Discussion of “What is Statistics”?

Deborah Nolan, Dept. of Statistics, UC Berkeley
Duncan Temple Lang, Dept. of Statistics, UC Davis

December 16, 2008

We want to congratulate Emory Brown and Rob Kass on an important, timely and compelling article. They challenge each of us to think very seriously about the future of statistics education and practice and our role in that evolution. Along with Brown & Kass, we too believe that it is essential for the health of the field to make significant changes and additions to the content and focus of statistics training at all levels in order to attract, retain, excite and inspire students to become statisticians. To us, it seems “obvious” that we should broaden our view of statistics education to incorporate, alongside existing mathematical content, the process of real-world data analysis, skills with computing and data technologies, and statistical experience in scientific contexts. Based on our experiences, we think statistical programs need to:

- focus on statistical experience, reasoning and applications *throughout* statistical training;
- recognize computing as an essential building block for statistical learning, creativity, exploration, and practice;
- design new courses and curricula to attract bright, motivated students to the field;
- change the culture of statistics training to engage students in active, participatory “effortful learning” in addition to critical study.

Our discussion continues by providing some details and our thoughts on these four important aspects of statistics education. We then describe some of the activities we have pursued in our research and our teaching on these topics, and suggest how these might provide possible practical solutions to some aspects of these significant challenges Brown and Kass indicate we face.

We have been thinking about and working on making changes in these directions for many years. We believe strongly that the field of statistics is at a crucial tipping point and bold measures of reform in statistics curricula are called for. The changes are necessary both in order to attract and prepare future statisticians, but also to keep pace with the rapidly changing "big science" fields. Our experiences over the last ten years have shaped our views on the subject. These experiences include:

- organizing a summer school that engages students in applied research projects of statisticians, with an aim to encourage undergraduates to apply to statistics graduate programs;
- designing and teaching new courses in statistical computing and data technologies;
- teaching faculty how to teach computing;
- revamping a graduate program to broaden the curriculum and the graduate student population; and
- exploring how to make research activities and results available through dynamic, reproducible, interactive documents.

Before elaborating on these four viewpoints, it behooves us to make explicit a few parameters in this discussion. When considering statistics training programs, there are several different levels and career goals to take into account: undergraduate preparation for the workforce, undergraduate preparation for graduate school, Masters preparation for the workforce, and Ph.D. preparation for academia and for careers outside of academia. With one exception related to an introductory course to attract freshmen into the major, our discussion is primarily focused on advanced undergraduate and graduate students, i.e. not the service-oriented introductory statistics class. Commonalities and differences can be found across these different levels of training. One important commonality is teaching data analysis. The collective perspective is that data analysis is taught in all statistics programs. However, the phrase “data analysis” has many connotations, and we believe that it is often the case that “data analysis” experience is simply illustrating a particular statistical method by applying it to a pedagogically chosen data set. We use the term quite differently to refer to formulating a statistically oriented approach to a scientific question, and this involves much more than just one or more statistical methods. Also, when we refer to “computing”, we do not mean simply programming or numerical algorithms, but rather the broader notions of computational concepts, ideas and skills for statistical inquiry and working with data. Both are core elements of statistical thinking.

Statistical experience For those learning statistics, the intuition and experience that are necessary for good statistical practice are the hardest things to learn (Wild and Pfannkuch, 1999) and to teach. They involve very different types of concepts, and a new *dimension* of both learning and thinking than are used in mathematical thinking. After years of studying mathematics and statistics from textbooks, statistics students have learned a toolbox of methods, and they have not necessarily learned how to frame a question in a meaningful way, e.g. balancing constraints, resources, and context. Students may know how to use one or more of these tools, but are not masters of the tools and often use them with trepidation. They need training and practice in mapping a scientific question into statistical approach and developing an understanding, experience and intuition about when and how to use statistical methodology in the scientific context. These are essential skills in statistical thinking that involve many more aspects than selecting and applying statistical methods to data. However, most courses focus explicitly on statistical methodology, either the theory or the “application”, and very few address the essential larger context. The result of our focus on techniques is to train the student as a confirmatory consultant rather than an engaged scientific collaborator. To add this important dimension to our programs, we advocate regularly teaching statistics, at all levels, from the vantage point of statistical concepts flowing from contextual problem solving with data. We know this is challenging, but that does not excuse us from avoiding it.

