Visualization Viewpoints

Editor:
Theresa-Marie Rhyne

Apply or Die: On the Role and Assessment of Application Papers in Visualization

Gunther H. Weber
Lawrence Berkeley National Laboratory

Sheelagh Carpendale
University of Calgary

David Ebert
Purdue University

Brian Fisher
Simon Fraser University

Hans Hagen
University of Kaiserslautern

Ben Shneiderman
University of Maryland

Anders Ynnerman
Linköping University

In his controversial 2004 paper "On the Death of Visualization," Bill Lorensen asked, "Can it survive without customers?"[1] Since then, concerns about the need to document the field's impact on applications has inspired outreach to application stakeholders and efforts to increase the presence and visibility of application-driven research at visualization conferences. Despite a general agreement that application-oriented papers are of value—application papers have won several honorable mentions, for example, at the 2016 Scientific Visualization (SciVis) and Visual Analytics Science and Technology (VAST) Conferences—it remains an open question as to what criteria to use when judging whether an application paper merits acceptance. This is all too clear given the range of evaluations that this paper type receives in conference reviews and the diversity of opinions voiced by reviewers. These discussions frequently focus on how best to achieve the goals of relevance to application problems, in addition to the more familiar criteria of technical excellence and methodological rigor that are expected for IEEE venues.

A panel at the 2016 IEEE Visualization (VIS) Conference in Baltimore, Maryland, titled "Application Papers: What Are They, and How Should They Be Evaluated?" brought together leading researchers in the three main visualization areas—information visualization (Carpendale and Shneiderman), scientific visualization (Hagen and Ynnerman), and VAST (Ebert and Fisher)—to explore these challenges. Our goal was to gain a better understanding of prevalent views in the visualization community and to start a discussion on how to evaluate application papers more consistently. This article summarizes points that arose during that panel discussion in an effort to derive a first set of conclusions from them and prompt further dialogue.

Current Trends Favoring Application Papers

Three major trends that favor an increased emphasis on application papers emerged during the VIS panel discussion:

- the need for bigger teams with an increased emphasis on integration and infrastructure in order to realize the ambitious visions outlined in proposals,
- changing research policies, and
- the increasing number of researchers involved in interdisciplinary research.

Realizing the "Grand Vision"

One of the main drivers of visualization research is to achieve the grand vision of enabling data understanding in science, engineering, and society.

Published by the IEEE Computer Society

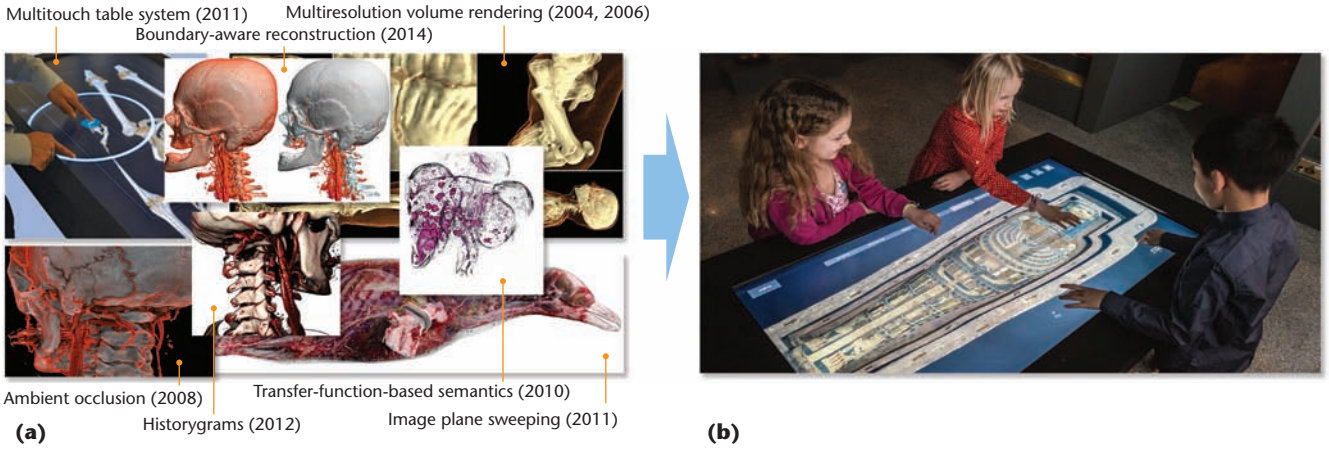

Figure 1. Visualization table. (a) Based on a progression of visualization research papers published over the past decade,[2–9] the visualization table developed at the Norrköping Visualization Center C, Linköping University, Sweden, impacts medicine as well as science communication. (b) Using the table, visitors to the Mediterranean Museum in Stockholm are interacting with a combined surface and volumetric visualization of the mummy Neswaiu.

