

Final Course Evaluation report for 36-617, Fall 2022

Brian Junker

12/06/2022

On page 3 below is a summary of your ratings for 36-617, and on the succeeding pages are your text comments and questions for me.

Once again, you were really thoughtful and helpful in your comments about the course! I will say a bit more about your comments below, but here I'll just say that I think you've hit on several important points about the course.

Here are a few comments and reactions to things that caught my eye in your ratings, comments and questions:

Ratings

- I'm pretty happy with your ratings of how well you understand things, and especially your interest in the course.
- Although many think the difficulty of the course was about right, roughly half the class think it was too hard ("a bit too hard", or "way too hard"). I think that's related in part to how many topics we cover, and in part to the fact that the course is fast-paced.
- With respect to pacing, a little more than half think the pace was about right, and almost half think it was a bit too fast. I see many of your written comments are about the course being a bit fast-paced too, and I think this is something to work on in the syllabus for future semesters.

One thing I didn't ask you to rate is whether you were especially stressed, either by this course or by the totality of things you had to do for the MSP program. Some of your written comments speak to that stress, and it is something we should work on for the program. Some stress is OK, and even healthy because it shows that you care, but too much stress makes it hard to think well and be happy.

Your Comments

- Once again, it's nice to see that the zoom recordings are useful, especially as the material has strayed farther away from things that you already sort of knew coming into the class. It's also good to see that office hours and Piazza are still really helpful.
- Generally people seemed to be happy to learn about nonparametric regression, multilevel models and Bayes. Several people would have liked more time on nonparametrics and Bayes.
- There was a mix of comments suggesting either that the pacing of the course was about right, or it went a bit fast. It is definitely a fast-paced course, and the impression of fast pace is exacerbated by the many topics we cover. There really isn't enough time to get in-depth about anything, just enough time to expose you to ideas and things that you will be able to look up in textbooks, on the web, etc., when you need them later in your professional career.
- There are a noticeable number of comments about being lost or somewhat lost in the Bayesian stuff. (On the other hand, most people seem to "get" the multi-level stuff pretty well, which is great!) In part that's related to the general goals of the course, to expose you to things and ideas that you can get more detail on later in your career when you need them. But it may also be that the idea of putting Bayes at the end of 617 just doesn't work very well, and that it would be better, in the future, to skip Bayes in 617 and offer Bayes as a separate semester (e.g. a spring "topics" course). That would give more time

to topics more directly related to regression in 617, and more time in the Bayes “topics” course to go into more detail about the reasoning behind Bayesian modeling, selection of prior distributions, how and why MCMC works etc.

- I tried to make most of the homeworks a more reasonable length in the second half of the semester, but hw09 was killer and I didn’t really give you enough time to do it. Several of your written comments speak to this. It was also challenging for you to have an assignment on Bayes due in the same week as your final IDMRAD paper.
- On a related note, while many people appreciated what they were learning from building an IDMRAD paper from the ground up, absolutely no one liked the rather compressed time schedule for the project. That’s really my fault: I should have started the whole process at least a week earlier, to give you more time to do and absorb the work in hw09, before trying to construct an IDMRAD paper around it. I just dropped the ball on how much time was left in the semester, after accounting for Thanksgiving break. This is definitely something to do better on, in the future!

Your Questions & Things You Learned

- People really seemed to appreciate learning about mixed-effects and multi-level models, which is great. I think people appreciated the exposure to Bayes as well but many people were hoping to learn more and/or in greater depth about Bayes. As I suggested above, we may need to think about splitting Bayes off into a separate course.
- Someone asked how much of what we learned transfers to other programming systems, like Python. Fitting `lm`, `lmer` and `glmer` models seems pretty far along in the Python world (see for example <https://www.slideshare.net/PyData/diamond-mixed-effects-models-in-python>) but I am less sure that a full set of diagnostic tools for regression and multilevel models exists in Python (you can always roll your own, but that’s not really a good solution for most people). See for example <https://towardsdatascience.com/going-from-r-to-python-linear-regression-diagnostic-plots-144d1c4aa5aand> and <https://stackoverflow.com/questions/46607831/python-linear-regression-diagnostic-plots-similar-to-r>. Stan can certainly be connected to python, but again I’m less sure about a Python analogue to `bayesplot`, for example.
- A couple people mentioned that they have a better sense of how to do a complete data analysis from raw data to final report, and/or that interpretability is the most important part of any analysis. I’m really glad to see that.


Remember, ordinary regression, generalized linear models and multilevel models (and even GAMs) have the great advantage of interpretability, so if you are working with someone who wants to understand and be able to communicate the mechanisms underlying the data you have, these models can be quite useful. If you are working on a problem where the main concern is good prediction, then methods with more of an ML flavor will generally do better, at the cost of interpretability. Sometimes your boss or client will want both, of course, and then you have to manage a tradeoff between predictive accuracy and interpretability. In any case, it is almost always worth trying a linear model or glm first, in part to see if a model with lots of interpretability will suffice to answer your boss’s or client’s question, and in part to serve as a “baseline” to try to beat with sexier methods that you may be asked to use.

- A couple people were confused about Bayes’ rule. I’m not quite sure what the confusion is, yet. Maybe it’s really confusion about how to specify a good likelihood and prior for any particular problem; people also asked about that, as well as more information about MCMC. For a general understanding of Bayes, I think Lynch Ch 3 is quite good. For a general understanding of MCMC I like Lynch Ch 4. (Both chapters are in the Week 13 and week 14 folders for you to read, and you can get the whole book for free from [springer link](https://www.springerlink.com) online.) Specifying good likelihoods and priors for complex data like our multilevel data is an art that takes practice.


