TEACHING STUDENTS AND STAFF CONSULTANCY SKILLS

Edward D. Rothman
University of Michigan, United States
erothman@umich.edu

Consultants are well aware of the skills we must impart. These include: an appreciation for the role of the work one does in the context of the research effort; an understanding of what we can and cannot do as statisticians; a general appreciation for modeling and visualization; and an understanding of the nuclear elements of modeling tools. All of these skills contribute to one's ability to communicate ideas to a client and, ultimately, to the reader of any findings. Changes in the toolkit for statisticians require that we continually learn. In this talk I describe a process that creates a continual improvement cycle for the mastery of these skills. By mastery, I mean that the student has demonstrated understanding through the application of a principle and taught what they have done to others, including non-professionals. The principles, based on modeling and visualization efforts, give rise to the nuclear elements that are the simple and communicable features of the tool.

INTRODUCTION

The skills required of a consultant are well known by many of us who have engaged in the activity. The first and foremost of these skills is communication. Included under this broad heading is both written and oral expression as well as the creation of an open dialogue with the client. The client must be sufficiently comfortable to reveal the theory examined, the design proposed, and some of the limitations of the study. This revelation and the client's understanding of statistics create unevenness between the consultant and client. Sensitivity to this difference is critical to the dialogue.

Establishment of a common language between a statistician and a client can be accomplished in many ways. The student who is new to consulting may simply learn the client language through course work and experience. This can be complimented with use of analogy described below. The idea is to use a language in one scientific area to accommodate the language elsewhere.

An ability to communicate without a deep understanding of what needs to be accomplished will leave a very limited impression on the client. Knowledge of statistical reasoning, modelling, limitations of inference and the ability to convey this to your client is the second skill set we will consider here. The approach I recommend includes an appreciation for "nuggets." By nuggets we mean displays---these can be formulae---or other summaries of the essential features of a modeling device. We illustrate these in a variety of contexts, especially multiple linear regression.

The product of our effort should reveal what we find from a variety of perspectives. This effort we call modeling, and it includes but is not limited to the statistical models that provide a simplified but useful way of understanding what was found. More generally what we label here as visualization will include statistical models, graphical displays and also alternative stories. To be sure any visualization captures only aspects of a situation. But it can provide statistical evidence supportive of what we see and in a format that the client can see too.

Mastery of any subject has a meaning that goes to the heart of achieving a purpose. Most systems of evaluation assess a student's understanding of the material, examples and theory, provided in the course. Our purpose includes the preparation for work as a consultant.

Statistics is essentially a practical subject, and one of the main advantages of using projects is that, for many students it is the first time that they are exposed to real data. The students then have to decide what data to collect, how to attain this data and what to do with this data. A project puts statistics into context and allows students to connect various pieces of knowledge together in a practical way. National competitions have been implemented in many countries and funded by various sources. Usually for each country there is some form of prize awarded as a reward for the best piece of work submitted for the competition. The competitions allow students from across each country to compete with each other for a national prize. In some cases the competitions have been international and, hence, students work is from many educational establishments from around the globe. The prize provides an incentive for the students to participate. There are many similarities on how countries run the competitions and the main educational aims are that the students see and use statistics in a practical environment.

Any project involves a student, or group of students, working under the guidance and direction of the teacher. A project involves planning and developing a schedule of study with outcomes to be achieved over a period of time. Implicit in the project is the deeper understanding of the subject and the bringing together of many statistical techniques. The students are not required to only produce a single outcome but are required to link multiple, often distantly related, prior learning tasks. The experience gained by the students and the teachers, involved with the projects, has been of enormous benefit to statistical education in those countries that have implemented them into their curriculum.

Though the ideas presented here have worked in specific context described, the ability to generalize has not been investigated. The nature of this work to-be-done is described in the discussion section.

COMMUNICATION

The use of analogy to create a language is what I teach, but when our students work in a specific industry they are expected to master that language. Here I illustrate the process through two examples.

A first order law such as Newton's $F=ma$, occurs virtually everywhere. Ohm's Law replaces the force, the push, with an EMF or Voltage and the acceleration, the reaction, with a current. The proportionality constant, the mass, m is then replaced by the resistance, r. Blood pressure is a force that creates cardiac output and is yet another version of the same law.

The law of mass action described in introductory chemistry courses specifies that the rate of reaction involving two substances A and B is proportional to the product of the two concentrations, [A] times [B]. Students readily see that the rate of spread of an epidemic can be modelled with "A" representing infected and B susceptible. And the rate of spread of a rumor may also be modelled with A, those who have heard the rumor and B those who may yet.

