

2019

Using the VAST Challenge in Undergraduate CS Research

Christopher P. Andrews
Middlebury College

R. Jordan Crouser
Smith College, jrouser@smith.edu

Follow this and additional works at: https://scholarworks.smith.edu/csc_facpubs



Part of the [Computer Sciences Commons](#)

Recommended Citation

Andrews, Christopher P. and Crouser, R. Jordan, "Using the VAST Challenge in Undergraduate CS Research" (2019). Computer Science: Faculty Publications, Smith College, Northampton, MA.
https://scholarworks.smith.edu/csc_facpubs/150

This Article has been accepted for inclusion in Computer Science: Faculty Publications by an authorized administrator of Smith ScholarWorks. For more information, please contact scholarworks@smith.edu

Using the VAST Challenge in Undergraduate CS Research*

Christopher P. Andrews¹, R. Jordan Crouser²

¹Computer Science Department

Middlebury College Middlebury, VT, 05753

candrews@middlebury.edu

²Computer Science Department

Smith College, Northampton, MA, 01063

jcrouser@smith.edu

Abstract

The Visual Analytics Science and Technology (VAST) Challenge is a yearly competition designed to push forward visual analytics research through synthetic, yet realistic analytic tasks. In this paper, we discuss the challenges and the successes we have experienced incorporating the VAST Challenge and associated datasets into undergraduate research programs at two liberal arts colleges. We advocate for increased undergraduate participation in this and similar competitions, arguing they afford unique opportunities for positive development in early researchers.

1 Introduction

There have been many benefits ascribed to undergraduate research (UR) programs, including enhanced domain knowledge and skills, honing of analytic skills, increased interest and retention, greater self-confidence, gains in critical thinking, and validation of post-graduation choices [14, 7, 9]. While Seymour et al. found that some of these may be more anecdotal than rigorously demonstrated [14], there is a general agreement that students benefit particularly from the apprentice model of research that is carried out in many small undergraduate institutions. In addition, for these gains to be realized, students

*Copyright ©2019 by the Consortium for Computing Sciences in Colleges. Permission to copy without fee all or part of this material is granted provided that the copies are not made or distributed for direct commercial advantage, the CCSC copyright notice and the title of the publication and its date appear, and notice is given that copying is by permission of the Consortium for Computing Sciences in Colleges. To copy otherwise, or to republish, requires a fee and/or specific permission.

must have **legitimate participation** in the research – they must be engaged, making decisions, and have personal stake in the research program rather than simply following directives [7].

Few undergraduate computer science (CS) programs offer (and even fewer require) courses in research methods, and many students begin their study of CS at the undergraduate level. As a result, students just starting out on a research project frequently lack foundational research skills, the remediation of which can consume valuable time [6, 17]. There is a need to balance getting students’ technical abilities to a point where they can pursue a meaningful research question and ensuring that sure they have a project in which they can take some ownership, while simultaneously continuing to advance faculty research agendas at a reasonable pace.

While multi-year projects that establish continuity and pull in new students who can grow to be contributors have been proposed as the ideal [17], the lure of many compelling internship opportunities has made recruiting and retaining CS students for multiple years increasingly difficult. Hadfield and Schweitzer suggest addressing these issues by distilling the research into small (but non-trivial) problems to promote a sense of ownership within each student, with the faculty member providing mentorship over the larger process [6]. Koeller suggests that, “the most efficient use of undergraduate resources is in implementation,” which should be accessible to all students who have progressed beyond data structures [8]. While these suggestions converge to a consistent message, practical and reproducible models which can be readily adapted by new faculty at primarily-undergraduate institutions (PUIs) are sorely lacking.

In this paper, we discuss our experiences using the annual IEEE Visual Analytics Science and Technology (VAST) Challenge [3] as a component of undergraduate research at two PUIs, and how it has helped us to address some of these challenges. We hope that this paper will serve as a resource for others in the VIS community wishing to increase participation of undergraduate students in their research programs, and that participation of undergraduate students in the VAST Challenge will continue to grow.

