

This article was downloaded by: [Virginia Polytechnic Institute and State University],
[Clifford A. Shaffer]

On: 19 April 2012, At: 06:26

Publisher: Routledge

Informa Ltd Registered in England and Wales Registered Number: 1072954 Registered
office: Mortimer House, 37-41 Mortimer Street, London W1T 3JH, UK



Computers in the Schools

Publication details, including instructions for authors and
subscription information:

<http://www.tandfonline.com/loi/wcis20>

The Role of Visualization in Computer Science Education

Eric Fouh^a, Monika Akbar^a & Clifford A. Shaffer^a

^a Virginia Tech, Blacksburg, Virginia, USA

Available online: 18 Apr 2012

To cite this article: Eric Fouh, Monika Akbar & Clifford A. Shaffer (2012): The Role of Visualization in
Computer Science Education, *Computers in the Schools*, 29:1-2, 95-117

To link to this article: <http://dx.doi.org/10.1080/07380569.2012.651422>

PLEASE SCROLL DOWN FOR ARTICLE

Full terms and conditions of use: <http://www.tandfonline.com/page/terms-and-conditions>

This article may be used for research, teaching, and private study purposes. Any
substantial or systematic reproduction, redistribution, reselling, loan, sub-licensing,
systematic supply, or distribution in any form to anyone is expressly forbidden.

The publisher does not give any warranty express or implied or make any representation
that the contents will be complete or accurate or up to date. The accuracy of any
instructions, formulae, and drug doses should be independently verified with primary
sources. The publisher shall not be liable for any loss, actions, claims, proceedings,
demand, or costs or damages whatsoever or howsoever caused arising directly or
indirectly in connection with or arising out of the use of this material.

The Role of Visualization in Computer Science Education

ERIC FOUH, MONIKA AKBAR, and CLIFFORD A. SHAFFER
Virginia Tech, Blacksburg, Virginia, USA

Computer science core instruction attempts to provide a detailed understanding of dynamic processes such as the working of an algorithm or the flow of information between computing entities. Such dynamic processes are not well explained by static media such as text and images, and are difficult to convey in lecture. The authors survey the history of visualization in computer science education, focusing on artifacts that have a documented positive educational assessment. Changes in how computing technology has affected the development and uptake of such visualization artifacts in computer science education, and how recent technology changes are leading to progress in developing online hypertextbooks are then discussed.

KEYWORDS *algorithm visualization, data structure visualization, program visualization, eTextbooks, hypertextbooks*

Computer science instruction has a dual nature. On one hand, computer science (CS) students constantly build artifacts of the practice. Creating software at small and medium scale is done from early in the first course. This is so ingrained in the culture that students often express that something is wrong with any early CS course that is not programming-intensive. Students from most other disciplines would be fortunate to get a fraction of the hands-on experience that is normal for CS instruction. However, this is only half of the story. CS students must also learn a body of knowledge. Most of this core knowledge relates to abstractions. One cannot see or touch an algorithm or data structure, let alone the higher order concepts that relate these algorithms and data structures to the underpinning theories of the discipline. To make the situation worse, much of the content involves the detailed workings of dynamic processes. Particularly at the intermediate

Address correspondence to Clifford A. Shaffer, Department of Computer Science, Virginia Tech, Blacksburg, VA 24061. E-mail: shaffer@cs.vt.edu

stage (between when students have properly learned the basic concept of “algorithm” or “recipe” for computation in the first year, and when they begin courses on specific subareas of CS that come in the junior and senior years), students spend significant time learning about dynamic processes. Sorting algorithms, search algorithms, the data structures that support them, and the analysis techniques used to assess their quality are the bread and butter of sophomore-level computer science. Even the more concrete aspects of CS instruction, such as programming languages and computer architecture, require an understanding of the dynamic flows of information that take place within the computer.

Dynamic process is extremely difficult to convey using static presentation media such as text and images in a textbook. During lectures instructors typically draw on the board, trying to illustrate dynamic process through words and constant changes to the diagrams. Most students have a hard time understanding such explanations at a detailed level. For this reason, there has long been great interest among CS instructors in using animation and visualization to convey dynamic concepts. With the ready availability and potential for educational use of our own fundamental artifact—the computer—it is natural to expect that CS educators would turn to using computers in hopes of finding new ways to convey our own discipline’s core knowledge. In this article the authors discuss the use of algorithm visualization (AV) and program visualization (PV) for CS education.

EARLY USE OF AV

AV has a long history in CS education, dating from the 1981 video “Sorting out Sorting” (Baecker & Sherman, 1981) and the BALSAs system (Brown & Sedgewick, 1985). Since then, hundreds of AVs have been implemented and provided free to educators, and hundreds of papers have been written about them (AlgoViz.org Bibliography, 2011). Good AVs bring algorithms to life by graphically representing their various states and animating the transitions between those states. They illustrate data structures in natural, abstract ways instead of focusing on memory addresses and function calls.

Early systems, such as BALSAs (Brown & Sedgewick, 1985), Tango (Stasko, 1990), XTango (Stasko, 1992), Samba (Stasko, 1997), and Polka (Stasko, 2001) suffered from the fact that their implementation technology was not easy to disseminate to a broad audience. This has often been a problem with earlier educational software. Massachusetts Institute of Technology’s Project Athena (Champine, 1991) sought not only to provide interactive courseware, but also to solve issues of dissemination through the X Windows display system (Gettys, Newman, & Scheifler, 1989). Many of the AV systems of the late 1980s and early 1990s were built on X Windows for this reason. However, the fact that many institutions did not have access

to sufficiently powerful workstations to run X Windows in an educational laboratory setting made access to this courseware difficult.

The first significant attempt to assess the pedagogical effects of AV was conducted by Hundhausen, Douglas, and Stasko (2002). They performed a meta-study of 24 prior small-scale experiments related to AV. They found that 11 of those experiments showed a “statistically significant difference between the performance of a group of students using AV technology and another group of students using an alternative AV technology or no AV technology at all” (p. 265). Ten of those experiments showed a nonsignificant result, two showed a significant result but the AV’s contribution was uncertain, and one showed a significant result where the AV appeared to reduce learning. Measuring the educational gain was found to be difficult due to lack of agreement on the type of knowledge being measured. Some experiments involving AV technology measure *conceptual* or *declarative knowledge* (comprehension of the concepts of an algorithm) as well as *procedural knowledge* (internalization of the step-by-step behavior and data manipulation of an algorithm). AVs in these studies were found to give the most significant results when measuring only the procedural knowledge.

AV systems reviewed in Hundhausen et al. (2002) had mixed results in regard to their pedagogical effectiveness. Use of XTango (Stasko, Badre, & Lewis, 1993) did not result in a statistical difference in posttest accuracy between students taught only by the textbook and those using the textbook and the visualization tool. Another experiment involving XTango and Polka (Byrne, Catrambone, & Stasko, 1999) as lecture aid resulted in better ability to predict the next step of the algorithm for students who viewed the animation in comparison to ones who only used the textbook. On programming accuracy, students who constructed and viewed algorithm animations using Balsa II (Brown, 1988) and Polka (Kann, Lindeman, & Heller, 1997) outperformed learners who did not use them. The authors note that the AV systems under study were all of the pre-Java variety, and the studies were conducted at a time when student access to computers (and especially to the AVs being used) would come only in a computer lab setting.