Computing Traditionally, education and research in statistics has relied primarily on mathematics. However, the enormous growth and the power of the computer in the past twenty years provides immense opportunities for the field for both the practice of statistics and also statistical education. Computing represents an alternative, complementary medium for students to understand and explore statistical concepts and methods. The ability to simulate and compute gives students and researchers a tangible laboratory for exploring statistical concepts to concretize mathematical abstraction, and it gives a forum for gaining insight and intuition about potential new methodologies. Through computing students actively engage in constructively framing instructions to do a particular task, e.g. designing experiments to explore or confirm their understanding of concepts. This is quite different from mathematical exposure to statistical concepts where the student is passively accepting the results of theorems or cautiously manipulating symbols to prove a concept which is known to be true. If students had as much background in computing as they do in algebra and calculus, we would be able to exploit this additional medium much more effectively.

In addition to leveraging the computer for pedagogical purposes, computing in its own right is an essential part of statistical training. Statisticians almost exclusively use the computer to access data, to filter, process and explore data, to iteratively model the data, and then report findings about the data. Each of these steps requires the computer, and in fact each requires very different computational skills. Computing is also the source of new research problems, such as stochastic algorithms, understanding computer networks, and software reliability. Further, it has changed the nature of other scientific problems by being the medium for acquiring and exchanging both data and statistical methods in areas such as computational biology, astrophysics, aeronautics, transportation engineering, and medicine. If one does not have computational skills, one simply cannot engage in the application and practice of statistics, regardless of one's knowledge of the concepts. In addition, one needs good foundations in the concepts of scientific and statistical computing and data technologies in order to be able to continue to adapt to the rapidly changing technologies. Since most statistics students go on to apply statistics rather than study it academically, computational skills are vital, but as with data analysis, it is a dimension that is omitted from many statistical curricula.

Attracting students We agree entirely with Brown & Kass that statistical thinking and inter-disciplinary interaction, or better, immersion is key for a statistics student to learn. Brown & Kass also say that we must present statistics as being deep, with serious content. Again, we wholeheartedly agree and also add that we must present it as vibrant and relevant in the modern world and for the future. The repeated focus throughout undergraduate and graduate courses on the same concepts at different levels of mathematical rigor presents the view that the important statistical ideas have all been developed. Indeed, many students, even graduates, do not encounter methods developed within the last decade or two within their courses. Further, the repetition of the classical material is not a compelling approach to attracting good students to the major. Similarly, in our experience at the graduate level, this approach does not attract students to advanced study nor does it prepare them for research. The traditional statistics curricula are based on the need to first present an intellectual infrastructure in order to understand the statistical method. Instead, statisticians need to lead with real and interesting scientific questions and show how statistics "saves the day". Early and continued exposure to statistics in this form will, we believe, excite and interest students. As a result, they will be eager to learn about the statistical theory and take the more traditional classes we offer on the fundamentals.

Changing the culture As Brown & Kass note, the culture of statistics is more one of confirming other people's work, and often criticizing it. The culture of changing the world, attacking the very hard problems and "dreaming big" is more with physics, computer science and engineering and seems to be quite removed from our field. Perhaps this "caution" is the nature of statistics and a good thing. However, being cautious and circumspect is quite different from a "can't do" attitude. As Brown & Kass note, we must instill in our students the self-confidence to immerse themselves in the subject matter discipline and work along-side the content experts. In our view, if statistics students are equipped with the unique skills of computing and experience with data, they will then acquire this confidence and be welcomed into scientific teams because they have something unique to contribute.