2 The VAST Challenge

The VAST Challenge is an annual competition in which visual analytics researchers, software developers, and designers pit their best tools and techniques against realistic analysis tasks. It has been held annually since 2006, providing synthetic tasks and datasets that mirror the real world in both scale and complexity. The data, tasks, and selected solutions are made publicly available through the Visual Analytics Benchmark Repository¹.

¹<https://www.cs.umd.edu/hcil/varepository>

2.1 Origin of the Challenge

The field of visual analytics has evolved into a complex hybrid of information visualization, interaction design, machine learning, and human psychology. Due to the dynamic interplay between human and machine analysis, standard benchmarking tasks and traditional measures of success have limited utility. The need for new evaluation mechanisms and realistic benchmark tasks provided the motivation for the challenge [3]. The challenge also provides researchers who would otherwise have no access to practicing analysts or data that is frequently sensitive or confidential with realistic tasks and datasets.

2.2 Data and Mechanics

The challenge typically consists of a collection of “mini-challenges” and an overarching “grand challenge” that ties the challenges together. Submissions are reviewed by the committee, professional analysts, and peer researchers. Each competing team has the opportunity to present a poster, participate in the challenge workshop held at the VAST conference, and submit a short two-page paper for inclusion in the workshop proceedings.

Tasks have included piecing together stories from text fragments (e.g., news articles, intelligence reports, and social media), finding patterns in GPS traces and sensor logs, and analyzing satellite images, video and sound recordings. Frequently, the challenges make use of heterogeneous data sets, requiring analysis that combines and relates different data sources and types.

Each challenge also comes wrapped in an accompanying story. Over the years, participants have had to contend with bio-terrorists, exotic animal smugglers, arms dealers, corporate hijinks, and polluters (to name just a few). The background stories serve as motivation to find the solution, as well as making the analysis fun and engaging [10, 13]. Finally, a key feature of each of the challenge data sets is that each one has an **embedded ground truth**, which allows for comparison between solutions, and motivates participants by indicating that a solution exists and that it is achievable by researchers and students [5].

2.3 Other Uses of the VAST Challenge Data

Due to the availability of the datasets through the benchmark repository, participation in the challenge is not required to make use of the data, and it is used fairly extensively [11]. Companies report using the data to test their analytic tools [13, 10], and researchers use the data sets to explore new analytic techniques and attempt to better understand the analytic process [5].

The challenge has also been used in a number of university courses as a source of projects [19]. Rohdantz et al. report that the novelty of the scenarios

“requires students to synthesize their learned knowledge to create novel approaches”. The students are reported to overwhelmingly find the scenarios to be engaging and useful in identifying requirements and evaluating the analytic tools they developed in the classes [12].

3 Participating in the Challenge

While the VAST Challenge data sets can be used in a variety of ways to support UR, it can be very rewarding to participate in the actual challenge. The reviewing process provides valuable feedback, and the opportunity to present a poster and write a short paper is a nice entry into conference participation. The students also get the experience of working on the same problems as graduate and industry researchers. However, assembling a submission is not without its own challenges. To date, Middlebury College has sent two official submissions and Smith College has sent one.

3.1 Challenges to Participation

A key challenge is timing. The challenges are typically released late in late April, early May, with submission due early in July, which does not line up well with the academic calendar. The reality for many research programs conducted at PUIs is that we often do not have multi-year student continuity in our research labs, and summer is the main time to engage in research. Thus, the start of summer typically needs to be devoted to onboarding new students who, in addition to being new to research, may also lack a background in many of the skills required to develop a submission. Producing a competitive submission requires a great deal of effort, and it is quite challenging to work toward an early July deadline under these conditions. The deadline roughly bisects the summer, which leaves time for a second project, but one that must be limited in scope based on the schedule.

3.2 Some Alleviating Techniques

Our hope is that our experiences show that these challenges are not insurmountable. One strategy is to try to place useful electives (e.g., information visualization, data science, software development, etc...) in the fall when planning course offerings, to create a pool of students with known skills.