Perhaps the most significant findings of the Hundhausen et al. (2002) meta-study were that (a) how students use AV, rather than what they see, has the greatest impact on educational effectiveness; and (b) AV technology is effective when the technology is used to actively engage students in the process. This view that engagement is key has had enormous impact on subsequent AV development.

AV IN THE INTERNET ERA

Much changed in the mid 1990s with the advent of the World Wide Web and the Internet as we know it. In large part due to the widespread availability

of Java, and to a lesser extent of JavaScript and Flash, AVs soon became available to many more instructors and students. At about the same time, more and more students gained access to their own computers on campus. During the late 1990s and early 2000s, a whole new crop of AV development systems based on Java came into being, and the older X Windows-based AV systems lost what little traction they had with the educational community.

Examples of the newer systems written in Java include JSamba (Stasko, 1998); JAWAA (Pierson & Rodger, 1998); JHAVÉ (Naps, Eagan, & Norton, 2000); ANIMAL (Rossling, Schuler, & Freisleben, 2000); and TRAKLA2 (Korhonen et al., 2003). One unifying feature of these systems is that they are intended to be complete toolkits for building AVs. This followed in the tradition of the pre-Java systems, which offered a programming framework for would-be AV developers. While the pre-Java systems were not easy to develop in, they at least made the development process feasible for those willing to invest considerable effort. In contrast, some of the Java-based development systems (such as JAWAA) focused particularly on making it easy for instructors to develop visualizations, albeit without much support for interactivity.

A key change that coincided with the use of Java was the advent of collections of AVs independent of any proposed AV development system. Java granted an AV developer independence from any given operating system or similar constraints, and the Internet allowed the software to be easily distributed. This changed how developers viewed their role. While JHAVE and ANIMAL follow the tradition of providing a system, their AV content is also significant. Many other projects have developed AVs as a stand-alone package with no pretensions of providing an AV system to be used by other developers. Examples of significant pre-packaged collections include Data Structure Navigator (DSN) (Dittrich, van den Bercken, Schafer, & Klein 2001), Interactive Data Structure Visualization (IDSV) (Jarc, 1999), Algorithms in Action (Stern, 2001), Data Structure Visualization (Galles, 2006), and TRAKLA2 (Korhonen et al., 2003).

Another effect of the Java revolution was the advent of one-off AV development by students and/or instructors. Building an AV is a great way for students to learn Java, something that they found attractive in the late 1990s. Unfortunately, most of these student-built AVs were poor in terms of pedagogical value (Shaffer et al., 2010). Fortunately, a number of particularly well-regarded AVs were also developed, usually as part of a sustained effort over several years to create individual Java applets to present specific lessons. Examples include Binary Treesome (Gustafson, Kjensli, & Vold, 2011), JFLAP (Rodger, 2008), University of Maryland's Spatial Index Demos (Brabec & Samet, 2003), and Virginia Tech Algorithm Visualization Research Group (2011).

THE IMPACT OF ENGAGEMENT

The early-to-mid 2000s saw a series of influential working group reports conducted as part of the annual ACM Innovation and Technology in Computer Science Education (ITiCSE) conference (Naps et al., 2003b, Naps et al., 2003a, Rossling et al., 2006). Naps et al. (2003b) extended the conclusions of Hundhausen et al. (2002) regarding the need for AVs to engage the learner in an active way. Their report defined a taxonomy that specifies the level and the type of engagement at which AVs can involve learners. They defined six engagement categories:

- *No viewing* indicates cases where no AV is used;
- *Viewing* involves both watching the AV execution and directing the pace of the AV;
- *Responding* involves answering questions about the visualization being viewed;
- *Changing* asks the learner to provide input data to the AV to direct its actions;
- When *Constructing*, students build their own visualization of the algorithm; and
- *Presenting* asks the learners to tell about the AV and gather feedback.

This taxonomy imposes a hierarchy in the learner engagement that is closely related to the Bloom hierarchy (Bloom & Krathwohl, 1956).

Urquiza–Fuentes and Velázquez–Iturbide (2009) performed a new meta-study to evaluate AVs in regard to their pedagogical significance. Their evaluation focused on areas where some educational benefit has been perceived. The educational benefit was measured against the taxonomy defined by Naps et al. (2003b). They found that *viewing* does not increase the learning benefit compared to *no viewing*. This is consistent with the prior findings of Hundhausen, et al. (2002). Unfortunately, Shaffer et al. (2010) reported that most AVs are still at the *viewing* level.

Urquiza–Fuentes and Velázquez–Iturbide found that students interacting at the *responding* level with an AV have significantly better knowledge acquisition than those using AVs with only *viewing* features. The *changing* level has better acquisition results than the *responding* level. *Constructing* an AV improves the learning outcomes more than *changing* (though at a significantly greater cost in time and effort). The *presenting* level has better acquisition results than *constructing*.

Among the successful AVs, Urquiza–Fuentes and Velázquez–Iturbide cited specific examples from various larger AV systems. JHAVÉ (Naps et al., 2000) AVs consist of a graphical interface plus textual information about the algorithm and some questions for the learners. JSamba (Stasko, 1998)

lets learners and instructors use a scripting language to customize the representation of data structures. TRAKLA2 (Korhonen et al., 2003) includes exercises where the learner demonstrates proficiency in a step-by-step execution of the algorithm, assessing the knowledge of the learner by asking him to predict the next step of the algorithm. Alice (Cooper, Dann, & Pausch, 2000) allows the learner to visualize objects in a 3D virtual world. jGRASP (Hendrix, Cross, & Barowski, 2004) lets instructors develop specific visualizations to satisfy an objective (how they want the learners to see the objects).

Urquiza–Fuentes and Velázquez–Iturbide noted that an element typically found in successful AVs is narrative content and textual explanations. Integrating an explanation with the steps of the visualization improves learning at the *viewing* engagement level. Feedback on student action is a feature at the *responding* engagement level, where the student is actually engaged by demonstrating his/her proficiency with the knowledge. Extra time using the AV can be achieved when the students are asked to construct or to customize an AV (the *constructing* engagement level). Good AVs also provide advanced features like a more versatile interface to manipulate and direct aspects of the visualization (often required to support the *responding* level and above).

Myller, Bednarick, Sutina, and Ben–Ari (2009) investigated the impact of engagement on collaborative learning. They extended the engagement taxonomy of Naps et al. (2003b) to better capture the differences between student behaviors. They added four engagement levels, including *controlled viewing* and *entering input*, falling between the original levels of *viewing* and *responding*. They claim that their extended engagement taxonomy (EET) can also be used to guide successful collaboration among students when using visualization tools. To probe their claim, Myller et al. (2009) conducted an experiment in which students were asked to work in pairs during lab session. The students interacted with the program visualization tools BlueJ (Sanders, Heeler, & Spradling, 2001) and Jeliot (Levy, Ben–Ari, & Uronen, 2003) at different EET levels. They observed and recorded all students' communication during the experiments. They found that EET levels are positively correlated with the amount of interaction. Interaction is one of the essential ingredients for a successful computer-supported collaboration along with coordination and motivation (Meier, Spada, & Rummel, 2007).