Technical visualization papers play an important role in accomplishing this mission by advancing the state of the art in algorithms and methods. These ultimately serve as building blocks or "dots" for achieving this grander vision. Realizing the overall vision also requires significant infrastructure-building and integration work to "connect the dots" in order to build a complete large-scale system. From this perspective, application papers are success stories that demonstrate that we are achieving our goal and accomplishing this grander vision.

Figure 1 shows a research project that illustrates how to combine and tailor the contributions of many technical visualization papers into a widely used system. The Norrköping Visualization Center C at Linköping University, Sweden, developed a visualization table for the medical domain that is also used in science communication. Based on a progression of research papers published over the past decade,[2–9] the visualization table is used in science communication. Figure 1b shows visitors at the Mediterranean Museum in Stockholm interacting with the visualization table, which is displaying a combined surface and volumetric visualization of the mummy Neswaiu. A fuller account of this project was published in the *Communications of the ACM*.[10] Papers of this kind provide valuable success stories to the visualization community.

Integrating visualization techniques into a system, including work with application professionals, the combination and refinement of individual methods, and detail work to create usable systems, requires large teams with dedicated researchers working at the intersection of visualization and application. These researchers bridge the gap between fundamental visualization technology research and the grander vision of data understanding. Application papers document these efforts, providing valuable knowledge to peers on how to build large-scale systems. Furthermore, such researchers who invest an immense effort in infrastructure building and integration need a clear career path, and in the traditional evaluation model, this requires publications. Application-oriented papers allow these researchers to earn recognition and prosper in their careers. The success of these researchers in turn supports the field, demonstrating that innovation in newly developed methods can be used in and will impact real-world applications.

Application Impact and Societal Relevance

Many funding sources for visualization, such as from the US National Science Foundation (NSF), focus on basic research. However, many visualization research groups also rely on funding sources with an application-driven mission. Government agencies such as the US Department of Energy (DOE) and the National Institutes for Health (NIH) as well as private foundations like the Stiftung für Innovation and the Keck Foundation have requirements that visualization must help to accomplish; they do not propose the development of new visualization techniques as the ultimate goal. Funding from these sources is increasingly dependent on a track record of refereed publications that demonstrate a research group's accomplishments and the ability of the of visualization field to help the funding organization accomplish its mission.

To remain competitive, groups funded by these institutions must publish papers that focus on the combination, adaptation, and refinement of existing techniques to address a particular problem. More importantly, application papers also provide the success stories necessary to show visualization's

potential in a range of application domains and to obtain funding in the first place. If we shun application papers from visualization venues, we rob ourselves and future generations of young visualization researchers of important funding opportunities. This trend is likely to be exacerbated because even funding agencies that focus on basic research are starting to encourage or even require collaborations with application scientists and professionals—for example, applications are one method of showing broader impact in NSF proposals.

Emergence of Interdisciplinary Researchers

Computation is playing an increasingly important role in many sciences, and the data deluge from simulations and experiments requires increasingly sophisticated computational analysis methods. As a result, a new generation of interdisciplinary students and researchers is emerging, as researchers in many disciplines become better versed in computational methods and computer science programs create new interdisciplinary degrees in computational science and engineering, such as computational chemistry and computational neuroscience. Many schools are also creating new interdisciplinary programs in data science.

This trend has grown to the point that some students even pursue doctoral degrees in two disciplines. These researchers are looking for appropriate venues to publish their research results, which can be extremely interesting and valuable to our community. Thus, we now have the choice between welcoming their contributions or pushing them to other places. The future of the visualization community—whether it will grow, stagnate, or even shrink—may depend on this decision.

Benefits and Contributions of Application Papers

Although the increasing importance of application papers is in part motivated by necessity—developing research infrastructure, funding, and adapting to new research and education trends—these papers also have a beneficial influence on visualization. Most importantly, application-driven work guides basic visualization research, ensuring it solves relevant problems. It can also accelerate the rate of scientific discovery due to the challenging complexity of real-world problems. Furthermore, it encourages us to evaluate new visualization research based on real data. Application papers build a tacit knowledge about problem domains and communicate successes and failures to peers and funding agencies.

