

International Journal of Computer Science and Mobile Computing



A Monthly Journal of Computer Science and Information Technology

ISSN 2320-088X

IJCSMC, Vol. 4, Issue. 4, April 2015, pg.467 – 471

RESEARCH ARTICLE

Programming Shell Scripts by Demonstration

Jyoti Yadav

Roll No.15902, CS, Department of Computer science, Dronacharya College of Engineering, Gurgaon-123506, India
Email: jyotirao1212@gmail.com

Jyoti Chandel

Roll No. 15039, CS, Department of Applied Computer science, Dronacharya College of Engineering, Gurgaon-123506, India
Email: chandeljyoti1994@gmail.com

Neha Gupta

Roll No. 15061, CS, Department of Applied Computer science, Dronacharya College of Engineering, Gurgaon-123506, India
Email: nehagupta81094@gmail.com

Abstract: Command-line interfaces are heavily used by system administrators to manage computer systems. Tasks performed at a command line may often be repetitive, leading to a desire for automation. However, the critical nature of system administration suggests that humans also need to supervise an automated system's behavior. This paper presents a programming by demonstration approach to capturing repetitive command line procedures, which is based on a machine learning technique called version space algebra. The interactive design of this learning system enables the user to supervise the system's training process, as well as allowing the user and system to alternate control of the learned procedure's execution.

Introduction

A recent study (Kandogan & Maglio 2003) has shown that most system administrators perform their management and troubleshooting tasks via command-line interfaces rather than using the variety of graphical user interfaces at their disposal. Command line interfaces (CLIs) have several advantages for system administration work over graphical interfaces (GUIs). First, they allow tasks to be automated, which is often necessary when performing the same task across multiple machines in a cluster, or when human error while performing a task could lead to costly downtime. Second, they preserve organizational knowledge about how to accomplish a task in a human-readable, executable form. Third, they allow administrators to easily share knowledge (in the form of copied and pasted shell commands) with their colleagues via instant messaging and email. On the other hand, the use of shell scripts to capture procedural knowledge has its drawbacks. Administrators in the study said that they reused old scripts, handed down from previous administrators, without fully understanding what the scripts did. The cost of authoring a

script may be too high, or require too much programming knowledge. The time and effort required to diagnose failures of an automated script may be prohibitive. This paper proposes a programming by demonstration approach to the problem of capturing and automating repetitive system administration procedures. A programming by demonstration system learns how to perform a procedure by observing the user perform the procedure one or more times, directly in the user interface. Given concrete examples of the procedure's execution, the system induces variables, conditionals, and loops in the underlying procedure. It then lets the user execute the learned procedure in order to repeat the task directly in the user interface. We have constructed a system, which we call SMARTshell, that learns Unix command-line procedures by observing the interactions between a user and a terminal. SMARTshell is an adaptive system that learns procedures from human-generated examples, and is capable of refining its behavior based on feedback from the user during the playback process. Underlying SMARTshell is a machine learning algorithm, based on version space algebra. We have chosen the machine learning algorithm carefully to enable the system to respond to user feedback, and provide a user experience in which both the user and the system collaborate to achieve the goal.

SMARTshell implementation

We illustrate the SMARTshell system on a representative scenario: testing and restarting a development server. Imagine a developer who is making changes to a server, and needs to repeatedly verify that his changes have not broken any of the tests in the test suite. Each time he builds a new server, he must run the following steps in a console:

- start up the server, and make note of the port number it chose to start up on
- run the test suite, passing the server's port number as an Argument
- bring up a process listing to determine the server's process id operates like a macro recorder.

Before performing the task, the developer types "smsh start" into a shell window to bring up the SMARTshell recording interface (Figure , lower window). From then on, every command he types into the shell window (top window in the figure) is recorded by SMARTshell. He has the option of annotating each step with human-readable text that explains the command, which will be displayed when the procedure is later played back. After he has completed the task, the developer closes the recording window and SMARTshell saves the procedure for later reuse. At a future point in time, when the developer has to perform the same task again, he starts up SMARTshell in playback mode (Figure). The system indicates that the procedure is four steps long and that the first step is to start up the server. When the developer clicks on the "step" button in the playback interface, the system automatically performs the command in the console above. The user also has the option of modifying the command before it is executed; although not yet implemented, the learning algorithm could use this as feedback that the command it predicted was not the user's intended command. Playback continues on the next steps in the procedure, running the test suite and bringing up a process listing. The kill command in the next step requires as its first argument the process id of the currently-running server, which was printed out as a result of the `ps` command. In this case, SMARTshell guesses that the user wants to run the command "kill

20068”, using the correct process id, despite the fact that when he had recorded the procedure, the user had typed “kill 20002”.