SOME SELECTED VISUALIZATION ARTIFACTS

In this section the authors describe in some detail several visualizations or systems that have supported documented pedagogical improvements. Those that showed a statistically significant difference in students' performance before and after using the visualization systems under a controlled experiment were selected, typically because they engaged learners at levels of the taxonomy above *viewing*.

TRAKLA2

TRAKLA2 is one of the most widely used AV collections, routinely used throughout Finnish universities. It provides an outstanding example of operating at the *responding* engagement level. TRAKLA2 allows the learner to control the visual representation of the data structures manipulated by the algorithm. Learners can “build” a data structure by dragging and dropping GUI elements. TRAKLA2 exercises ask learners to determine a series of operations that will change the state of the data structure to achieve some outcome. For example, students might build a tree data structure by repeatedly dragging new values to the correct locations in the tree. Alternatively, the learner can practice to gain understanding by examining a step-by-step execution of the algorithm (called the model solution). Many TRAKLA2 exercises include some tutorial text along with pseudocode to explain the algorithm, but their main purpose is to provide an interactive proficiency exercise.

Laakso et al. (2005) reported on the use of TRAKLA2 exercises in data structures and algorithms courses at two universities in Finland. TRAKLA2 exercises were incorporated as classroom (closed lab) exercises, and supplemented lecture hours and classroom sessions. TRAKLA2 exercises were also incorporated into the final exam (one out of five questions) and midterm (they replaced half of a question with a TRAKLA2 exercise in a total of five questions). A year after the introduction of TRAKLA2 in the curriculum, the activity of students increased in all aspects of the course, not only the part involving TRAKLA2 exercises. The average performance in classroom exercises rose from 54.5% to 60.3% (number of exercises completed). Student opinions of TRAKLA2 were collected through an online survey. Eighty percent of the students reported having a good opinion of TRAKLA2. After a year of using TRAKLA2, the opinion of students on its suitability for instruction rose significantly. Ninety-four percent of students agreed that TRAKLA2 exercises helped in their learning process. Their opinion did not change regarding their preference on the medium used to access and perform the exercises: They preferred a mixed (online and traditional classroom) experience. The later result confirms recommendations on integrating AVs with existing class materials from Levy and Ben-Ari (2007), discussed later.

Usability plays a big role in the success of an AV in regard to the learning objectives. Usually learners spend more time-on-task when using an AV system. But it is not clear if this is because the learners have to get used to the system due to its lack of usability (Pane, Corbett, & John, 1996). The usability of TRAKLA2 has been studied through a series of questionnaires and observations in a usability lab (Laakso et al., 2005). The observations did not find any critical usability issues, and revealed that 80% of the time was used to solve exercises while 14% of the time was used to get acquainted with the system interface. The students had a total of 15 minutes to complete all the exercises. This is an important finding since it is typical that any given AV is intended to be used from only a few minutes up to a couple of hours.

TRAKLA2 has been used to assess the impact of the EET on student performance. Laasko, Myller, and Korhonen (2009) studied whether students using a visualization tool in pairs at EET-level *changing* outperformed pairs using it on EET-level *controlled viewing*. First-year students were divided into two groups, both groups using the same textual materials. The treatment group (EET-level *changing*) used TRAKLA2 exercises while students in the control group (EET-level *controlled viewing*) viewed AVs containing equivalent information. All the students took an individual pretest and were asked to freely form pairs. They then had 45 minutes to study the learning materials and solve exercises together (using paper and pencil). All the students were also asked to take an individual posttest. The students were videotaped during the experiment. The first analysis of pretest and posttest scores did not show a statistically significant difference between the treatment and the control group; but a second analysis of the video revealed that some students in the treatment group were not using TRAKLA2 as expected: They did not solve TRAKLA2 exercises but instead just watched the model solution. Thus they were interacting at the EET-level *controlled viewing*. After regrouping the students by creating a third group for the students in the treatment group interacting on the *controlled viewing* level, the analysis of the scores showed that the students in the treatment group interacting on a *changing* level outperformed the *controlled viewing* groups.

JHAVÉ

JHAVÉ (Naps et al., 2000) is an AV development system aimed at making it relatively easy for AV developers to create animated slideshows with built-in assessment through pop-up questions to users. JHAVÉ's interface includes a pane containing the visualization, a pane displaying the pseudocode, and often a brief text giving a tutorial about the algorithm. The information page allows importing of images and has been used for example to represent the algorithm's flow chart. JHAVÉ now includes a large collection of AVs, and has seen wide use.

Lahtinen and Ahoniemi (2009) developed "Kick-Start Activation," a method for teaching CS1 courses to learners with no programming experience using visualization. Their method relies on three requirements:

1. A real programming problem as the starting point of the lesson. Data structures, pseudocode, flow charts, and programming concepts are explained based on the solution of the problem. Visualization tools are used as class materials.
2. Use of pseudocode and flow charts to represent the algorithmic solution of the problem. The pseudocode looks similar to a description in English.
3. The student should interact with the algorithm (testing, debugging, etc).

Lahtinen and Ahoniemi (2009) argued that most AV systems presume that learners are already familiar with the programming language used by the system. Since they were targeting novice learners with no or little programming experience, they needed an AV system that was syntax free and flexible enough to represent the flow chart of the algorithm within the system. JHAVÉ was chosen as the visualization tool to supplement their method. In Kick-Start Activation, students were asked to design an algorithm at an early stage of the class, called “the premature algorithm.” They then tested and visualized their solution and eventually amended it. The result was “the mature algorithm,” and it is visualized in JHAVÉ. When using this method, learners are engaged with the AV system at the *viewing*, *responding*, *changing*, *constructing*, and *presenting* levels. The evaluation of Kick-Start Activation/JHAVÉ revealed that 86% of the students found the visualization useful for learning. Students with no programming experience found it more useful than ones with previous experience. Interesting aspects of this experiment included integration of the AV into a pedagogical method, and an innovative way of using the AV tool: pseudocode close to English and construction of the algorithm’s flow chart using images, all packaged with the Java applet for the algorithm. This was possible because of the flexibility of JHAVÉ. Students used JHAVÉ throughout the entire class.

JHAVÉ-POP (Furcy, 2007) is a JHAVÉ extension for practicing elementary pointer and linked-list operations. The user can type code snippets in C++ or Java, and JHAVÉ-POP automatically generates a step-by-step visualization of the contents of memory as each program statement is executed. Student feedback on JHAVÉ-POP collected through a questionnaire (Furcy, 2009) indicated that “JHAVÉ-POP deepens student understanding of pointers by making them concrete” and it “helps students not only understand but also debug their programs” (p. 39).

ALVIS

ALVIS (Hundhausen & Brown, 2005) is a program development environment where programs are written using a pseudocode-like language called SALSA. It includes features to support storyboarding. Hundhausen and Brown (2008) developed pedagogy with the goal of involving their students at all five engagement levels within a first semester programming class studio experience. In this experiment, students worked in pairs to create their own visualization of algorithms studied in class. They presented and discussed their solutions with classmates and instructors. The students used ALVIS to develop their solution. To evaluate their system, Hundhausen and Brown divided the learners into two groups, one using ALVIS and the other using art supplies (pen, paper, tape, etc.) to build their visualizations. All students developed their solutions in SALSA, using either a text editor (art supplies group) or ALVIS.