Improving the Quality of Basic Research

Traditional research models posit a linear path that separates basic research from applications. This assumes that knowledge flows from basic curiosity-driven research into applied research. To a certain degree, there is disdain in academic circles for applied research. This belief was promoted in Vannevar Bush's 1945 manifesto on "Science: The Endless Frontier."[11] However, numerous critics have pointed out the flaws in this conceptualization, citing substantial successes from integrated application-oriented work.[12,13] For example, applied research labs, such as the Bell Labs and US national laboratories, contribute to both basic and applied research. These applied projects have produced Nobel Prize–winning research. Similarly, the Fraunhofer Gesellschaft in Germany as well as the automobile and aerospace industries successfully produce basic research that is driven by applications.

It is possible to measure the positive impact of close collaboration between academic and applied researchers. For example, at the 2014 ACM Knowledge Discovery and Data Mining (KDD) Conference, out of 1,036 submissions, the papers written in collaboration between academic and industry researchers received a statistically significantly higher rating than pure academic or pure industry papers.

There are many ways in which applications promote basic research that accelerates the development of new visualization techniques. First, applications provide challenging problems for visualization, which often generate new research questions that then trigger new directions in basic research (mathematics, computer science, and visualization). Real-world problems have strict requirements on solution quality and ensure that a developed theory or technique can handle meaningful data, providing an alternative form of empirical test of a basic research hypothesis. The result is a positive feedback loop, where applied and basic research cross-fertilize each other.

EventFlow (hcil.umd.edu/eventflow) exemplifies the integration of basic and applied research (see Figure 2).[14] EventFlow is a novel tool for event analytics that provides an efficient means to visualize databases of records containing time-stamped categorical event sequences. Using a theory-driven development approach, collaborations with healthcare professionals on real problems have made it possible to develop appropriate control panels and innovative visual designs that present data in ways that support problem solving.

