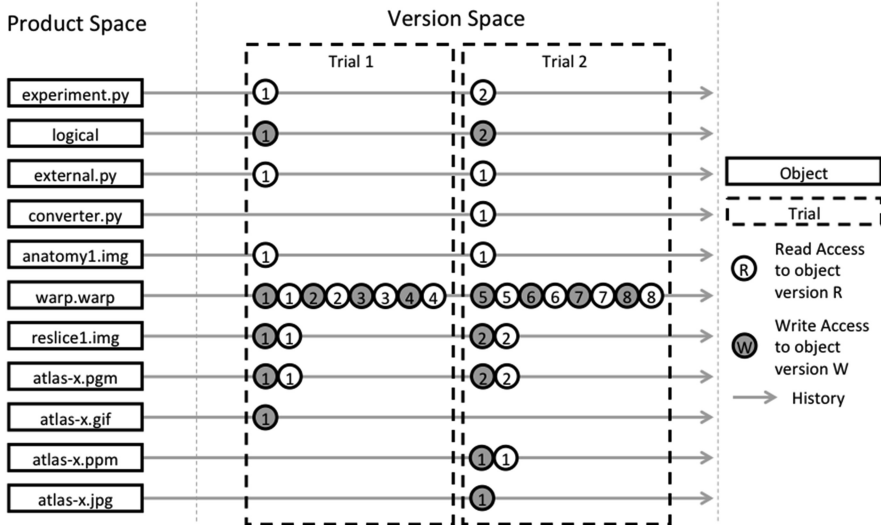

Fig. 1. Version model example

As mentioned before, users can use trial versions to restore states of the experiment. The main goal when restoring a trial is for reproducing it. For this reason, restoring Trial 1 would only restore the files "experiment.py", "external.py", and "anatomy1.img"

(all at version 1). In addition, it would remove "warp.warp", "reslice1.img", "atlas-x.pgm", and "atlas-x.gif", because these files did not exist prior to Trial 1. However, restoring Trial 2 would restore "experiment.py" (at version 2), "external.py" (at version 1), "converter.py" (at version 1), "anatomy1.img" (at version 1), "warp.warp" (at version 4), "reslice1.img" (at version 1), and "atlas-x.pgm" (at version 1); and it would remove "atlas-x.ppm" and "atlas-x.jpg". Note that it would not touch "atlas-x.gif", since Trial 2 has not accessed it. Note also that it would restore "warp.warp", "reslice1.img", "atlas-x.pgm" because the state of these files before Trial 2 is equal to the state after Trial 1.

Trial versions not only identify the state of an experiment, but also track its evolution. In the example of Fig. 1 we can see that Trial 2 is an evolution of Trial 1, because it was an execution of "experiment.py" after Trial 1. If the user executes a new script, "experiment2.py" (that is in the same directory as "experiment.py"), she would have a new trial, with version 3, but it would not be an evolution of Trial 2. However, if she executes again "experiment.py", she would have Trial 4 based on Trial 2.

We also provide a special type of trial version to avoid losses on the restore operation. If a user changes the content of "experiment.py" but instead of running a new trial using the modified script, she restores Trial 2, she would lose all changes. To avoid these losses, we create a special "backup" trial with the current content of all file objects in the last version (i.e., file objects edited after Trial 4). In this case, we would have Trial 5 as a backup trial, with contents of "experiment.py", "external.py", "converter.py", "anatomy1.img", "warp.warp", "reslice1.img", "atlas-x.pgm", "atras-x.ppm", and "atlas-x.jpg". At least one of these files should be different from the ones of Trial 4 for the backup trial to be created.

After restoring Trial 2, if a user runs Trial 6, it would be based on Trial 2. We keep track of this information by storing the base version of each trial. Before Trial 6, we had the base version restored to 2. After running Trial 6, we update the base version to 6. This allows our version model to track the evolution in a non-linear way. In fact, by considering the evolution of "experiment.py", as presented in Fig. 2, it is possible to see two *branches* of Trial 2: one that goes from Trial 2 to Trial 4, and another that goes from Trial 2 to Trial 6. A branch is a sequence of trials that were executed in parallel to other sequences of trials. Branches can have either a common ancestor to other branch or no ancestor at all. In this case, Trial 2 is the common ancestor of both branches, and Trial 4 and 5 belong to the same branch.