Students were videotaped during studio sessions. All artifacts used during the sessions were collected afterward. All students answered a series of questionnaires to describe their experience, and the students using art supplies were interviewed at the end of their sessions. For both session groups, students spent the majority of their time coding their solution. Students using art supplies discussed more with the teaching assistant in comparison to students using ALVIS. Codes developed by the ALVIS students had approximately half as many errors per solution/algorithm as the ones in the art supplies group. Hundhausen and Brown contended that developing an algorithm using either approach prepared the student for discussion and critical thinking. At the end of the experiment, ALVIS was found to help students code faster, with less semantic errors, and to foster audience participation during presentation, yielding more technical discussion about the algorithm.

Virginia Tech Hashing Tutorial

The Virginia Tech Hashing Tutorial (Virginia Tech Algorithm Visualization Research Group Web site, 2011) goes beyond a simple AV, providing a complete textbook-quality treatment of an important topic in CS, the concept of search by hashing. A series of AVs (in the form of Java applets) are embedded in text presented in HTML pages. In addition to showing the basics of the hashing process, additional applets allow students to explore the relative performance of various approaches discussed in the tutorial.

In 2008 and repeated again in 2009, students in separate sections of a sophomore-level course on data structures and algorithms were taught about hashing using two separate treatments. One section was given standard lecture and textbook for one week, similar to what had been done previously in the class. The other section spent one week of class time working through the online tutorial with AVs to present the same material. Since the tutorial used text from the course textbook, it was an exact match to the material being presented in the control section. However, the online tutorial heavily supplemented this text with AVs.

In each of the trials, the two sections were given a quiz on hashing at the conclusion of the week. The results were positive: Significant differences in performance were obtained in favor of those who used the online tutorial versus standard lecture and textbook. This means that not only can an online tutorial be as effective as a lecture (with important implications for distance learning), but that providing proper interactivity allows computerized instruction to be better than lecture-based (passive) instruction. However, at least one major consideration might influence the results of this study: How much impact did the controlling structure of coming to class and doing the tutorial in a lab setting have on the results? The outcome could be quite different for a student just reading the material and working through the AVs

on his/her own, where self-discipline might not be sufficient to provide the necessary amount of time and attention. Likewise, the controlled environment of attending lecture before reading the textbook on one's own is also likely to have a major difference compared to simply reading the book on one's own.

ALViE

Like JHAVÉ, ALViE (ALViE, 2011) is a post-mortem algorithm visualization tool since it allows the learner to see the resulting visualization after an algorithm execution. This means that the visualization is driven by a script that could be generated in a number of ways, including from an instrumented program executing a complex algorithm or simulation. It is written in Java and uses XML to describe the data structures and the relevant events of the algorithm execution.

Crescenzi and Nocentini (2007) described the results of two years of teaching data structures and algorithms using ALViE intensively. They designed the courses so that their students could go through the maximum number of engagement levels. The first year of the experiment involved the engagement levels *no viewing*, *viewing*, *constructing*, and *presenting*. ALViE was systematically used during class sessions to teach students about algorithms and data structures. The homework consisted of implementing an algorithm without the help of the visualization tool. Each student was assigned a final project to build an animation for a specific algorithm using ALViE and then present it to his/her peers and the instructor. The students were asked through a survey to evaluate their experience with ALViE. Seventy percent of the students reported *positive* and 30% *very positive* opinions on the use of ALViE for the final exam. Ninety percent of the students agreed that viewing the animation improved their understanding of the algorithm's behavior. All students agreed that the use of visualization was educationally valuable.

During the second year, the authors made some changes: They published a (printed) textbook (Crescenzi, Gambosi, & Grossi, 2006) and they added the *changing* engagement level. The interesting point of the textbook is that it was designed to be an extension of ALViE, with the textbook describing the algorithm and the data structures. All illustrations in the textbook came from the ALViE GUI. Readers are invited to visualize the execution of all algorithms presented in the book. The textbook and ALViE have been adopted by several universities in Italy.

Alice

Alice (Alice, 2011) is a 3-D interactive programming environment. It allows users to create virtual worlds by dragging and dropping objects in the main

windows. Users can write scripts to control the objects' behavior. Alice's goal is to expose novice learners to object-oriented programming using a simple, graphical interface. Its ease of use and support for storytelling has led to wide adoption within the CS community. It has been used to introduce students at a wide variety of levels to programming, including CS majors, nonmajors, high school students, and middle school students.

Moskal, Luri, and Cooper (2004) conducted a study to find if exposure to Alice improves CS1 students' performance. Specific questions follow: Is Alice effective in improving the retention rate of students in the CS major? Does exposure to Alice improve students' confidence in their ability to succeed in CS? All students involved in the experiments were CS majors. They were divided into three groups: (a) a treatment group consisting of students with no or little programming experience and who enrolled in the Alice-based course; (b) control group 1, consisting of students with no or little programming experience and who did not enroll in the Alice-based course; and (c) control group 2, consisting of students with previous programming experience and who did not enroll in the Alice-based course. Students in the treatment group earned an overall GPA of 3.0 ± 0.8 in the class. Students in control group 1 earned an overall 1.9 ± 1.3 GPA. Control group 2 earned an overall GPA of 3.0 ± 1.2 . Historically, students who reach college with little or no prior programming experience do far worse in introductory CS courses and have much lower retention rates than students who come to college with programming experience, so this is potentially an important result. The retention rate over the two years of the experiment was 88% for the treatment group, 47% for control group 1, and 75% for control group 2. Evaluation of student attitudes revealed that students with no prior programming experience who did not take the Alice-based course had "more negative attitudes with respect to creativity in computer science after CS1" (p. 78). The authors concluded that Alice was effective in improving student performance, that it can help increase the retention rate in CS programs, and that it can better the attitude of students toward CS.

Jeliot

Jeliot (Levy et al., 2003) is designed to help teach Java programming to high school students. In Jeliot the learner can visualize program source code along with an automatically generated animation showing the behavior of the program. Thus, Jeliot is an example of a PV tool, since it visualizes the behavior of an actual program rather than a hand-crafted presentation of an abstract algorithm. PVs are able to generate visualizations that give insight into the behavior of actual programs with little additional effort on the part of the programmer beyond writing the program. They have the disadvantage that they cannot so easily abstract away irrelevant details and

focus on higher order or key selected aspects of the underlying algorithm. The interface allows for step-by-step execution of the program.

Jeliot visualizations were used to present conflicting animations by Moreno, Sutinen, Bednarik, & Myller, (2007). They argue that conflicting animations can be used to expose learners' misconceptions about programming concepts. Conflicting animations extend the *responding* engagement level by asking the student to detect errors in the animation that is being executed. The *changing* level requires students to fix the erroneous part of the animation. The student might construct a conflicting animation at the *constructing* level and present it to their classmates, thus extending to the *presenting* engagement level. Moreno, Joy, Myller, and Sutinen (2010) developed Jeliot ConAn, a Jeliot-based AV system to generate animations that contain known bugs.