There is no substitute for an understanding of the client models and language, but use of analogy provides a tool we can use to begin communication with our client.

MASTERY

The learning process we most commonly use focuses, in my eyes, on the past. It involves a marble exchange with the student. The instructor provides a set of marbles—principles, facts, information—then examines how much of these have been retained. The latter step often involves a return of marbles through papers or examinations. My purpose is to have students apply these principles to problems and situations that are new. The process of teaching this requires a deep understanding of the principles to be sure, but also a demonstration of a capability to apply what was learned to entirely new situations. This process I call mastery.

Mastery requires three elements: an understanding of the principle, a demonstration of the understanding through an application to an entirely new situation, and finally a demonstration of an ability to teach the principle to an educated layperson.

The first step may be viewed as what we have done in the past. Evaluation of this step should be by someone who has mastered the principle and cannot be accomplished through multiple choices. When has a physician asked you for a differential diagnosis? We want the "experiment" to look as much like the "real-world" situation as possible. In the consulting class, the client should generally be real. Hypothetical examples may provide a satisfactory model of what the consulting experience will be, but use of real clients would be preferable.

The application should be new. Each new application, as those of us who have consulted even for a short time can attest, presents new challenges. Our aim, to prepare the student to be useful, requires exposure to these new situations. In a sense, these applications take learning from the particular to the general.

The teaching element of mastery reflects two aspects of learning we have all encountered. The first is that deep-seated feeling that we understand a principle better after we have taught it than when we were students learning it ourselves. The second is that we must communicate our findings to the client. This activity is akin to teaching. Indeed, I require applicants for consulting positions with CSCAR to present to our staff. The topic is not seen as an issue, but the presentation and the management of questions are important in our assessment of the applicant.

As will be evident in the course described next, mastery plays a critical role in the design and assessment of projects.

THE COURSE

The consulting course includes all students who want to complete the requirements of the applied masters program as well as several PhD candidates in statistics and related disciplines. These students have had at least two semesters of core material in mathematical and applied statistics. They have been exposed to basic linear models as included in Faraway (2005) and applied multivariate analysis. The theory class would include likelihood theory, Bayesian ideas, the exponential family of models, and some large sample theory. Approximately 20 students are enrolled in the course during the fall semester.

I believe students should not take the course too early in the program. They must be capable of implementation of standard linear models and diagnostic checks. In addition, they should be able to read and use models that include mixed models and models that allow dependency. One option is to expand the course to allow students design and analysis opportunities.

There are 42 one-hour classes in a typical semester here. Each student must engage in three projects on teams with two other students. The rationale is that each student should have an opportunity---and the responsibility---to manage a team. I have found that when group assignments are used, some students may seek the path of least effort. This is not a lesson I want to convey.

The teams are determined in part by student interest, but an attempt is made to mix the partners. I have had some of the same students on two teams and this will work but may be less effective than mixing. As you may imagine, we have a responsibility to educate the student and also to help the client. The ultimate burden is on the instructor to ensure the project moves forward.

Approximately one-half of the lecture time is devoted to the formulation of the problem and the final set of classes to presentation. The project load—three per student—and the material make for a hectic schedule. Much of the substantive work is conducted out of class and visits to office hours are quite common. In addition to office hours, where the entire team will meet with the instructor, the students will also meet to plan the work and prepare the presentation. Students are also expected to complete homework assignments based on the material described briefly in class. They may also meet with the client---with the instructor's agreement--- to ensure that client needs are met.

The lecture allocation is as follows: approximately 20 minutes of the first 20 classes introduces the clients and their problems to the class. The remainder of these initial 20 classes detail key concepts—usually connected to the problem—but more broadly relevant. The approach---limited by the 50-minute class length---describes nuggets. The last twenty classes allow 20 minutes for student presentations and feedback. The instructor must be careful to manage the time. I advise clients, in advance, that they will have just fifteen minutes to present the purpose and provide a modest level of background material. This is not unlike how I allocate the time I use for a one-hour appointment at CSCAR. The tendency to continue to speak exhaustively on a topic you find exciting is one that most clients have.

After the student team has had an opportunity to work on the problem, with appropriate levels of input from the instructor, they then prepare a presentation for the class. The student presentations always include a Power Point summary of their work. This is not the final report. Before the end of the term, the students revise this presentation for the client. The revised presentation and a text document constitute their final product.