Challenges, Experiences, and Solutions

The role of introductory courses Over the last decade, many educators have focused attention on improving the introductory statistics courses. These courses service thousands of students who take only one statistics course, typically to fulfill some general education requirement of the university or their degree program. However, rather than providing basic statistical literacy to the masses, the introductory course can be viewed as a recruitment opportunity. We believe the field and the students would be significantly better off by showing the challenges and applicability of statistics to important policy and scientific decision making in many contexts, and teaching students how to think statistically and creatively in these contexts. How can we present the role of statistics in addressing "big science" problems in introductory courses? One possibility is to develop an "honors course" for a small group of students that is creative and bold in the sorts of research-like experiences it provides.

Our experiences with developing and running a summer program in statistics with Mark Hansen (Hansen et al., 2006) is a source of ideas for how such a course might work. In the summer program undergraduates, with limited backgrounds in statistics and computing, are exposed to important, topical scientific research problems presented by statisticians working on a scientific team. The program was held at UCLA in 2005 and 2006, and was funded primarily by the Institute for Pure and Applied Mathematics at UCLA. Recently the NSF awarded a grant to continue this program for four summers, beginning in 2009 at UC Berkeley and then moving to the National Center for Atmospheric Research (NCAR), Columbia University, and UCLA.

The core of the program consists of three data analysis projects spread over six days. Each project is lead by a research statistician who organizes two days of activities around an applied project. The researcher presents the scientific problem and explains its importance, provides data, and prepares short talks and computer investigations where the students are introduced to the material in stages. At each stage, the instructor guides a discussion in which the students come up with with different approaches for the sub-problem and then they work in groups to follow up on one or more of these approaches, and return to discuss their findings. Students use R (R Development Core Team, 2006) to explore and visualize the data. The aim is to keep the interaction fluid and make it easy to move from individual and small group activities to a short presentation on a topic by the speaker to informal class discussion and group presentations. Overall, we found that the students were captivated and engaged by their interactions with the researcher. With the help of numerous instructors and TAs, they quickly mastered the computing tools, and were excited about using them to uncover the basic structure of the data, get to the statistical problems the data present, and gain a sense of how statisticians approach large, complex problems. Also, the program has been successful in attracting a broad spectrum of students. For example, in 2005 & 2006, half the participants were female (24/49) and one fifth (11/49) were from under-represented minority groups.

Teaching computing While statistics students must learn practical details of computing such as programming language syntax and the names of useful functions, we must strive to teach higher level concepts of computational thinking that enable students to approach computational tasks intelligently. This includes the ability to discuss and reason about computational problems precisely and clearly. Further, as computing and data technologies continue to evolve rapidly and as we enter the era of mainstream parallel and distributed computing for scientific computing, it is essential that statistics students are in a position to continue to learn new aspects of computation based on a good foundation rather than a thin memorization of specifics and ad hoc tricks. Statistics programs must prepare the student for their future, which undoubtedly involves computing.

Since 2004, we have been developing and teaching an upper-division course in our respective departments. The two courses are similar and have been developed in close collaboration. The over-arching topics are data technologies and statistical and scientific programming. Although the course has no statistics prerequisites, students work with topical and relatively large data sets where they perform exploratory data analyses using advanced data technologies and "modern", computationally intensive, statistical methods that they typically do not see in other classes. These methods (e.g. CART, k^{th} nearest neighbor methods, Naieve Bayes classifiers, hierarchical clustering, and spline smoothers) are intuitive and relatively easy to describe and give students a sense of the power of modern statistics. We have observed amazing transformations in our classes as students, who initially were unsure of their abilities in computing or otherwise reluctant to work with the computer, gain the confidence and skills to tackle a wide variety of data problems. It is empowering because they are involved, active participants. The students find it interesting because the data are available for compelling topical questions, and many find it refreshingly different from more traditional classes. We have also found that the course has attracted many students from other majors and graduate students from other disciplines. For example, at Berkeley the course is now taught every semester with enrollments of about 75.