To enhance learners' mental models, Ma, Ferguson, Roper, Ross, and Wood (2009) proposed a four-stage model to detect and fix bad mental models regarding programming concepts. A model's stages are as follows:

- *Preliminary*: Find the erroneous pre-existing mental mode;
- *Cognitive conflict*: Dispute the students' pre-existing mental model and then induce adaptation/change in the model;
- *Model construction*: Use visualization to aid students in the process of constructing a correct mental model; and
- *Application*: Test the newly constructed mental model (typically by solving programming problems).

To assess mental models, the student logs into a system displaying a roadmap of all the programming concepts that must be understood. Each concept is attached to a series of exercises. Each exercise is linked to a cognitive question, its associated Jeliot visualization materials, and a final question to test understanding of the concept. If a student's answer to the cognitive question is incorrect, then he/she is required to run the visualization attached to the question to identify the discrepancy between his/her model and the program's execution. During this phase the student can be helped by the instructor to improve his/her understanding of the algorithm. The student is tested again using a similar program implementing the same concept with different data values. The authors evaluated their models on three content areas: Conditionals and loops, scope and parameter passing, and object reference assignment. Out of 44 students who participated in the experiment on conditionals and loops, 23% had a correct understanding of the concept at the beginning of the experiment, and 73% had a correct understanding at the end of the experiment. Only 19% of the learners had a good comprehension of scope and parameter passing prior to the experiment, and 85% after. Ma et al. (2009) found that their model works well on simple concepts, but

its benefits are hardly perceptible on the more complex ones (e.g., object reference assignments).

ViLLE

ViLLE (Rajala, Laasko, Kaila, & Salakowski, 2007) is a program visualization tool developed at the University of Turku in Finland. It is designed to be used as a lecture aid or for independent learning to visualize the execution of example programs created by instructors or students. The user can follow changes in program states and data structures. The key features of ViLLE are its support for multiple programming languages (including Java, C++ and an extensible form of pseudocode) and its built-in editor for interactive quizzes and tests displayed as pop-up windows.

The effectiveness of ViLLE was assessed through an experiment on students in their first programming class (Rajala, Laasko, Kaila, & Salakoski, 2008). The students were divided into two groups, one using ViLLE and the other using only the textbook. The experiment took place during a two-hour computer lab session. Students in both groups took a pretest where they were asked to predict the output or the state of a program. They then read a programming tutorial with ViLLE as lecture aid for one group and only the tutorial for the other group before taking the posttest. The posttest had the same question as the pretest plus two extra questions about completing a piece of program with given statements. The experiment did not reveal a statistically significant difference between the ViLLE group and the control group as far as learning acquisition. However, within the ViLLE group, there was a significant difference in knowledge improvement for students with no previous programming experience over those with prior experience, thus reducing the gap between novice and experienced programmers. In contrast, a significant difference in the total points earned in the posttest remained between students with prior programming experience and those with no prior experience in the control group. Another experiment was conducted using high school students (Kaila, Rajala, Laakso, & Salakoski, 2009). All class materials were accessible via the Moodle course management system with links to ViLLE for the treatment group. At the final exam, students using ViLLE performed better than the control group (statistically significant). ViLLE was more effective for program execution tracing and learning program writing skills.

jGRASP

jGRASP (Hendrix et al., 2004) provides a complete program development environment with synchronized “object viewers” that visualize objects and data structure states. It has been used in lab sessions when teaching CS2 courses using Java. A series of experiments were performed by Jain, Cross,

Hendrix, and Barowski (2006) at Auburn University to evaluate the pedagogical benefits of jGRASP. The experiments were intended to find if using jGRASP data structure viewers results in students being able to produce code with fewer bugs, and whether students are able to detect and fix non-syntactical bugs more accurately. Two experiments were conducted to answer those questions. In each experiment, students were divided into two groups. One used only the JGRASP debugger while the other used both the debugger and the object viewers. The students were observed during in-lab activities. They were required to implement four basic operations with singly linked lists. This experiment showed a statistically significant difference between the two groups in regard to accuracy. The mean accuracy was 6.34 points for the treatment group and 4.48 points for the control group. The authors concluded that in most cases (95%), jGRASP can be used to “help increase the accuracy and reduce time taken for programs implementing data structures” (pp. 35). During the second experiment, students were asked to detect and fix non-syntactical errors in a Java program (containing nine errors) implementing a singly linked list with four functions. This experiment showed that the students using object viewers detected on average 6.8 bugs versus 5.0 for the control group. Students using the visualizers were able to correct an average of 5.6 bugs and introduced only 0.65 new errors in the program while the control group fixed 4.2 bugs and introduced 1.3 new errors on average. However, students using the object viewers spend slightly more time to complete the task: 88.23 minutes on average versus 87.6 minutes for the other students.

JFLAP

JFLAP (Java Formal Languages and Automata Package) (Rodger, 2008) is designed to help teach formal languages and automata theory. This is traditionally senior-level material, though JFLAP has been used to teach aspects of formal languages to first-semester students. Users can build a variety of automata in JFLAP, then simulate and visualize their behavior. JFLAP can also be used to create and parse strings in grammars and convert a nondeterministic finite automaton to a deterministic finite automaton. JFLAP effectiveness was assessed through a two-year study (Rodger et al., 2009) involving 12 universities the first year and 14 the second year. The goal of the experiment was to answer these questions: How effective is JFLAP in enhancing the learning process? and What additional value might JFLAP add to the Formal Languages and Automata course? During the first year, most of the instructors used JFLAP for homework and for demonstrations in class; no instructor used it for examinations. The results after the first year of the experiment showed that 55% of the students used JFLAP to study for exams. Ninety-four percent said that they had enough time to learn how to use JFLAP. Sixty-four percent said that using JFLAP helped them to earn better grades. Eighty-three percent

said that it was easier to use JFLAP than paper and pencil. However, 50% thought that they would have performed as well in the class without using JFLAP. During the second year of the experiment, the pretest and posttest were shortened. The survey revealed that when preparing for exams, 32% of the students used JFLAP during more than 20% of their study time, and 29% of the students used it to study additional problems. JFLAP made the course more enjoyable for 63% of the students, and 72% said that using JFLAP made them feel more engaged in the course.

HOW OFTEN ARE AVS USED IN CS COURSES?

Periodic surveys of instructor interest in AVs and their level of use have been conducted at CS education conferences and on CS education listservs for more than a decade. Naps et al. (2003b) reported on three surveys conducted in 2000. Collectively, they indicate a strong positive view in favor of AV use (over 90%). However, only about 10% of respondents at that time indicated frequent use of AVs in data structures courses at their institutions, while half to three quarters indicated occasional use of AVs in such classes.

At SIGCSE'10 (a major CS education conference, held in Milwaukee during March 2010), the authors conducted a new survey to determine instructor attitudes toward AVs and their use in the classroom. For details, see Shaffer et al. (2011). The survey was designed in the spirit of the 2000 surveys reported in Naps et al. (2003b) and the findings are consistent with those earlier results. In 2010, 41 of 43 respondents agreed or strongly agreed that AVs can help learners learn computing science, with two neutral. However, just over half had used AV in a class within the past two years. It is not clear that all of the “yes” answers refer to what we might consider to be AVs. However, they probably all refer to some sort of dynamic software visualization run on a computer, as was the case in the 2000 survey. The authors did not ask about frequency of use of AVs in class.