Figure 2 presents an evolution history bigger than what we described so far. In the figure, Trials 1, 2, 4, 5, 6, and 7 are related to "experiment.py" and Trials 3, 8, 9, and 10


Fig. 2. Evolution history. Nodes represent trial versions (Color figure online)

are related to "experiment2.py". We represent trials that did not finish (i.e., halted due to an error) as red nodes and backup trials as yellow nodes. According to the Figure, Trial 7 did not finish and Trial 5 is a backup trial. In addition, after getting an error on "experiment.py" execution (i.e., Trial 7), the user executed "experiment2.py" (Trial 8). Then she restored Trial 3 and executed "experiment2.py" again, creating a new branch. Finally, she restored Trial 8 and executed "experiment2.py", generating Trial 10.

Note that we have two branches of "experiment.py" and two branches of "experiment2.py" in the end. Users can use branches to try different processes for their experiments and to execute their experiment on the same code base, but with different input files or parameters.

Figure 3 presents an UML representation of our version model. The gray classes, *FileObject* and *LogicalObject*, belong to the product space. The white classes, *FileObjectVersion, LogicalObjectVersion, TrialVersion, RestoreVersion, SourceCode Version,* and *FileAccessVersion*, belong to the version space. Note that a *TrialVersion* has one or more *FileObjectVersion*. This composition represents all file object versions accessed (read or written) in a trial. However, when restoring a trial version, only a subset of them is actually overwritten. We identify these by the RestoreVersion association class. Note also that a trial version always has at least one file object version (and corresponding file object): its main script.
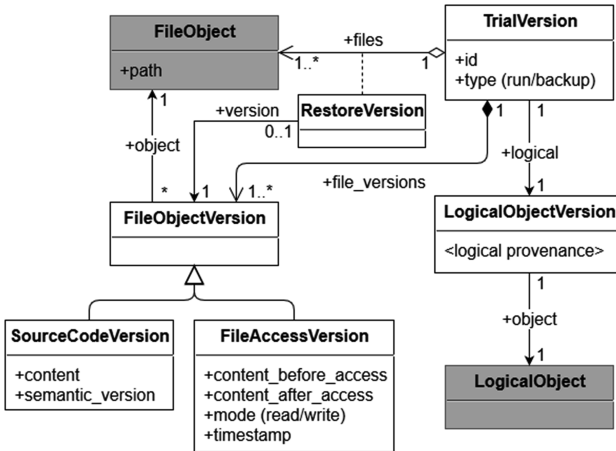

Fig. 3. UML representation of the version model

3.2 Comparing Trials

After tracking the evolution history of an experiment through its provenance, we can compare trials in the history. We compare provenance of file objects in two trials by comparing all their file object versions, with their extra attributes. For example, when comparing source codes, we check not only if their content has changed, but also if the declared version has changed. This way, we can identify that a module content has changed because the user upgraded it from version "1.0.1" to "1.0.2". During the

comparison of changes, we ignore attributes that are always different, such as the moment of opening input and output files.

In addition to reporting changes on file object versions that exist in both trials, we also report file objects that exist in the first trial but do not exist in the second one as *removals* and file objects that exist in the second trial but do not exist in the first one as *additions*. Hence, when we compare Trial 1 and Trial 2 of Fig. 1, we have "atlas-x.gif" as a removal, because Trial 2 did not access it. We also have "converter.py", "atlas-x. jpg", and "atlas-x.ppm" as additions because only Trial 2 accessed these files. Finally, we have "experiment.py" as a change, because it has changed to import "converter.py".

In the previous example, we also have the former versions of "warp.warp", "reslice1.img", and "atlas-x.pgm" (the ones on Trial 1) as removals; and the later versions of the same files (the ones on Trial 2) as additions. This occurs because input and output files can have more than one state (file object version) during a trial and it is not possible to identify them only by their path. Thus, we identify them considering also their content before and after the access. Since these files did not exist before Trial 1, we identify them as different file objects versions than before Trial 2, as at that moment their content is the last version written on Trial 1 (e.g., content just before version 5 of "warp.warp" is its content at version 4).

As our version model groups the entire logical provenance in a special object (the logical object), comparing it is specific for each implementation. Thus, we describe our logical provenance comparison in Sect. 4.