Faculty experience For many faculty there is a large divide between their computational training and what today's students are expected to do. Some faculty have kept up and learned how to "compute", but many haven't and many have done so in an ad hoc manner, which conveys to the students that computing is not important. This is very unfortunate as it means that new students do not get the opportunity to learn it either. So, they are in the same position as previous generations. That is, students are left to learn computing by themselves, and the results are typically quite poor with students having significant misconceptions, limited abilities and confidence. How do we break out of this cycle and provide the opportunity for students to learn this material? One approach we have pursued is to develop workshops specifically to teach faculty, and foster Internet discussion groups for instructors.

In addition to developing new computing courses, we have also worked to develop expertise among faculty and graduate students at other institutions so that they can teach this important material to the current and next generation of students. To do this, we (along with M. Hansen (UCLA) and R. Peng (Johns Hopkins)) are organizing workshops to assist faculty in acquiring the knowledge, skills, and teaching practices in these new areas. The NSF provided us with funds for a series of three workshops. The first, held in 2007, brought together computing specialists and industry consultants (people who have employees in statistical roles) to advise us in preparing material for course and curriculum development plans. Two additional workshops (one held in 2008 and another in summer 2009) focus on providing the background and skills for instructors to teach statistical computing courses, and examples of how to include modern data analyses projects in their courses. The materials produced for these workshops and resources from our classes are available on the Web (Nolan et al., 2007). We have also created electronic mailing lists, discussion boards, and a wiki for continued discussion and assistance and to build a community of educators interested in incorporating computing into the statistics curriculum and sharing course materials.

Course Materials Finding interesting and topical scientific problems with accompanying data in a form accessible to instructors who want to teach in this experiential way can be difficult. The Internet provides a great resource for data, but often falls short in supplying analysis and context. Articles that present applications are plentiful in research journals, but the analysis is typically presented as a completed work and the pedagogically important thought-process that led to the conclusions and approaches are omitted. Where will educators find a wealth of materials that are suited to this approach to teaching statistics? Vehicles for getting the experience with data from the working statistician to the student are needed.

A project we are experimenting with (Nolan and Temple Lang, 2007) offers a novel approach for providing students with statistical experience. The idea is to enable researchers to document their entire computations and analysis process so that they can be reproduced in entirety for both themselves and their peers (Gentleman and Temple Lang, 2007), e.g. reviewers, editors, bosses. Researchers would work in an environment that captures their writings, computations and thought process in an electronic notebook. In essence, this "lab notebook" would be a database of all the activities within the data analysis or simulation study, and it can be projected into different "views" (e.g. code, the final paper, various "dead ends") to make the information it contains available for different audiences. An important consequence of this approach is that these rich documents provide a flow of materials from statistics researchers involved in scientific applications to the education community. These documents would provide resources to instructors to assist them in teaching in new ways because they would open up the thought process and experience behind a data analysis both to the instructor and the students. Further, these documents can be displayed with interactive controls allowing the reader to explore different analysis choices, e.g. changing the values of nuisance parameters, discarding outliers. This technological approach would support a model for passive cooperation between statisticians active in research and consulting and the community of statistics educators. Instructors would then have libraries of real case studies that include data analysis projects and current research methodologies that show how statisticians think and work.

Adjustments Making such fundamental changes to the training of statisticians will not follow a prescribed, straight path. At most institutions, the training process has been running fairly smoothly for twenty-plus years or more. We cannot anticipate all that will happen to our programs as a result of such modifications. Even the question of where to begin is not easily answered. Changes of this magnitude will have repercussions, and it is important to make adjustments, continue on, and not turn back to the old system that supposedly "worked". How do we begin? How do we ensure that students on different pathways don't slip through the cracks? It takes a concerted effort, along with perseverance, to make significant changes to a program.