The third survey question asked what respondents see as the greatest impediments to using AVs in courses. Roughly half the responses relate to finding suitable AVs to use. The other half of the responses relate to issues in integrating AVs into the courses. These results are roughly the same as reported in Naps et al., 2003a, Naps et al., 2003b, and Rössling et al., 2006. Such issues are much harder to deal with, and are representative of well-known problems for adoption of educational technology in general (Hew & Brush, 2007). While it is easy to give students pointers to AVs as supplemental material, it is much harder to integrate AVs into lectures, labs, and assignments. This comes as a result of factors like lack of time on the instructor's part to make course changes, lack of time within the course to spend additional time on a given topic, and inconsistencies between the course textbook and the AV.

AV use in the classroom is highly dependent on the instructor's willingness to integrate AVs into his/her class material. Levy and Ben-Ari (2007) studied adoption of Jeliot by instructors. They grouped instructors into four categories. One group used Jeliot regularly during a school year, and two groups were asked to use Jeliot at least once in an exercise. All instructors were acquainted with Jeliot but had no prior experience with using an AV in their classroom. They classified instructors' experience using AVs into four categories, each category defined by a specific attitude toward the use of AVs in the classroom:

- *Internalization*: Instructors who found Jeliot useful and consistent with their teaching approach. They used the system frequently and integrated it into their class materials and activities.
- *By-the-book*: Instructors who found Jeliot useful but did not integrate it well with the class. These teachers relied on the textbook and used the AV only occasionally.
- *Rejection*: Jeliot was perceived by these instructors as having limited educational benefits, thus they did not use it in their class.
- *Dissonant*: Instructors who experienced Jeliot with mixed feelings. They professed enthusiasm about the system but were disinclined to use it in their classroom, and these instructors used it only when required to do so.

To better the experience of instructors using AVs, Levi and Ben-Ari advocated for a higher integration (by designers) of AVs into existing class materials like textbooks. This was also recommended by Naps et al. (2003b), who encouraged AV designers to create supporting Web sites along with the AV to ease the integration process.

An important resource for instructors seeking to use AVs is the AlgoViz Portal (AlgoViz.org, 2011). AlgoViz provides a gateway for users and developers to AV-related services, collections, and resources. It includes a catalog of over 500 AVs, and a large bibliography of related research literature. AlgoViz attempts to address issues raised in the instructor surveys regarding lack of information about AV availability and use. AlgoViz allows the community to add value beyond information embodied in a simple AV catalog. This value comes as a byproduct of direct communication between community members. While forums are one obvious method for this, there is a deeper communication involving community ratings and recommendations for content entries and sharing experiences on how to make the best use of content. A major concern for instructors according to the surveys is deciding whether a given educational resource is of good quality, and knowing how to use it. Therefore, user feedback on resources is as important as the catalog entries. AlgoViz also provides the *field report*, which gives a convenient way for instructors who have used an AV in a class to share their experience. Field reports supplement the evaluation data in the AV catalog, since catalog

AV ratings are not typically based on real-world classroom experiences. A field report can thus provide empirical evidence to strengthen or qualify a recommendation.

Important technological factors contribute to making AVs easier to use in the classroom than in prior years. More AVs are available than ever before (Shaffer et al., 2010), backed up by improved research studies and improved access both through general Internet search and at the AlgoViz portal. Increased access to the Internet by students and instructors, both in and out of the classroom, makes AV use more practical. For example, at Virginia Tech, while it has been possible to project Internet material from a computer in class for over a decade, it has only been in the past few years that such access has become both ubiquitous and reliable in all of our classrooms. This makes a huge difference in the confidence of mainstream instructors for using such technology, in contrast to early adopters (Hew & Brush, 2007). From discussions with instructors and from comments on the surveys, we know that even now not all schools have sufficient classroom support infrastructure to allow broadcasting of Internet presentations in typical classrooms. This is still an evolving technology.

Another potential factor in favor of increased use of AVs is ubiquitous availability of laptops and mobile devices. For example, all engineering majors at Virginia Tech are required to own a tablet PC, and most also have mobile devices including smartphones, eBook readers, or iPads. However, there is also a downside to the non-PC devices, in that they have various technology limits on how courseware content can be provided. So while there is ubiquitous access to the Internet among our target audience, there are still limits. For example, Java applets cannot be displayed on most such devices. However, recent widespread adoption of the HTML5 Internet standard makes it easier than before to develop interactive content. Content developed under the new standards will run on all major browsers without the need for specialized plugins, and will automatically be adapted to viewing on a variety of mobile devices. This means unprecedented ability for AV developers to reach a broader audience along with easier adoption of such innovative courseware by that audience.

THE FUTURE: ELECTRONIC TEXTBOOKS

Many in the CS community share the dream to create *hypertextbooks*, on-line textbooks that integrate AVs, assessment exercises, and traditional text and images. Efforts to create hypertextbooks have been ongoing for two decades. Goals include improving exposition through a richer collection of technologies than are available through print textbooks, and increasing student engagement with the material to get them to learn at a higher level in Bloom's taxonomy. See Rössling et al. (2006), Ross and Grinder (2002), and

Shaffer, Akbar, Alon, Stewart, and Edwards (2011) for background on efforts to define and implement the hypertextbook.

Hopefully the technology is in place to finally achieve these goals in the near future. As noted, HTML5 sets the stage as never before. The real game-changing potential that the computer brings to education is its ability to allow the student to interact in non-passive ways with information, and the ability to provide immediate feedback to the student based on these interactions. The vision goes beyond simply creating an interactive eTextbook. Many in the community embrace the concept of a true creative commons for textbook creation, not just in terms of licensing but in terms of allowing for community creation of artifacts. The goal is to allow instructors to modify existing eTextbooks by adopting major portions and then changing sections, or taking chapters from different books and combining them. This aspect of an eTextbook creative commons is already in place with the Connexions (2011) project. However, existing eTextbooks available from Connexions are not heavily interactive in the way that advocates of AV hope to see.

The survey results presented herein lead us to the conclusion that it is easier to integrate a complete block of instruction (either a complete topic or even a semester course) than it is to fit a new piece of instruction such as an AV into an existing presentation on that topic. Instructors are used to the concept of adopting a new textbook for new courses that they teach, and often welcome lecture notes and other class support artifacts. Even for courses taught previously, instructors will adopt new textbooks and new presentations for various topics in the course. A key aspect is that adopting a new chunk of content allows the instructor to completely replace or populate a block of course time, as opposed to squeezing new content or presentation techniques on top of existing material. In the past, most AV developers have implicitly taken the approach that their AV will be integrated into existing presentations. A typical example might be an AV presenting a sorting algorithm, with no tutorial explanation of the algorithm. The idea seems to have been that this visualization can be slipped into lecture or used by students to supplement the textbook for self-study. But this approach has led to the problems noted above, particularly dissonance with other course materials. In contrast, a complete unit of instruction (including AVs) can more easily replace an unsatisfactory existing presentation of the topic.