Evaluation does not include this initial presentation. The goal is to make clear the framing of the presentation, appropriateness of the methods, and the clarity of the work. Other students in the class and the instructor provide feedback to the team making the presentation.

Each student on the team will speak for approximately seven minutes and then they may respond to any feedback.

The feedback session lasts approximately 30 minutes and includes comments on presentation and content. One goal here is to teach the student to avoid the hammer-nail syndrome; when you hold a hammer every problem looks like a nail. There are many approaches to each problem, and we want the student to see alternative approaches. Not all questions raised at this session are answered during this time period. Alignment with the client's purpose is key. The students will all have a statement of purpose among the first few slides. They will include a table of contents, make extensive use of graphs, describe detailed calculations and background materials in appendices, and have a clear statement of conclusions. The students may also recommend follow-up questions and issues arising in the report.

PROJECTS

In a typical semester, approximately 20 projects are required. As the Director of the consulting service at the University of Michigan, I have access to many clients who want more than consultation. Typically, at CSCAR we offer free advice on how to approach a problem, but the client must carry through the work. Some clients have had little statistical training and want us to do the work for them. This gives them an incentive to present to our class, because we not only provide them with ideas on how they might look at the questions they raised but also we will do much of the analysis and presentation of findings. We will not do data entry, nor do we guarantee a favorable result.

In addition to the projects that arise from CSCAR clients, I have access to projects from external clients—from my own consulting practice---as well as projects from previous semesters. The later category allows me to balance the selection of topics. For instance, one project should involve the design/analysis of sample survey, one should involve some qualitative data, and another should have some statistical computing aspect such as use of the EM algorithm. These topics do not always present in a timely fashion. The non-university consulting allows me to include a discrimination case or some other legal context. This allows me to make the class material broad and also more interesting.

To provide a view of some of past projects, the following list might help:

1. Performance Gap: Investigate why some student scores are lower on standardized tests (Principal components, regression)
2. Linguistics: The frequency of rarely used words in Mandarin, Korean and Thai manuscripts
3. Comparison between two measures of blood flow (Time series)
4. Survival analysis of left ventricular assist devices
5. Risk of operative procedure (logistic regression model and comparison with a subjective measure)
6. Seeking intra-operative indicators of death and complications
7. Are certain vehicle-passenger combinations that are less likely to give rise to death? (FARS dataset combined with other information on driver habits.)
8. Estimation of the number of children who will enter K-8th grade in a new housing development (EM algorithm)
9. Are people with less medical insurance coverage also less likely to use the prescribed drugs?
10. Do different levels of management use different types of communication skills?
11. Did the police department, City of Detroit, undermine the civil rights of individuals arrested for homicide? (The appearance of a higher arrest to number of homicides per year in Detroit was the plaintiff's data)
12. Can you estimate the change in the gross domestic product using a sample of firms listed on the stock exchanges? (Sample survey)
13. Sample size needed for a certain level test (various situations that allow students to examine issues of sample size, power, confidence interval levels and other related issues)

NUGGETS

Nuggets include ideas on question type and order, the measurement system, alternative models of association, dependence-importance and approaches, enumerative studies, linear models, dimension reduction, observational studies both retrospective (case-control) and prospective (cohort), matching and propensity scores, latent variables, hierarchical linear models, and generalized linear models.

To illustrate how we might introduce each of these items in just a few minutes consider a multiple regression problem with (two) standardized predictors. The generalization from two to many is straightforward and two is sufficient to produce the four points we want. The standard error of the estimate of a regression coefficient can be assigned to the students (in this special case) as a homework assignment. The variance of the estimator will be seen as a function of four factors:

1. The sample size appears in the denominator and increasing the size will clearly reduce the variance in our estimator.
2. The variance in the response (dependent variable), represents the variability of the response for a fixed value of all of the independent variables, appears in the numerator so we want this small. Omission of an important predictor will make this number large. We can ask the client whether other potential predictors have been left out of the model.
3. The variance in the predictor appears in the denominator. The greater this variance, the smaller the standard error of the estimate. We can ask the client if in the study we have sufficient variability in our potential predictors.
4. The variance inflation factor—the inverse of (1- R^2), with R^2 the squared correlation between the predictor at hand and all others appears in the denominator. Specifically, we want to know whether we need to look at alternative modeling tools such as principal components or ridge regression estimates.