Another technique is to be strategic about which portions of the challenge to undertake. The Smith entry was done by a pair of rising sophomores after only a single computer science course. They focused on the 2016 Mini-Challenge 1, a design challenge that required participants to design an innovative interactive visual interface which could enable security investigators to conduct real-time

analysis of multimodal streaming data about a resort property [4]. The completed design, rendered with PowerPoint was not particularly novel in the context of more advanced submissions to the challenge, but the development of a multiple-display, coordinated system was novel *to the students themselves*, and gave them the opportunity to learn valuable skills such as rapid prototyping in the context of interface design.

The two Middlebury entries were for the 2015 Mini-Challenge 1 (MC1), which dealt with patterns of behavior at an amusement park [18], and 2016 Mini-Challenge 2 (MC2), which required participants to analyze environmental sensors and “prox-card” records for the main building of the fictitious GASTech company [4]. In both cases, the solutions were built up as a series of connected visualizations using Middguard, our in-house framework. By pursuing designs consisting of simple linked visualizations and without extensive “computational analytics” required, students could be given ownership over small, self-contained pieces that become useful tools when connected together.

The Middlebury entries could be a little more ambitious both because the students were more advanced (a rising senior on MC1 and a rising senior and a rising junior on MC2), but also because we made use of an apprentice model with a faculty member working alongside the students. Hunter et al. note that “apprentice-style UR fits a theoretical model of learning advanced by constructivism” [7]. This style of research allows the student to make their own meaning through interaction with and support from their faculty mentor. This leads to a number of direct benefits to the students: the students gain new technical knowledge and skills, learn how research is conducted and disseminated in their field, and develop their ability to approach complex problems. In all cases, by the end of the challenge, the students were equal contributors in terms of both implementation and analysis. This model also makes it easier to keep them from losing a lot of time following unpromising paths by steering their efforts. In addition, working alongside a faculty member on the same problem (and one for which the faculty member doesn’t already have an answer) is a novel experience for most students — one that seems to give them more of a stake in the process and the outcome.

In the same vein, we have found it useful for the faculty member to start work on the challenge well in advance of the students. This allows time to strategize and pick the most relevant challenge, and to clean up the realistically “messy” data. The heterogeneous nature of the data sets also means that some early work can focus on finding ways to correlate the data. While these are very realistic data problems, solving them ahead of time allows the students to be able to focus on the actual analytic problem.

4 Balancing Education and Research Agendas

A major challenge in supporting UR is in balancing the sometimes orthogonal goals of student education and furthering faculty research agendas. While participating the challenge makes a good, self-contained project that teaches students about one aspect of research, it is more difficult to shape them into a more sustained, productive research agenda, especially in light of the way the challenge bisects the summer research time. In this section we will discuss some of the other ways that we have engaged with the VAST Challenge.

4.1 Onboarding New Undergraduate Researchers

Between 6 and 8 students conduct research in the Human Computation and Visualization Lab at Smith College each year, with a subset of these positions being reserved specifically for undeclared students in their first two years. During the first week of the program, we use the VAST Challenge as an opportunity for these mixed groups of students to work together toward a common goal, flex their programming muscles, and run into roadblocks that require them to investigate unfamiliar techniques. We have found that opportunities to interact with peers within the low stakes / high energy context of the VAST Challenge are useful in helping the adviser to identify strengths and growth areas within the cohort, and working collaboratively can help new students integrate into the lab culture.

4.2 Engaging Students in Open Source Projects

At Middlebury, our first attempt at the challenge was essentially an onboarding exercise – an attempt to both get a new lab up and running and a new undergraduate researcher thinking analytically. While we did not have time to put together an official submission, the difficulties we faced building the system that we envisioned inspired us to build Middguard².

Middguard is a web-based framework for building bespoke visual analytics tools requiring a collection of heterogeneous, but connected visualizations – the basic target being the development of tools to solve new, heterogeneous data challenges, such as those poised by the VAST Challenges. The framework allows developers to create visualization modules that can be bundled together and deployed to create more complex tools without the need for the individual tools to have knowledge of one another. The framework is responsible for connecting data models to the visualizations, fetching data from the server, and broadcasting selection events to all visible modules so that the views can work together.

²<https://github.com/AnalyticArtsLab/middguard>

In addition to making a contribution to the lab’s research agenda, the project became an important part of the education of our lead developer, as he built upon it for both a senior project and senior thesis.