4   Implementation in noWorkflow

We implemented the proposed approach on top of noWorkflow [14]. noWorkflow transparently captures provenance from Python scripts by running `now run <script>`. After running the script, it creates a unique *trial id* to identify the collected provenance and stores the provenance in two databases: a content database for storing file objects and a relational database for storing logical objects and metadata of file objects. noWorkflow uses hash codes to associate metadata of file objects in the relational database to their actual content in the content database.

noWorkflow captures the main script, imported modules, and file accesses as *file object provenance*. As *logical provenance*, noWorkflow captures trial start time, finish time, command line, success status (i.e., indication if the trial finished successfully), environment variables, function activations (calls) with parameters, returned values, duration, caller, variables, and variable dependencies.

We support restoring previous trial versions through the command `now restore <trial_id>`. This command restores the trial version as described in Sect. 3.1. Even though noWorkflow captures source code of external modules, this command only restores local modules to avoid breaking the Python installation. It is possible to filter the restore command to restore only the main experiment script, input files, or local modules.

For visualizing the evolution history, we offer the command `now history`. It supports filtering by experiment script or trial status (e.g., finished, unfinished, or backup).

Trials can be compared by the command `now diff <trial_id1>` `<trial_id2>`. This command has options to specify what should be compared. For instance, `-f` compares file access to input and output files. We use the techniques described in Sect. 3.2 to compare file objects. For comparing equality of contents, we use only hash codes, instead of looking for all differences within files. To understand differences between file object versions, users can run external diff tools over the file versions. The diff command also compares logical provenance. Since most trials have at least start time, finish time, command line, and success status as logical provenance, we always compare these attributes when running this command. With the option `-e`, we support comparison of environment attributes (i.e., part of logical provenance) through a similar process that identifies changes, additions, and removals. Figure 4 presents an excerpt of a brief diff between file accesses from Trial 1 and 2. Note that before presenting file access diff, it presents the diff of these attributes.

```
$ now diff 1 2 -f --brief
[now] trial diff:
  Start changed from 2016-02-11 04:49:09.008354
                  to 2016-02-11 04:49:09.898675
  Finish changed from 2016-02-11 04:49:09.536409
                  to 2016-02-11 04:49:10.276422
  Duration text changed from 0:00:00.528055 to 0:00:00.377747
  Code hash changed from cd1be11a2308ab217327a7d361138cb7f6c25106
                    to 2f637ec102961a7677e3f629ab88612d8875f04f
  Parent id changed from <None> to 1

[now] Brief file access diff
[Additions]           | [Removals]              | [Changes]
(rb) atlax-x.ppm      | (w) atlax-x.gif (new)   |
(w) atlax-x.jpg (new) | (w) atlax-x.pgm (new)   |
(w) atlax-x.pgm       | (w) reslice1.img (new)  |
(w) atlax-x.ppm (new) | (wb) warp.warp (new)    |
(w) reslice1.hdr      | ...                     |
(wb) warp.warp        |                         |
...                   |                         |
```

Fig. 4. Brief diff between file access from Trial 1 and 2

The process of comparing function activations is a bit more complex. First, noWorkflow exports function activations of both trials to a graph format. Next, it transforms both graphs into lists of nodes. Then, it applies the longest common subsequence (LCS) algorithm [9] over the lists. Finally, it recombines nodes into a graph that displays common nodes, additions, and removals. The idea behind using LCS is that activations are in sequence and the generated graph keeps the activation order at some degree. Thus, it is possible to use the LCS and match common nodes.

Currently, we do not compare function activations with the diff command. For comparing them, we provide a visualization tool that can be accessed by running the command `now vis`. The visualization tool also presents the history graphically (shown in Fig. 2). It is also possible to use Jupyter Notebook to visualize the diff and history [17].

Figure 5 presents activation graphs of Trial 1 and Trial 2 and their comparison. Nodes represent function activations and their colors represent their duration in a traffic light scale, where red fills represent the slowest activations and green fills represent the