This particular nugget may take more than a 30 minute period but it describes many of the key issues we face in explaining to a client why a variable they believed important may not appear—in light of the data---significant. A similar discussion of suppression, though not quite as frequently useful, can also be helpful.

Another nugget involves sample survey design. The Horowitz-Thompson Estimator provides a way to describe how—with knowledge of the distribution of the feature of interest---we can have a small standard error in an estimate with only a modest sample size. The variance of the H-T estimator can be expressed as a sum of products of two types of terms: $(Y_i/p\mathrm{i}(i) - Y_j/p_i(j))^2$ and the term—$(p_i(ij)$-$p_i(i)*p_i(j))$ where Y_i is the value of the feature of interest in unit i, and $p_i(i)$ is the probability that the ith unit is selected and $p_i(ij)$ is the probability that both the ith and jth units are included in the sample.

From this expression we can readily see that if we had access to information proportional to the unknown values of the Y's sought, we could obtain an estimator with quite small variance. For instance, if family income is close to proportional to taxes paid on a house we could sample houses in proportion to the taxes paid. Further, if we can stratify the population so that there is plenty of variation between strata and hopefully little within, then when the first factor is large we can chose to make the second factor zero. Independent samples in the two strata do just what we want. Finally, if we arrange clusters so that individuals in the same cluster are quite different—the first factor is large---we can make the second factor negative. Obtaining a census of a particular cluster does just this. Use of all NBA basketball teams---and the feature of interest the height of the players---provides us with examples. The players' positions create strata. The teams provide examples of clusters.

With a homework assignment to compute the variance of the H-T estimator in the format you give them, I can cover this material in just three 30 minute segments.

ASSESSMENT

Evaluation of the final projects includes a self-assessment. Each student involved in a project describes their work and a rates the contribution of the other members of the team. For instance, a five-point scale with one as the lowest score and five the high score would be awarded to each of the other two members. The students must also prepare a journal. The journal includes the following items:

1. An organized list of "nuggets" they have acquired from me and through the projects they worked on.
2. A glossary of terms with clear and concise definitions
3. Copies of the final report for each of their projects. This is usually a word document. The Power Point presentation edited to reflect the comments received.
4. The datasets from each of the three projects and metadata. The meta-data reflect the operational definitions of the fields and details reflecting the collection of these data. Unusual data are identified and described.
5. Solutions to homework assignments

The journal provides the student with an organized portfolio of their work. It may even be useful when they seek employment. The primary purpose is to collect together what they have learned in a format they can readily access.

Each student is also evaluated through an oral examination. The questions relate directly to the projects. The purpose of this portion of the evaluation is to ensure that the student has mastered all aspects of the project—not just the subset they specifically contributed.

DISCUSSION

The prototype of a course detailing the teaching and mastery of aspects of the consulting process has been used with success here at the University of Michigan with one instructor. Student evaluations attest to their appreciation of the program. And CSCAR has hired two of the graduates of the program---a small number to be sure and certainly not a representative sample---but at least the results seem to be excellent. However, evaluation of the process more broadly is necessary. The robustness of the method with other instructors, other University settings (e.g., students), and levels of preparation have not been fully investigated.

The principle that guides the use of real projects rather than those prepared projects seems critical to achieve our aim. It is clear, however, that a prepared project will be easier for the instructor to accumulate, can be chosen to illustrate a specific set of points, and can be transferred from one instructor to another. The disadvantages are also evident. The niceties of the prepared project provide little insight into the difficulties of dealing with a real project. Real clients have time limits. They communicate their purpose in terms that only become clear with probing questions. Data may have typographical errors and other oddities that need to be addressed. And most interesting is that the formulation of an approach may be linked to the format of the presentation of the question. The excitement generated by the enthusiasm of the client will also be lost.

Consider a client who has devoted several years to a new drug, say. The usefulness is not found to be significant. The client, in a real setting, may suggest that we re-examine the response—look at other features. They may insist that there are subsets of the data where efficacy was achieved. A data point or two may—when omitted—yield a significant result and the client may want us to look at such a subset. All these issues do arise in a consulting framework and simulation may not convey the passion of the situation.

Our consultant might find they need to manage multiple responses, subsets of data, optional stopping, and a general intra-ocular traumatic test (stop when something hits you between the eyes). The ethical challenge should not be addressed in the context of a job or the constraints of a mortgage payment but in the safe environment of the classroom.

REFERENCES

Faraway, J. (2005). *Linear Models with R*. New York: Chapman and Hall/CRC texts, Statistical Science Series; volume 63.