4.2.1 Data Generation

In 2019, Smith partnered with Pacific Northwest National Laboratory (PNNL) to generate a challenge dataset. The 2019 VAST Challenge centers on a natural disaster. A group of 28 volunteer Smith undergraduate students joined 1 Smith professor and 8 PNNL researchers in a role-playing exercise, interacting on a social media feed to generate over 2,100 realistic posts covering the time before, during, and after an earthquake. These exercises are particularly interesting because they flip the **consumption model** frequently adopted in undergraduate data science education, engaging undergraduate students as legitimate co-creators not only of data that might benefit students like themselves, but which would be used to advance the state of the art in visual analytics research. This is a powerful shift as it allowed the students to feel they were making direct contributions to real research.

5 Reflections on student engagement and outcomes

One tangible measure of success for interventions such as these is the generation of research products during the undergraduate research experience. By that metric, participating in the challenge was a success. The Middlebury students have made contributions to multiple short papers [1, 2], and the work on Middguard led to a poster [15], and a thesis [16]. Further, Middlebury won an award for the 2015 submission, indicating that our model could produce a competitive team and that working at the undergraduate level is not in and of itself a barrier to success (two further undergraduate teams have since won awards as well).

Of course, it is also important to consider other metrics in our assessment of the utility of the VAST Challenge in promoting undergraduate scholarship. A commonly stated motivation for supporting undergraduate research is to increase the pipeline to graduate school [7, 14]. However, while enrollment in computer science courses is healthy and the number of majors continues to rise, the current economic climate means that only a small fraction of computer science majors go on to attend graduate school immediately after completing their undergraduate degree. Indeed, none of the students described above were on a path toward doctoral programs in CS, and their experience working on the challenges did not change that. Instead, the student who have graduated have gone on to successful positions in industry (at Google, Morgan Stanley, and Epic, to name a few).

Of course, we mentioned a number of less tangible effects attributed to UR. Our sample size is too small to replicate any of the earlier results, however, our informal conversations with the students afterward revealed signs of many of these. One student reported, “I think I gained a weird sense of confidence through this challenge. The comments are very informative and sincere so they oddly encouraged me to try new problems and take on challenges that were beyond my scope.”

As for the challenge itself, a common sentiment was that it “is very fun because the challenge simulates real world problems” and the realism inspired “a sense of interest and excitement”. They talked about all of the important skills they gained from visualization and design, to working with databases and servers, and working on existing codebases. The students also appreciated the different perspective on visualization they gained, with one describing how it encouraged her to be more mindful of users: “visualization requires you to be an empathetic problem solver with an eye for design”.

6 Conclusion

Our goal in writing this paper is to encourage the use of the VAST Challenge and its datasets in UR. While principally designed to engage researchers and developers at research institutions and businesses, the problems are within reach for undergraduates. The scenarios have many compelling features: they model real analytic problems, they are sufficiently complex so that they can’t be solved by in a couple of hours, they have a ground truth so results can be evaluated, and the scenarios themselves are engaging and fun.

The diversity of the data also provides a number of different challenges and could provide interesting material for research outside of the visual analytics domain. While the challenges are obviously designed to push forward research in visual analytics, we feel that engagement with the challenge datasets need not be limited to visual analytics or even data science researchers. The data could also be fodder for undergraduate projects in a number of different areas within computer science such as natural language processing, machine learning, and image processing.

We have laid out some of the potential benefits and challenges of participating in the challenge itself, as well as discussing some of the other ways we have incorporated the datasets into our research. We would also like to encourage more undergraduate teams, especially from PUIs, to participate in the challenge itself³. An ideal outcome would be the development of a community of faculty members working to bring visual analytics research into the undergraduate sphere.

³We encourage teams from PUIs to mark their submissions as such

7 Acknowledgments

This work was made possible by undergraduate researchers at Smith College and Middlebury College, in particular the contributions of Dana Silver '16, Julian Billings '16, Lily Taub '17, Shannon Ovitt '18, Ji Won Chung '18, Zheng (“Alice”) Mu '19J, Tiffany Xiao '20J, and Ananda Montoly '22. The authors also wish to thank the Smith Undergraduate Research Fellowship Program, the Middlebury Undergraduate Research Office and the Bicentennial Fund for Research Partnerships.