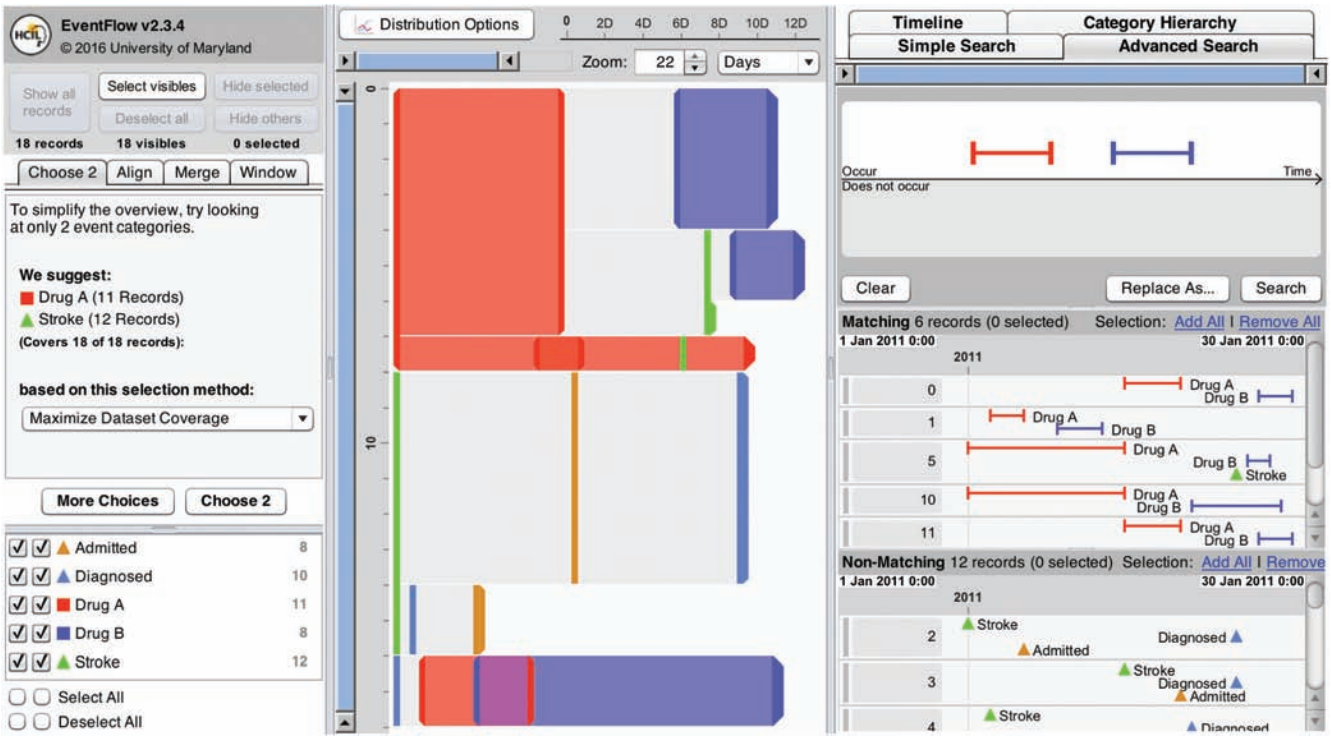

Figure 2. EventFlow event analytics tool. This figure illustrates EventFlow with a healthcare example, showing a small database of 18 patient records with three point events (admitted, diagnosed, and stroke) and two interval events (drugs A and B). The overview in the center shows that seven patients received drug A, after which four shifted to drug B, and three of them had a stroke. For eight patients, the first event in their record is a stroke. On the right, a search for patients who received drug A followed by drug B has located six matches.

Application professionals also provide real datasets for the development and evaluation of visualization methods, and application-oriented papers make this data and the insights to be gleaned from it available to the larger visualization community. For example, the datasets used by the IEEE SciVis contest and VAST Challenge are now often used as examples when describing and evaluating new algorithms in technical papers. Finally, discussions with application professionals can be inspiring and illuminating because they provide a fresh perspective. Seeing how new techniques are used to gain new knowledge is also an extremely rewarding experience.

The types of contributions that an application paper can make to improving basic research include the following:

- *Contribution 1* (*C1*): Put previous research contributions in an application context, and describe a combination of methods to accomplish a visualization/analysis objective.
- *C2*: Present valuable application-specific contributions, such as tailoring existing methods.
- *C3*: Provide a foundation for future visualization research, and pose research challenges.
- *C4*: Provide new means (datasets and evaluation criteria) for evaluating visualization methods.

Building Deep Knowledge

A good application paper will go beyond proposing a method for visualizing data from a given application domain and will demonstrate support for the work processes, including cognitive work, that experts currently use. The best application papers go even further, beyond current work practices to propose new technological capabilities for domain experts that will enable them to improve on their current practices. For these projects, visualization designers must conduct a deep analysis of both tacit and explicit knowledge in the domain. The contributions from these papers include the following:

- *C5*: Convey the "language" of the application domain and assessment methods for particular situations, thus facilitating deep collaborations with application professionals.
- *C6*: Enable the visualization community to learn how researchers in other communities approach problems as individual analysts and in collaboration with others.
- *C7*: Document successes and failures of visualization approaches for a given application domain, and derive lessons learned from the visualization expert's perspective.
- *C8*: Discuss both design methodologies and in-

dividual design decisions that have proven to successfully address the users' needs.

Papers that make contributions at the level of cognitive and collaborative work practices will more deeply impact how the visualization community approaches visualization for a given application domain.

Providing Success Stories

In addition to providing success stories for funding agencies, application papers also communicate the successful application of visualization in many different fields. These success stories can help visualization researchers identify the methods that work and those that need improvement as well as see how their techniques and algorithms are applied to real-world data. The latter is helpful, for example, in steering future work on visualization techniques. Finally, success stories help attract new collaboration partners and create new ideas for visualization techniques.

Contributions in this area include the following:

- *C9*: Create awareness in application domains of the importance of visualization.
- *C10*: Document a success story for dissemination, both inside and outside of the visualization community.
- *C11*: Describe innovation processes and commercial impact.
- *C12*: Contribute to a widespread understanding of visualization science.

No Set Template

An application paper can thus serve many different purposes, so it is not possible to describe one template for how to document these different contributions. Quite the contrary, several panel participants argued that one size doesn't fit all and that there should be no fixed rules for writing such a paper. Shoehorning an application paper into rigid guidelines could prevent the paper from "singing the way it should."

Nevertheless, there was consensus that a paper should make contributions to the visualization community and not just to the application domain. That is, an application paper should be more than an "instruction manual." Although the paper should be targeted at an audience whose core competencies consist of designing new visualization and analysis methods and evaluating them, there are many ways in which an application paper can contribute to the state of the art in visualization research.

Assessment in the Review Process

This variance among application papers means they are particularly difficult to review and assess. They certainly need to be evaluated, but reviewers must consider the whole spectrum of possible contributions, some of which we have already mentioned. It would, however, be an interesting exercise to look at each of these contributions and see how they would best be described and put forward in author guidelines.

Weighing Technical and Other Contributions

Application papers contribute in many different areas, and reviewers should not expect an application paper to present a new technique or algorithm or focus too much on any single area. Instead, they should ask these overarching questions: Does this paper have an impact? What is this impact? Does it contribute new knowledge to the visualization community? Does it have common ground with the visualization community and core capabilities that can inform other areas?