fastest ones. The trial script is an activation itself and it is pointed out by a straight arrow. In this case, "experiment.py" is the trial script. In the graph, black arrows represent the start of activations; blue arrows represent sequence of calls within activations; and dashed arrows represent returns. In the graph comparison, nodes and arrows with black borders exist in both trials; nodes and arrows with red borders exist only on Trial 1; and nodes and arrows with green borders exist only on Trial 2. Note that "convert" activations exist only on Trial 1, while "pgmtoppm" and "pnmtojpeg" activations exist only on Trial 2. Trial 2 has also an activation representing the import "convert.py". Moreover, nodes that exist in both trials show colors side-by-side to easy comparison. For instance, one can easily notice that slice_convert was slightly faster in Trial 1 than in Trial 2.
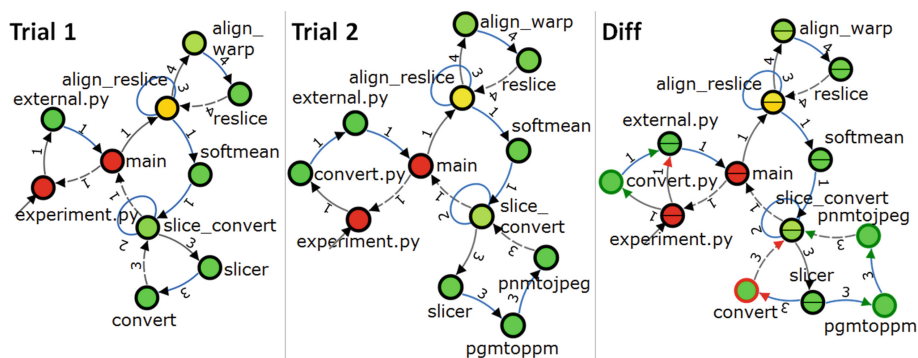

Fig. 5. Activation graphs of Trial 1 and 2, and their comparison (Color figure online)

5 Evaluation

We evaluate our approach by presenting how noWorkflow answers the questions related to provenance evolution listed in Sect. 1. We answered those questions using the example described in Sect. 3.1. This example is in fact the workflow of the first Provenance Challenge implemented in Python with procedures implemented as "dummies". The full history of this experiment can be obtained on noWorkflow by running `now demo 3`.

Q1: if a scientist has executed an experiment twice, but has replaced some procedures in the second trial, what are the trial differences? Q2: comparing multiple executions according to their parameters, what are the differences on execution behavior? Figure 5 presents the comparison of Trial 1 and Trial 2 activation graphs. It is possible to see that "convert" was replaced by "pgmtoppm" and "pnmtojpeg". To compare execution behaviors according to parameters, we can compare trials that share the same code base, but have different parameters.

Q3: how differences in the input data relate to differences in the values? We can use the `now diff -f` command to compare file accesses of trials (as shown in Fig. 4). This

command compares input data, output data, and arguments. Thus, it is possible to get the differences on inputs, and compare them to output values by restoring them.

Q4: using historical provenance, which parts of the execution fail frequently? A SQL query can look for failures on all trials. If we specify that the return value "−1" of a function activation represents a failure, the following query would return the most frequent failures on all trials combined:

```
SELECT name, count(name) AS c FROM function_activation
WHERE return_value = "-1" GROUP BY name ORDER BY c DESC;
```

Q5: which trials are related to a given trial? Q6. a given trial was derived from which trial? Q7. what are the available trials, and what are their durations? Q8. how many trials are associated to a given source code? Q9. how many trials present failures? Looking at the Evolution History (as shown in Fig. 2), it is possible to see both the ancestor of a given trial and all trials that derive from it. The evolution history also presents all available graphs. To get their duration, a user can activate tooltips on `now vis` or Jupyter Notebook and access trial information, including its duration. To get all trials associated to a given source code, we can filter the history to a specific script. Finally, the history graph presents trials with failures as red nodes.

6   Conclusion

In this paper, we presented a novel approach for tracking and analyzing the evolution of provenance collected from scripts. With our approach, a scientist can precisely record all provenance information related to each trial, switch between trials and adapt/reproduce specific trials, and compare trials. We implemented our approach as an extension to *noWorkflow*, which is available as open source software at https://github.com/gems-uff/noworkflow.

While the proposed version model is suitable for any tool that collect multiple versions of files during the execution of a trial, it may impact the execution time of experiments. This occurs because our version model requires the collection to be performed at runtime, reading file contents multiple times during a trial. Additionally, our current implementation captures and stores provenance versions at fine-grain. On the one hand, this provides a powerful support for further analysis. On the other hand, this is known to compromise scalability in terms of execution time and storage space [15]. In particular, storing many different versions of fine-grained data can be wasteful in some cases. This motivates the need for optimization techniques that attempt to balance storage and re-computation costs. We plan to address this issue in the future. Another limitation of the implementation is that we restore only local modules during the restore operation. If the user updates an external module, the experiment reproduction may produce different results. We intend to use virtual environments to avoid this issue.