In the past several years at Berkeley, the Statistics Department has undertaken a major reform of its Ph.D. program. A task force of faculty and graduate students reviewed the program, paying particular attention to the first two years and whether students were adequately prepared for research. The goals of broadening our graduate students' education and broadening our graduate student population acted as the impetus for this reform. The task force recommended that the program (1) broaden the traditional first year course requirements of two year-long courses in two of the three areas of probability, theoretical statistics, and applied statistics to include other courses, such as the new course "Probability for Applications" as well as courses from other disciplines; (2) require students to embark on a short-term research project, internship, or other research activity during the first summer of the program. In order to accomplish these recommendations, two additional, key changes needed to be made: (3) replace the preliminary exams, which were held in the summer between the first and second years in the Ph.D. Program, with the requirement of satisfactory progress in the first one to two years of graduate course work; (4) develop individual course plans for incoming students with the graduate advisor and a faculty mentor. The transition to this new program was not without problems and does involve effort and resources. Naturally, the effect of such significant changes to the program were not all anticipated from the start, and the program continues to evolve. Currently the general sentiment is that the program encourages more cross-disciplinary research and that these changes are attractive to graduate students.

Conclusion

Brown & Kass' discussion of statistical thinking is very important. It is what most of us recognize as the essence of statistical contributions. Yet too often, the educational focus remains on techniques and mathematical presentation of concepts because it is convenient and familiar. Perhaps the problem is that most academic statisticians have not had the experience that Brown & Kass speak of and the "anachronistic conception" is being passed on through the generations. At a time of great change for science and statistics, statistics education is not evolving at a rapid enough rate. Educators are mostly doing the same things over and over again with minor extensions and there are few forces to cause us to change in response to general changes in science. This is not any one individual's fault, and there are many truly vibrant and novel statistics educators in academia, but as Brown & Kass mention, this status quo is the result in the aggregate and has us concerned and frustrated. Can statisticians take on the challenge to find bold, new ways to teach statistical thinking and practice? From where will the impetus come? Senior statisticians can step up to this challenge and create a community that supports this change, including encouraging and enabling more junior statisticians who are in the midst of this sea change to take important roles in the process.

In summary, we agree wholeheartedly with most of the ideas that Brown & Kass espouse and we are grateful that eminent statisticians have taken the time to write this paper that challenges our field. Unfortunately, these types of papers often get tacit agreement, but little or no action. Again, many individuals will be enthusiastic about the opportunity for change, but in the aggregate change will be difficult. This is especially true if university programs must change, and especially at a time when budgets are being squeezed. However, the topic is clearly important, and we feel vital for our field. We must find a way to effect change. Perhaps this should come from an organization such as the ASA. We must focus on changing the "anachronistic conception of statistics" of Ph.D. students and recent graduates. We must encourage senior statisticians to seriously challenge their own perspectives and encourage and facilitate junior faculty to design new statistical programs that emphasize statistical thinking and reasoning. We should pool teaching resources, perhaps hold workshops to foster new ways of teaching, and develop case studies for teaching. We might even train graduate students nationally to rapidly teach important topics, such as computing. Inter-disciplinary science, computing, and the digital world present a change point for the field of statistics and require us to think about what a modern statistics curriculum would look like if we had both the freedom to change and resources to implement it. For too long, the field of statistics has acted more passively to such change points and responded by merely adding topics to classes and not seeking, considering, and embracing new paradigms.

Acknowledgments The authors were supported in part by NSF DUE 0618865 and DMS 0636667.

References

Robert Gentleman and Duncan Temple Lang. Statistical Analyses and Reproducible Research. *Journal of Computational and Graphical Statistics*, 16(1): 1–23, 2007.

Mark Hansen, Deborah Nolan, and Duncan Temple Lang. Undergraduate summer statistics program, 2006. `http://summer.stat.ucla.edu`.

Deborah Nolan and Duncan Temple Lang. Dynamic, interactive documents for teaching statistical practice. *International Statistical Review*, 75(3):295–321, 2007.

Deborah Nolan, Duncan Temple Lang, and Mark Hansen. Wiki: Computing in statistics: Model courses and curricula, 2007. `http://www.stat.berkeley.edu/twiki/Workshop/CompCurric`.

R Development Core Team. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria, 2006. URL `http://www.R-project.org`. ISBN 3-900051-07-0.

C. Wild and M. Pfannkuch. Statistical thinking in empirical enquiry (with discussion). *International Statistical Review*, 67(3):223–265, 1999.