Although this recommendation is already mentioned to some degree in the VIS Conference review guidelines, some reviewers still reject papers because they do not introduce a new method. The visualization community should continue to embrace the diversity of applications and aim for an inclusive review process. Review forms for conferences already contain questions that guide a reviewer in evaluating a paper's strengths and weaknesses—such as about novelty, reproducibility by graduate students, and appropriate evaluation. Perhaps the visualization community can agree on a set of questions—based on the list of 12 possible contributions of application papers we have listed here—that can guide reviewers similarly, ensuring that reviewers consider a spectrum of contributions. This questionnaire could be a separate list, although a separate document may be more easily overlooked and not as effective as questions included on the review form. Events such as "Visualization in Practice" can serve as a forum to facilitate this process.

Evaluating Developed Methods

The VIS panel also discussed another contentious aspect of the review process. Currently, the visualization community has an extremely strong focus on formal, often laboratory-based, user studies for the evaluation of new visualization techniques. However, the variability of application papers means that a user study is not always a possible or even an appropriate means for evaluation. One important "dimension of variability" that influ-

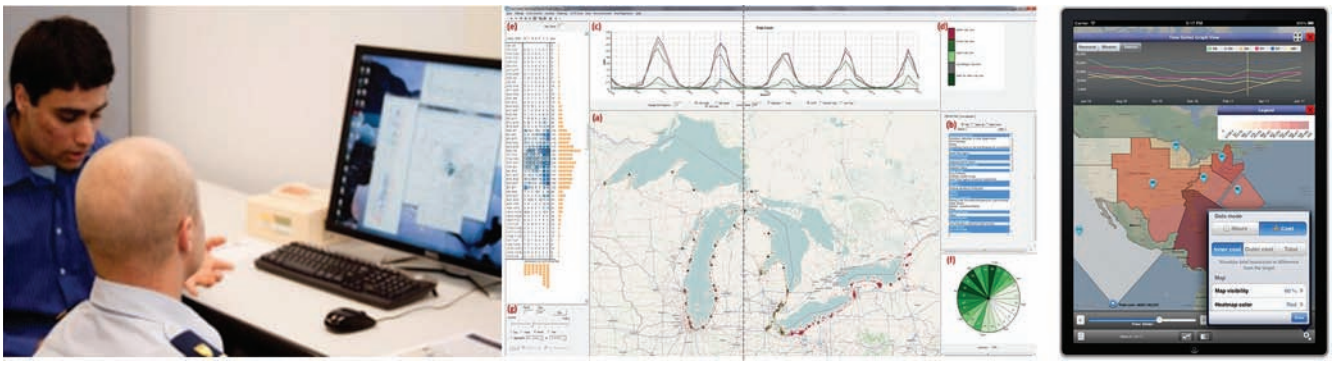

(a) (b)

Figure 3. The Coast Guard Search and Rescue (cgSARVA) tool. (a) The cgSARVA tool provides visual analytics capabilities for decision makers and analysts within the US Coast Guard, where the potential users vary from a small few to a hundred people. (b) The interactive Operations Performance and Assessment Report (iOPAR) tool is intended for use by the two Coast Guard vice admirals in charge of the Atlantic and Pacific area operations.

ences the utility of user studies is the size of the user community, which may range from a few specialists to every citizen. For larger user communities, user studies are an important and appropriate evaluation metric, but application papers describing a system for a few specialists are also valuable.

One example of this scenario from VAST brought up during the panel is the Coast Guard Search and Rescue (cgSARVA) tool (see Figure 3).[15] The cgSARVA tool provides visual analytics capabilities for decision makers and analysts within the US Coast Guard, where the potential users vary from a small few to a hundred people. In fact, the tool's interactive Operations Performance and Assessment Report (iOPAR) tool is intended for just two users: the Coast Guard's vice admirals in charge of the Atlantic and Pacific area operations. Analysts consider the tool useful for making important decisions, yet the small sample size and time constraints of the expert users make it difficult to impossible, and not always appropriate,[16] to perform typical user studies. A formal user study, while a laudable goal for empirical research, may thus be impractical in cases with small, busy, expert user communities.

Joseph McGrath has discussed how precision (finding an answer to a question), generalizability (finding an answer that will apply beyond the participants studied), and realism (finding out how some factors relate to the real world) are all equally valuable empirical goals.[17] However, no study method has yet been devised that addresses all three. When choosing the best empirical method for application papers, the most common goal will be realism. This will lead to empirical approaches, such as the following:

- observational field studies and more controlled field experiments, including both long- and short-term studies[17];
- diary studies, including autobiographical,[18] single participant,[19] and small group case studies;
- technological probes[20,21];
- focus groups;
- expert reviews[22]; and
- various forms of participatory design practices.