The final piece of the vision is to have students also participate in the eTextbook creative commons. They could do so by supplying comments, tagging, and rating of content at a detailed level such as for a particular topic, subtopic, or exercise. By interacting with each other and their instructors, they can give direction to improving course materials by drawing attention to what does and what does not work. This enables both a higher level of engagement by the students, and a huge motivating force for improvement of instructional material.

REFERENCES

- AlgoViz.org Bibliography. (2011). Annotated bibliography [of AV literature]. Retrieved from <http://algo-viz.org/biblio>
- Alice Web Site. (2011) Retrieved from <http://www.alice.org>
- ALViE Web Site. (2011). ALViE 3.0 [software]. Retrieved from <http://alvie.algoritmica.org/>
- Baecker, R., & Sherman, D. (1981). *Sorting out sorting*. Retrieved from <http://video.google.com/videoplay?docid=3970523862559774879>
- Bloom, B. S., & Krathwohl, D. R. (1956). *Taxonomy of educational objectives: The classification of educational goals. Handbook I: Cognitive domain*. Harlow, England: Longmans.
- Brabec, F., & Samet, H. (2003). Maryland spatial index demos Web site. Retrieved from <http://donar.umiacs.umd.edu/quadtree/>
- Brown, M. H. (1988). Exploring algorithms using Balsa-II. *Computer*, 21(5), 14–36.
- Brown, M. H., & Sedgewick, R. (1985). Techniques for algorithm animation. *IEEE Software*, 2(1), 28–39.
- Byrne, M. D., Catrambone, R., & Stasko, J. T. (1999). Evaluating animations as student aids in learning computer algorithms. *Computers & Education*, 33(4), 253–278.
- Champine, G. A. (1991). *MIT project Athena: A model for distributed campus computing*. Newton, MA: Digital Press.
- Connexions Web site. (2011). Retrieved from <http://cnx.org/>
- Cooper, S., Dann, W., & Pausch, R. (2000). Alice: A 3-D tool for introductory programming concepts. *Journal of Computing Sciences in Colleges*, 15(5), 107–115.
- Crescenzi, P., Gambosi, G., & Grossi, R. (2006). *Struttura di dati e algoritmi*. Upper Saddle River, NJ: Pearson Education, Addison–Wesley.
- Crescenzi, P., & Nocentini, C. (2007). Fully integrating algorithm visualization into a CS2 course: A two-year experience. In J. Hughes, D. R. Peiris, & P. T. Tymann (Eds.), *Proceedings of the 12th Annual SIGCSE Conference on Innovation and Technology in Computer Science Education*, ITiCSE'07 (pp. 296–300), Dundee, Scotland.
- Dittrich, J.-P., van den Bercken, J., Schafer, T., & Klein, M. (2001). *DSN: Data structure navigator*. Retrieved from <http://dbs.mathematik.uni-marburg.de/research/projects/dsn/>
- Furcy, D. (2007). *JHAVEPOP*. Retrieved from <http://jhave.org/jhavepop/>
- Furcy, D. (2009). JHAVEPOP: Visualizing linked-list operations in C++ and Java. *Journal of Computing Sciences in Colleges*, 25(1), 32–41.
- Galles, D. (2006). *Data structure visualization*. Retrieved from <http://www.cs.usfca.edu/galles/visualization>
- Gettys, J., Newman, R., & Scheifler, R. (1989). *The definitive guides to the X Window system* (Vol. 2) (Xlib Reference Manual for Version 11). Sebastopol, CA: O'Reilly.
- Gustafson, B. E., Kjensli, J., & Vold, J. M. (2011). *Binary treesome Web site*. Retrieved from <http://www.iu.hio.no/~ulfu/AlgDat/applet/binarytreesome>
- Hendrix, T. D., Cross, J. H., II, & Barowski, L. A. (2004, March 3–7). An extensible framework for providing dynamic data structure visualizations in a lightweight IDE. In D. T. Joyce, D. Knox, W. Dann, & T. L. Naps (Eds.), *Proceedings of*

- the 35th SIGCSE Technical Symposium on Computer Science Education* (pp. 387–391), Norfolk, VA.
- Hew, K., & Brush, T. (2007). Integrating technology into K12 teaching and learning: Current knowledge gaps and recommendations for future research. *Educational Technology Research and Development*, 55, 223–252.
- Hundhausen, C. D., & Brown, J. L. (2005). What you see is what you code: A radically dynamic algorithm visualization development model for novice learners. *Proceedings of the 2005 IEEE Symposium on Visual Languages/Human-Centric Computing* (pp. 163–170), Dallas, TX.
- Hundhausen, C. D., & Brown, J. L. (2008). Designing, visualizing, and discussing algorithms within a CS1 studio experience: An empirical study. *Computers & Education*, 50, 301–326.
- Hundhausen, C. D., Douglas, S. A., & Stasko, J. T. (2002). A meta-study of algorithm visualization effectiveness. *Journal of Visual Languages and Computing*, 13(3), 259–290.
- Jain, J., Cross, J. H., Hendrix, T. D., & Barowski, L. A. (2006). Experimental evaluation of animated-verifying object viewers for Java. In M. Burnett & S. Diehl (Eds.), *Proceedings of the 2006 ACM Symposium on Software Visualization* (pp. 27–36), Brighton, UK.
- Jarc, D. J. (1999). *Interactive data structure visualization*. Retrieved from <http://nova.umuc.edu/jarc/idsv>.
- Kaila, E., Rajala, T., Laakso, M.-J., & Salakoski, T. (2009). Effects, experiences and feedback from studies of a program of course-long use of a program visualization tool. *Informatics in Education*, 8(1), 17–34.
- Kann, C., Lindeman, R. W., & Heller, R. (1997). Integrating algorithm animation into a learning environment. *Computers & Education*, 28(4), 223–228.
- Korhonen, A., Malmi, L., Silvasti, P., Nikander, J., Tenhunen, P., Mård, P., . . . Karavirta, V. (2003). *TRAKLA2*. Retrieved from <http://www.cs.hut.fi/Research/TRAKLA2/>
- Laakso, M.-J., Myller, N., & Korhonen, A. (2009). Comparing learning performance of students using algorithm visualizations collaboratively on different engagement levels. *Educational Technology & Society*, 12(2), 267–282.
- Laakso, M.-J., Salakoski, T., Grandell, L., Qiu, X., Korhonen, A., & Malmi, L. (2005). Multi-perspective study of novice learners adopting the visual algorithm simulation exercise system TRAKLA2. *Informatics in Education*, 4(4), 49–68.
- Lahtinen, E., & Ahoniemi, T. (2009). Kick-start activation to novice programming—A visualization-based approach. *Electronic Notes on Theoretical Computer Science*, 224, 125–132.
- Levy, R. B.-B., & Ben-Ari, M. (2007). We work so hard and they don't use it: Acceptance of software tools by teachers. *Proceedings of 12th Annual SIGCSE Conference on Innovation and Technology in Computer Science Education* (pp. 246–250), Dundee, Scotland.
- Levy, R. B.-B., Ben-Ari, M., & Uronen, P. A. (2003). The Jeliot 2000 program animation system. *Computers & Education*, 40, 1–15.
- Ma, L., Ferguson, J., Roper, M., Ross, I., & Wood, M. (2009). Improving the mental models held by novice programmers using cognitive conflict and Jeliot visualisations. *SIGCSE Bulletin*, 41, 166–170.