References

- [1] C. Andrews and J. Billings. Middguard at DinoFun World. In *Visual Analytics Science and Technology (VAST), 2015 IEEE Conference on*, pages 129–130, October 2015.
- [2] C. Andrews, L. Taub, and S. Ovitt. Middguard at gastech. In *Visual Analytics Science and Technology (VAST), 2016 IEEE Conference on*, Oct 2016.
- [3] Kristin Cook, Georges Grinstein, and Mark Whiting. The vast challenge: History, scope, and outcomes, 2014.
- [4] R Jordan Crouser, Kristin Cook, John Fallon, Georges Grinstein, Kristin Liggett, Danko Nebesh, Diane Staheli, Mark A Whiting, and Kirsten Whitley. VAST Challenge 2016: Streaming Visual Analytics. Technical report, MIT Lincoln Laboratory, October 2016.
- [5] Carsten Görg, Zhicheng Liu, and John Stasko. Reflections on the evolution of the Jigsaw visual analytics system. *Information Visualization*, 13(4):336–345, October 2014.
- [6] Steven Hadfield and Dino Schweitzer. Building an undergraduate computer science research experience. pages 1–6. IEEE, October 2009.
- [7] Anne-Barrie Hunter, Sandra L. Laursen, and Elaine Seymour. Becoming a scientist: the role of undergraduate research in students’ cognitive, personal, and professional development. *Science Education*, 91(1):36–74, January 2007.
- [8] Andreas Koeller. Experiences with Student Research at a Primarily Undergraduate Institution. page 7, 2005.

- [9] D. Lopatto. Survey of Undergraduate Research Experiences (SURE): First Findings. *Cell Biology Education*, 3(4):270–277, December 2004.
- [10] Pascale Proulx and Casey Canfield. The beneficial role of the VAST Challenges in the evolution of GeoTime and nSpace2. *Information Visualization*, 14(1):3–9, January 2015.
- [11] Swetha Reddy, Manas Desai, Catherine Plaisant, and Jean Scholtz. A usage summary of the VAST Challenge Datasets. Technical Report, HCIL, University of Maryland and PNNL, October 2010.
- [12] Christian Rohrdantz, Florian Mansmann, Chris North, and Daniel A Keim. Augmenting the educational curriculum with the Visual Analytics Science and Technology Challenge: Opportunities and pitfalls. *Information Visualization*, 13(4):313–325, October 2014.
- [13] Jean Scholtz, Catherine Plaisant, Mark Whiting, and Georges Grinstein. Evaluation of visual analytics environments: The road to the Visual Analytics Science and Technology challenge evaluation methodology. *Information Visualization*, 13(4):326–335, October 2014.
- [14] Elaine Seymour, Anne-Barrie Hunter, Sandra L Laursen, and Tracee DeAntoni. Establishing the benefits of research experiences for undergraduates in the sciences: First findings from a three-year study. *Science education*, 88(4):493–534, 2004.
- [15] Dana Silver. MiddGuard: A Framework for Collaborative and Extensible Visual Analytics. In *Proceedings of CCSCNE-2016*, page Poster #5049, April 2016.
- [16] Dana Silver. *Middguard: Collaborative and Extensible Visual Analytics*. PhD thesis, Middlebury College, May 2016.
- [17] Elaine Wenderholm. Challenges and the Elements of Success in Undergraduate Research. 36(4):3, 2004.
- [18] Mark Whiting, Kristin Cook, Georges Grinstein, John Fallon, Kristen Liggett, Diane Staheli, and Jordan Crouser. VAST Challenge 2015: Mayhem at dinofun world. In *2015 IEEE Conference on Visual Analytics Science and Technology (VAST)*, pages 113–118, Chicago, IL, USA, October 2015. IEEE.
- [19] Mark A. Whiting, Chris North, Alex Endert, Jean Scholtz, Jereme Haack, Carrie Varley, and Jim Thomas. VAST contest dataset use in education. pages 115–122. IEEE, 2009.