Appropriate evaluation of application-oriented papers must continue to be an evolving practice that involves considering previously successful methods that have been used in visualization,[23] making use of McGrath's discussion of empirical approaches and the tradeoffs between them,[17] and adapting methods from other empirical sciences.[24]

Another factor limiting the effectiveness of user studies is that is often impossible to reduce analysis to simple tasks or that tasks are not known a priori.[25,26] For example, in some cases visualization tool development is an iterative process—application professionals learn what types of visual analysis are possible while working with visualization researchers. Therefore, the problem statement must be iteratively refined during development.

The takeaway message from this is that reviewers should not expect a user study with 20 to 40 participants for each paper. When evaluating a paper's contribution, reviewers should be more inclusive. An evaluation should ask the following: Has the method produced relevant results? Are experts using the method—that is, is there expert buy-in?

For many of these questions, reviewers must have some level of expertise in the application area or should consult those who do. The visualization community should actively solicit additional reviews from application domain experts who can determine more authoritatively whether the developed methods are useful in the domain. A reviewer pool could consist of visualization experts evaluating the soundness of the visualization methods and

domain experts evaluating the impact and insight gained in the domain area.

While writing this article, we came to the realization that contrasting applied and basic research imposes unnecessary limitations on our field. Rather than discussing tradeoffs and compromises between basic and applied work, we should see them as mutually reinforcing streams of research. Basic research may take place independently for some time, but in the end, effective visualization techniques will inevitably be incorporated into applications. Applied research may continue in some isolation until a fundamental research question, new assessment method, or groundbreaking implementation inspires visualization researchers to explore the implications for basic research. We conclude that by exposing the visualization community to real-world problem settings, application papers will improve the quality of fundamental visualization research as well as build a bridge to real-world applications and potential funding sources. Application papers tell our success stories, demonstrating that we are accomplishing our goal of enabling data understanding in science, engineering, and society. Visualization research needs these success stories to secure continued funding, attract interdisciplinary researchers, and reward researchers that build the infrastructure to ensure the use of newly developed visualization techniques.

A healthy dose of application contributions is key to building a lively, relevant, and vibrant visualization research community for decades to come.

The diversity of contributions and the spectrum of target audiences make the evaluation of application papers particularly difficult. There is no single mold or template for application papers. As a consequence, both authors and reviewers should be flexible and open minded when it comes to application papers. Authors should ask themselves, how can a fellow visualization researcher benefit from reading our paper? How does the paper contribute knowledge to the visualization community at large? Authors can help reviewers by clearly articulating relevant contributions.

In turn, reviewers should have an open mind on how a paper can benefit a visualization audience. Instead of following a checklist and expecting a paper to make contributions in familiar and narrowly defined areas, such as new techniques, or evaluating new methods using familiar methods such as formal user studies, they should ask themselves higher-level questions: Does this paper present something that is new and useful to the visualization community? Does the paper present plausible evidence that the new system is useful in the application domain?

Of course, the challenge for this approach is that these higher-level questions can lead to increased subjectivity in the review process. Therefore, the visualization community should work to develop common ground among reviewers and objective criteria for evaluating application papers. Developing specialized events, such as "Visualization in Practice," that capitalize on the domain knowledge of smaller, more integrated program committees and reviewer pools will help the visualization community to understand the unique needs and situations of use of our stakeholders. Colocating these events with our main conferences will provide opportunities for outreach to these communities, with the goal of building a more diverse and impactful field. These colocated events will inspire and guide application development, core visualization methods, and innovative scientific methods for studying visually enabled reasoning in analytics, decision making, and operational management. Similarly, the IEEE SciVis Contest and VAST Challenge could help guide the development of usable and constant assessment criteria. For example, the SciVis Contest already considers both scientific relevance and overall visualization approach when choosing winning entries. It also involves application scientists and professionals during assessments—an idea that may be beneficial to the main conference as well.

The VIS panel participants also agreed that the focus on acceptance rate as the measure of quality is counterproductive and it harms our community. Although this focus affects all papers—including technique papers—it seems to influence application paper reviews the most. Computing has a lower acceptance rate and lower impact factor than many other fields, in part because a strict focus on acceptance rate reduces the exposure of innovative early-stage work, resulting in slower progress for the field as a whole.

The visualization community can learn from the experiences of other, more application-focused computing conferences. For example, the KDD Conference has increased industry participation

successfully with its conference track of industry and government invited talks.