In this case, SMARTshell has correctly guessed that the user wants to extract the process id from the output of the previous command, and uses that as the argument to the current command. The bottom of the playback window shows

If the user accepts this command, he can continue on to the next step in the procedure. If for any reason, this is not the right command, he may either modify it in the playback interface, or ask SMARTshell to try another guess by clicking on the “Try another” button in the interface.

SMARTshell maintains a set of candidate hypotheses, and trying other guesses produces the next most likely hypotheses in that space. Whichever command is accepted, this information can be used to update the learning algorithm for future invocations of the procedure. Once the user is fairly certain that the system has learned the correct procedure, he can automate the remainder of the procedure by clicking on the “Run” button in the playback interface.

Learning shell scripts by demonstration

We have formalized the problem of learning shell scripts by demonstration as a machine learning problem using version space algebra (Lau, Domingos, & Weld 2000; Lau *et al.* 2003). Version space algebra is a method for modeling a machine learning problem by decomposing it into smaller, independently-learnable parts, and combining the results of the subproblems into an answer for the complete. Although version spaces were originally used for binary classification, we note that the above formulation holds for any functions that map from an input to an output. In this case, a hypothesis is a mapping from the domain to the range of the function space, and the version space consists of functions that correctly produce the output label given the input label.

Version space algebra consists of a framework for combining simple version spaces into composite ones, through the three operators *union*, *join*, and *transform*. The union allows multiple simpler version spaces to be combined into a single version space that contains the union of their hypotheses. The join enables the cross product of two spaces to form a new space; an example of a join is sequencing two commands together, where all possible combinations of a Each unix command is further broken down into the stub (the first word on the command line, typically a program like `ps` or `gcc`), and the individual arguments. As with commands, the number of arguments is lazily determined when the first example of this procedure step is observed. Each argument may either be a constant, a filename, a Unix user identifier, or some function of the previous command’s output. A filename could either be a literal, constant string, or a wildcard whose value changes each iteration through a loop (for example, when touching all the files in a director in sequence).

The CommandOutput version space contains hypotheses that describe different ways to extract information from the output of the previous command. The information to be extracted is defined as an extent, or a range of characters spanning two locations in the output. The Index version space, for example, contains hypotheses that describe locations in the character string based on offset from the beginning of the string. The RowCol version space describes locations based on row and column position in the previous command’s output, and can be used (for example) to define a substring that begins at the start of the third row of output. The version space diagram in Figure 4 describes the structure of the version space. We next describe how to update the

structure in response to training examples. With each new example, the version space is pruned to only contain hypotheses that are consistent with the observed example. Examples are decomposed into sub-components that are used to train the corresponding version spaces. For example, the trace of Unix commands is separated into a number of individual Unix commands, and each command is used to update the corresponding version space. Each unix command is parsed into stub and arguments, and the values used to update the corresponding version spaces. For example, the stub `ls` must match the stub observed in previous examples of this step, otherwise the Stub version space collapses (contains no constant-string hypotheses that are consistent with all training examples). Version space collapse is one indicator that the target procedure does not lie within the bias expressed in this version space; in future work, we plan to dynamically extend the bias to consider additional hypotheses in this event.

Discussion

Our experience with the SMARTshell system has led us to formulate several desiderata for machine learning algorithms that incorporate human input and enable the user to take control at specific points during the execution of the learned process. One of the largest barriers to the adoption of automated systems for system administrators is trust that the system will do the right thing in the right situation. Our desiderata thus reflect the need to establish trust between the human operator and the adaptive system. Specifically, we believe that:

- The system must be able to learn incrementally as users provide examples;
 - Learning must happen in real-time, so that the results are immediately accessible to the user;
 - The system must be able to explain its inferences, such as a user-understandable representation of the proposed hypothesis; and
 - The user should be able to understand the the learning algorithm at a high level.
 - The user and the system ought to be able to take turns performing steps in the procedure.
- These desiderata will certainly inform our future work in the area of automating tasks through programming by demonstration, and may apply to other areas where supervisory control of machine learning systems is required.

Conclusions

In summary, we have described an approach to learning shell scripts by demonstration. We have cast the problem as a machine learning problem and formalized it using the version space algebra framework. Our SMARTshell system provides an interface for users to interact with the system to demonstrate new examples, refine the current hypothesis, and execute the procedure either automatically or under the user's control.

References

Kandogan, E., and Maglio, P. P. 2003. Why don't you trust me anymore? Or the role of trust in troubleshooting activities of system administrators. In *CHI 2003 Workshop: System Administrators are Users, Too*.

Lau, T.; Wolfman, S. A.; Domingos, P.; and Weld, D. S. 2003. Programming by demonstration using version space algebra. *Machine Learning* 53(1-2):111–156.

Lau, T.; Domingos, P.; and Weld, D. S. 2000. Version space algebra and its application to programming by demonstration. In *Proceedings of the Seventeenth International Conference on Machine Learning*, 527–534.

Mitchell, T. 1982. Generalization as search. *Artificial Intelligence* 18:203–226