- Meier, A., Spada, H., & Rummel, N. (2007). A rating scheme for assessing the quality of computer-supported collaboration processes. *International Journal of Computer Supported Collaborative Learning*, 2(1), 63–86.
- Moreno, A., Joy, M., Myller, N., & Sutinen, E. (2010). Layered architecture for automatic generation of conflictive animations in programming education. *IEEE Transactions on Learning Technologies*, 3, 139–151.
- Moreno, A., Sutinen, E., Bednarik, R., & Myller, N. (2007). Conflictive animations as engaging learning tools. In R. Lister & Simon (Eds.), *Seventh Baltic Sea Conference on Computing Education Research (Koli Calling 2007)* (pp. 203–206), Koli National Park, Finland.
- Moskal, B., Lurie, D., & Cooper, S. (2004). Evaluating the effectiveness of a new instructional approach. *SIGCSE Bulletin*, 36, 75–79.
- Myller, N., Bednarik, R., Sutinen, E., & Ben-Ari, M. (2009). Extending the engagement taxonomy: Software visualization and collaborative learning. *Transactions on Computing Education*, 9(1), 1–27.
- Naps, T. L., Cooper, S., Koldehofe, B., Leska, C., Rossling, G., Dann, W., . . . McNally, M. F. (2003a). Evaluating the educational impact of visualization. *SIGSE Bulletin*, 35(4), 124–136.
- Naps, T. L., Eagan, J. R., & Norton, L. L. (2000). JHAVÉ—An environment to actively engage students in Web-based algorithm visualizations. *SIGCSE Bulletin*, 32, 109–113.
- Naps, T. L., Rossling, G., Almstrum, V., Dann, W., Fleischer, R., Hundhausen, C. D., . . . Velazquez-Iturbide, J. (2003b). Exploring the role of visualization and engagement in computer science education. *SIGCSE Bulletin*, 35(2), 131–152.
- Pane, J. F., Corbett, A. T., & John, B. E. (1996). Assessing dynamics in computer-based instruction. In R. Bigler, S. Guest, & M. J. Tauber (Eds.), *Proceedings of the ACM Conference on Human Factors in Computing Systems* (pp. 197–204), Vancouver, BC, Canada.
- Pierson, W. C., & Rodger, S. H. (1998). Web-based animation of data structures using JAWAA. *SIGCSE Bulletin*, 30, 267–271.
- Rajala, T., Laakso, M.-J., Kaila, E., & Salakoski, T. (2007). ViLLe—A language-independent program visualization tool. In L. Raymond & Simon (Eds.), *Seventh Baltic Sea Conference on Computing Education Research (Koli Calling 2007)*, Koli National Park, Finland.
- Rajala, T., Laakso, M.-J., Kaila, E., & Salakoski, T. (2008). Effectiveness of program visualization: A case study with the ViLLe tool. *Journal of Information Technology Education*, 7, 15–32.
- Rodger, S.H. (2008). *JFLAP Web site*. Retrieved from <http://www.jflap.org>.
- Rodger, S. H., Wiebe, E., Lee, K. M., Morgan, C., Omar, K., & Su, J. (2009). Increasing engagement in automata theory with JFLAP. In S. Fitzgerald, M. Guzdial, G. Lewandowski, & S. A. Wolfman (Eds.), *Proceedings of the 40th ACM Technical Symposium on Computer Science Education* (pp. 403–407), Chattanooga, TN.
- Ross, R., & Grinder, M. (2002). Hypertextbooks: Animated, active learning, comprehensive teaching and learning resources for the Web. In S. Diehl (Ed.), *Software visualization* (Lecture notes in Computer Science 2269) (pp. 269–284). Berlin, Germany: Springer.

- Rössling, G., Naps, T. L., Hall, M. S., Karavirta, V., Kerren, A., Leska, C., . . . Velazquez-Iturbide, J. (2006). Merging interactive visualizations with hypertext-books and course management. *SIGCSE BULLETIN*, 38(4), 166–181.
- Rössling, G., Schuler, M., & Freisleben, B. (2000). The ANIMAL algorithm animation tool. In J. Tarhio, S. Fincher, & D. Joyce (Eds.), *Proceedings of the 5th Annual SIGCSE/SIGCUE ITiCSE Conference on Innovation and Technology in Computer Science Education* (pp. 37–40), Helsinki, Finland.
- Sanders, D., Heeler, P., & Spradling, C. (2001). Introduction to BlueJ: A Java development environment. *Journal of Computing Sciences in Colleges*, 16(3), 257–258.
- Shaffer, C. A., Akbar, M., Alon, A., Stewart, M., & Edwards, S. H. (2011). Getting algorithm visualizations into the classroom. *SIGCSE Bulletin*, 11, 129–134.
- Shaffer, C. A., Cooper, M., Alon, A., Akbar, M., Stewart, M., Ponce, S., & Edwards, S. H. (2010). Algorithm visualization: The state of the field. *ACM Transactions on Computing Education*, 10, 1–22.
- Shaffer, C. A., Naps, T. L., & Fouh, E. (2011). Truly interactive textbooks for computer science education. In G. Rössling (Ed.), *Proceedings of the 6th Program Visualization Workshop* (pp. 97–103), Darmstadt, Germany.
- Stasko, J. T. (1990). Tango: A framework and system for algorithm animation. *Computer*, 23(9), 27–39.
- Stasko, J. T. (1992). Animating algorithms with Xtango. *SIGACT News*, 23, 67–71.
- Stasko, J. T., Badre, A., & Lewis, C. (1993). Do algorithm animations assist learning? An empirical study and analysis. In S. Ashlund, K. Mullet, A. Henderson, E. Hollnagel, & T. N. White (Eds.), *Proceedings of the INTERACT '93 and CHI '93 Conference on Human Factors in Computing Systems* (pp. 61–66), Amsterdam, The Netherlands.
- Stasko, J. T. (1997). Using student-built algorithm animations as learning aids. In C. M. White, C. Erickson, B. J. Klein, & J. E. Miller (Eds.), *Proceedings of the 28th SIGCSE Technical Symposium on Computer Science Education* (pp. 25–29), San Jose, CA.
- Stasko, J. T. (1998). *JSamba*. Retrieved from <http://www.cc.gatech.edu/gvu/softviz/algoanim/jsamba/>
- Stasko, J. T. (2001). *Polka animation system*. Retrieved from <http://www.cc.gatech.edu/gvu/softviz/par-viz/polka.html>
- Stern, L. (2001). *Algorithms in action*. Retrieved from <http://www.cs.mu.oz.au/aia/>
- Urquiza-Fuentes, J., & Velazquez-Iturbide, J. (2009). A survey of successful evaluations of program visualization and algorithm animation system. *ACM Transactions on Computing Education*, 9(2), 1–21.
- Virginia Tech Algorithm Visualization Research Group Web site. (2011). Retrieved from <http://research.cs.vt.edu/AVresearch/>