As a call to action, we encourage the visualization community to increase the presence, traction, and appreciation of science applications in several ways:

- invite application leaders to visualization conference program and review committees;
- engage application communities with more outreach to their leaders, events, and publications;
- include application keynotes, panels, presentations, and exhibits;
- feature application successes with media releases and awards for best cross-sector collaborations;
- encourage the governing bodies of journals and conferences to adapt review criteria that accommodate the list of 12 application contributions we provided here;
- discuss and expand our initial list of ways in which application papers can make contributions to visualization; and
- develop criteria for how application papers should be presented to make contributions clear and assessable.

We feel that a healthy dose of application contributions is key to building a lively, relevant, and vibrant visualization research community for decades to come. Consequently, we call upon the community to develop more inclusive guidelines for application papers that will invigorate the field with increased diversity of uses of visualization and visual analytics.

Acknowledgments

Gunther H. Weber was supported by the director of the Office of Science, Office of Advanced Scientific Computing Research, of the US DOE under contract DE-AC02-05CH11231 through the project "Towards Exascale: High Performance Visualization and Analytics Program," program manager Lucy Nowell. Sheelagh Carpendale and the InnoVis Research Team thank Canada's Natural Science and Engineering Research Council, Alberta Innovates Technology Futures, and SMART Technology for research support and the University of Calgary for encouraging cross-disciplinary collaborations. David Ebert thanks the US Department of Homeland Security Center of Excellence program for funding this work and supporting this application-driven research (award 2009-ST-061-CI0003). Brian Fisher thanks Boeing for its collaboration, research funding, and contributions to the Vancouver Institute for Visual Analytics (VIVA). VIVA is a joint research institute of Simon Fraser University, the University of British Columbia, and the British Columbia Institute of Technology that promotes the value of visual analytics; provides training in visual analytics tools, techniques, and theory; and connects university research with industry needs. Ben Shneiderman and the University of Maryland EventFlow team thank Oracle and Adobe for their research sponsorship, which also brought them in touch with knowledgeable practitioners who challenged them to deal with significant real problems. Anders Ynnerman thanks the Knut and Alice Wallenberg Foundation, Swedish eScience Research Center, and Swedish Research Council for generous support of the projects described in this article. Examples and images were provided by collaborators at Linköping University.

The authors also thank the audience that contributed to the panel discussion and whose views helped shape the opinions presented in this article. Special thanks to Meghan Haley for ensuring the panel video was accessible in time for writing this Viewpoint and to Bernd Hentschel for valuable discussions and comments on the narrative. Finally, we thank the anonymous reviewers for their comments that helped to improve this article.

References

1. B. Lorensen, "On the Death of Visualization," position paper, *NIH/NSF Proc. Fall 2004 Workshop Visualization Research Challenges*, 2004; erie.nlm.nih.gov/evc/meetings/vrc2004/position_papers/lorensen.pdf.
2. E. Sundén, A. Ynnerman, and T. Ropinski, "Image Plane Sweep Volume Illumination," *IEEE Trans. Visualization and Computer Graphics*, vol. 17, no. 12, 2011, pp. 2125–2134.
3. S. Lindholm et al., "Boundary Aware Reconstruction of Scalar Fields," *IEEE Trans. Visualization and Computer Graphics*, vol. 20, no. 12, 2014, pp. 2447–2455.
4. S. Lindholm et al., "Spatial Conditioning of Transfer Functions Using Local Material Distributions," *IEEE Trans. Visualization and Computer Graphics*, vol. 16, no. 6, 2010, pp. 1301–1310.
5. F. Hernell, P. Ljung, and A. Ynnerman, "Local Ambient Occlusion in Direct Volume Rendering," *IEEE Trans. Visualization and Computer Graphics*, vol. 16, no. 4, 2010, pp. 548–559.
6. D. Jönsson et al., "Historygrams: Enabling Interactive Global Illumination in Direct Volume Rendering Using Photon Mapping," *IEEE Trans. Visualization and Computer Graphics*, vol. 18, no. 12, 2012, pp. 2364–2371.
7. C. Lundström et al., "Multi-touch Table System for Medical Visualization: Application to Orthopedic