We also intend to explore alternatives on detecting file object changes, and to work on better algorithms to compare activation graphs. We already started looking for existing graph matching techniques [10]. Additionally, we plan to work on a semantic versioning for trials that encodes the intention of evolution, and to improve *logical*

provenance comparison on noWorkflow to compare not only activation graphs and environment variables, but also variables, variable dependencies, parameters, and return values on activations. Finally, we foresee the elaboration of different formats for provenance visualization that would help on comparing trials.

References

1. Altintas, I., et al.: Kepler: an extensible system for design and execution of scientific workflows. In: International Conference on Scientific and Statistical Database Management (SSDBM), Santorini, Greece, pp. 423–424 (2004)
2. Angelino, E., Yamins, D., Seltzer, M.: StarFlow: a script-centric data analysis environment. In: McGuinness, D.L., Michaelis, J.R., Moreau, L. (eds.) IPAW 2010. LNCS, vol. 6378, pp. 236–250. Springer, Heidelberg (2010)
3. Bochner, C., Gude, R., Schreiber, A.T.: A python library for provenance recording and querying. In: Freire, J., Koop, D., Moreau, L. (eds.) IPAW 2008. LNCS, vol. 5272, pp. 229–240. Springer, Heidelberg (2008)
4. Callahan, S.P., et al.: VisTrails: visualization meets data management. In: ACM SIGMOD, Chicago, USA, pp. 745–747 (2006)
5. Collberg, C., et al.: A system for graph-based visualization of the evolution of software. In: ACM Symposium on Software Visualization (SoftVis), New York, NY, USA, p. 77–ff (2003)
6. Conradi, R., Westfechtel, B.: Version models for software configuration management. ACM Comput. Surv. **30**(2), 232–282 (1998)
7. Davison, A.P.: Automated capture of experiment context for easier reproducibility in computational research. Comput. Sci. Eng. **14**(4), 48–56 (2012)
8. Freire, J., et al.: Provenance for computational tasks: a survey. Comput. Sci. Eng. **10**(3), 11–21 (2008)
9. Hunt, J.W., Szymanski, T.G.: A fast algorithm for computing longest common subsequences. Commun. ACM **20**(5), 350–353 (1977)
10. Koop, D., et al.: Visual summaries for graph collections. In: IEEE Pacific Visualization Symposium (PacificVis), pp. 57–64 (2013)
11. Lerner, B.S., Boose, E.R.: Collecting provenance in an interactive scripting environment. In: Workshop on the Theory and Practice of Provenance (TaPP), Cologne, Germany (2014)
12. Mattoso, M., et al.: Towards supporting the life cycle of large scale scientific experiments. Int. J. Bus. Process Integr. Manag. **5**(1), 79–92 (2010)
13. McPhillips, T., et al.: YesWorkflow: a user-oriented, language-independent tool for recovering workflow information from scripts. Int. J. Digit. Curation **10**, 1 (2015)
14. Murta, L., Braganholo, V., Chirigati, F., Koop, D., Freire, J.: noWorkflow: capturing and analyzing provenance of scripts. In: Ludaescher, B., Plale, B. (eds.) IPAW 2014. LNCS, vol. 8628, pp. 71–83. Springer, Heidelberg (2015)
15. Murta, L.G.P., et al.: Odyssey-SCM: an integrated software configuration management infrastructure for UML models. Sci. Comput. Program. **65**(3), 249–274 (2007)
16. Pimentel, J.F., et al.: Fine-grained provenance collection over scripts through program slicing. In: Mattoso, M., Glavic, B. (eds.) IPAW 2016. LNCS, vol. 9672, pp. 199–203. Springer, Heidelberg (2016)

17. Pimentel, J.F.N., et al.: Collecting and analyzing provenance on interactive notebooks: when IPython meets noWorkflow. In: Workshop on the Theory and Practice of Provenance (TaPP), Edinburgh, Scotland (2015)
18. Stamatogiannakis, M., Groth, P., Bos, H.: Looking inside the black-box: capturing data provenance using dynamic instrumentation. In: Ludaescher, B., Plale, B. (eds.) IPAW 2014. LNCS, vol. 8628, pp. 155–167. Springer, Heidelberg (2015)
19. Tariq, D., et al.: Towards automated collection of application-level data provenance. In: Workshop on the Theory and Practice of Provenance (TaPP), Boston, MA, USA (2012)