Surgery Planning," *IEEE Trans. Visualization and Computer Graphics*, vol. 17, no. 12, 2011, pp. 1775–1784.
8. P. Ljung et al., "Full Body Virtual Autopsies Using a State-of-the-Art Volume Rendering Pipeline," *IEEE Trans. Visualization and Computer Graphics*, vol. 12, no. 5, 2006, pp. 869–876.
9. P. Ljung et al., "Transfer Function Based Adaptive Decompresion for Volume Rendering of Large Medical Data Sets," *Proc. IEEE/ACM Symp. Volume Visualization*, 2004, pp. 25–32.
10. A. Ynnerman et al., "Interactive Visualization of 3D Scanned Mummies at Public Venues," *Comm. ACM*, vol. 59, no. 12, 2016, pp. 72–81.
11. V. Bush, "Science: The Endless Frontier," *Trans. Kansas Academy of Science*, vol. 48, no. 3, 1945, pp. 231–264.
12. B. Shneiderman, *The New ABCs of Research: Achieving Breakthrough Collaborations*, Oxford Univ. Press, 2016.
13. V. Narayanamurti and T. Odumosu, *Cycles of Invention and Discovery: Rethinking the Endless Frontier*, Harvard Univ. Press, 2016.
14. M. Monroe et al., "Temporal Event Sequence Simplification," *IEEE Trans. Visualization and Computer Graphics*, vol. 19, no. 2, 2013, pp. 2227–2236.
15. A. Malik et al., "A Visual Analytics Process for Maritime Response, Resource Allocation and Risk Assessment," *Information Visualization*, vol. 13, no. 2, 2014, pp. 93–110.
16. S. Greenberg and B. Buxton, "Usability Evaluation Considered Harmful (Some of the Time)," *Proc. ACM SIGCHI Conf. Human Factors in Computing Systems*, 2008, pp. 111–120.
17. J.E. McGrath, "Methodology Matters: Doing Research in the Behavioral and Social Sciences," *Readings in Human-Computer Interaction: Toward the Year 2000*, 2nd ed., Morgan Kaufman, 1995, pp. 152–169.
18. C. Neustaedter and P. Sengers, "Autobiographical Design in HCI Research: Designing and Learning through Use-It-Yourself," *Proc. ACM Designing Interactive Systems Conf.*, 2012, pp. 514–523.
19. U. Hinrichs, S. Forlini, and B. Moynihan, "Speculative Practices: Utilizing InfoVis to Explore Untapped Literary Collections," *IEEE Trans. Visualization and Computer Graphics*, vol. 22, no. 1, 2016, pp. 429–438.
20. A. Thudt et al., "Visual Mementos: Reflecting Memories with Personal Data," *IEEE Trans. Visualization and Computer Graphics*, vol. 22, no. 1, 2016, pp. 369–378.
21. B. Gaver, T. Dunne, and E. Pacenti, "Design: Cultural Probes," *Interactions*, vol. 6, no. 1, 1999, pp. 21–29.
22. M. Tory and T. Möller, "Evaluating Visualizations: Do Expert Reviews Work?" *IEEE Computer Graphics & Applications*, vol. 25, no. 5, 2005, pp. 8–11.
23. H. Lam et al., "Empirical Studies in Information Visualization: Seven Scenarios," *IEEE Trans. Visualization and Computer Graphics*, vol. 18, no. 9, 2012, pp. 1520–1536.
24. T. Hogan, U. Hinrichs, and E. Hornecker, "The Elicitation Interview Technique: Capturing People's Experiences of Data Representations," *IEEE Trans. Visualization and Computer Graphics*, vol. 22, no. 12, 2016, pp. 2579–2593.
25. C. North, "Toward Measuring Visualization Insight," *IEEE Computer Graphics & Applications*, vol. 26, no. 3, 2006, pp. 6–9.
26. R. Chang et al., "Defining Insight for Visual Analytics," *IEEE Computer Graphics & Applications*, vol. 29, no. 2, 2006, pp. 14–17.

Gunther H. Weber *is a staff scientist/engineer in the Computational Research Division at the Lawrence Berkeley National Laboratory and an adjunct associate professor in the Department of Computer Science at the University of California, Davis. Contact him at ghweber@lbl.gov.*

Sheelagh Carpendale *is a professor at the University of Calgary. Contact her at sheelagh@ucalgary.ca.*

David Ebert *is the Silicon Valley Professor of Electrical and Computer Engineering at Purdue University. Contact him at ebertd@purdue.edu.*

Brian Fisher *is a professor at Simon Fraser University. Contact him at brianfisher@ieee.org.*

Hans Hagen *is a professor at the University of Kaiserslautern. Contact him at hagen@informatik.uni-kl.de.*

Ben Shneiderman *is a distinguished university professor of computer science at the University of Maryland and a member of the US National Academy of Engineering. Contact him at ben@cs.umd.edu.*

Anders Ynnerman *is professor of scientific visualization at Linköping University and director of the Norrköping Visualization Center C. Contact him at anders.ynnerman@liu.se.*

Contact department editor Theresa-Marie Rhyne at theresamarierhyne@gmail.com.