I really enjoy teaching 617 and getting to work with the MSP students each year, and this is no exception. Thank you for a great semester!

–BJ


My Understanding is (1, 2, 3, 4, 5)




My Interest is (1, 2, 3, 4, 5)...



The Difficulty is (-2,-1, 0, 1, 2)...



The Pace is (-2,-1, 0, 1, 2)...



Comments

I feel pretty lost on the Bayes and Stan. I feel like the material became significantly harder since after the Fall break. Professor Junker is really helpful during office hours and on Piazza. Lorenzo was helpful during the office hour. I like having zoom recordings so that I can rewatch them when I don't understand something. The homeworks are really long and it takes the 3-4 days to complete the homework. The quizzes feel fair. The take-home midterm was really long, same with hw09. The timeline for completing the project felt really rushed and I feel like I didn't have enough time to work on the project. I like the times of the office hours. I feel much better now that the deadline for the homeworks are pushed to Wednesday so that we have two office hours before submitting the homework. Gradescope works well for submitting homeworks. Thank you for the wonderful semester!

- I think the materials are not very difficult, but I really hope they could present more specific example of code.
- The instructor and TA are kind and nice, I really appreciate you all!
- The in-class lectures, zoom recordings and the textbooks are clear and in detail, I think that is very great.
- The homework is a little hard for me, and it has a huge workload.
- The take-home midterm is a little hard for me, also I didn't have enough time to precisely finish it.
- The IDMRAD project is great, but I think it might be more useful for students who want to apply the PhD, maybe most students want to find a job after graduation.
- The office hours are really useful! Especially for the mathematical theory deduction, even if I majored in statistics in the past and I've done the mathematical theory deduction for whole four years of university, about some steps during the process of deduction, they still might cause some difficulties for me. So I think office hours really help.

Thank you for you all!

Overall, I think that the course material was slightly rushed. I enjoy the full coverage of the variety of topics, such as getting into MCMC using Stan, but the tradeoff was definitely the lack of time to practice such material. I thought the homeworks in the second half of semester were a good challenge and the right length, but I thought more practice to enforce those skills would've been nice. This is not a fault of the class but I thought it was worth mentioning. I thought the in class lectures were helpful, and having the zoom recordings were nice if I needed to go back aside from looking back at lecture notes. The textbook and textbook problems were fine. The quizzes were quicker ways of finding out if I understood the material, which was definitely helpful. Having piazza was nice to ask questions, and I really appreciate the quick response times. The due dates and grace periods were really nice to have, as there were many assignments and tasks going on during the semester.

Overall things have been run well. There is good documentation and explanations offered when needed. One thing that is a bit stressful is figuring out the whole STAN installation issue during our last week of the semester, but once I figure it out I'm sure it'll be fine:)

The IDMRAD project is good. It helps me learn how to organize my research output and is useful for writing reports in future work.

I think the difficulty of this course is hard for me. When talking about mcmc, I'm a bit lost there. The instructor and TA are great, we can contact them through Piazza and they reply to piazza very fast. The instructor is super helpful during office hours. I think the lecture sometimes is a bit dry when getting into the coding part. The take-home midterm format is great, it helps us better understand IDMRAD project in final.

it was all good except the timeline was a little bit off. I was able to pull it off though

(continued)

Comments

This is probably the most difficult course (for me) in this semester by far, but I think it has a fitting difficulty level. The content of this course is challenging and interesting enough to keep me awake in the afternoon right after lunch time, and I can't stress enough how much I enjoyed the lectures of this class. Professor's lecture style is also exactly what I wanted to be and the pace of the course feels much better in the second half of the semester.

The homework of this course is equally challenging, but I always found the office hours helpful whenever I am stuck on a homework problem.

I personally adore the take-home midterm, it's gave me a stress-free opportunity to test my understanding of this course, if anything the midterm could have been a little harder than the usual problem sets.

The IDMRAD project is another thing that I enjoyed in this course. Working through a real-life problem without a "correct" answer feels much more rewarding and engaging for me. Although the deadline is a bit tight due to the extra problem set, the project is flawless.

I have rarely used Piazza to ask questions, but when I do, the response is clear and it's a helpful platform since you can see everyone else's questions.

I never experienced any difficulties submitting things online.

Overall, this is an adequately challenging course and I am more interested in statistics than ever after taking it.

I think the take-home midterm was a bit coding intensive. I think it would be more useful if we focused more on the practical takeaways from our models and what these takeaways suggest to the problem's context.

At times I feel like I and other students are writing the code but only have just enough understanding of the concepts to vaguely defend what they are doing. So I think we should focus more on the foundational reasons behind whatever method we are learning before delving into the details. Because in the real world, I could be asked to explain what MCMC is generally but I probably wouldn't be asked to explain Gibbs sampling or other particular versions of MCMC.

I suppose my final recommendation is to really focus in and emphasize the fundamental reasons of 1. what we are doing, 2. why we are doing it

I think the class material and code provided in class is very useful and the office hour really helps a lot. The IDMRAD project and the peer review helps me a lot to improve my own draft and learn from other people.

I felt as if the course schedule at the end of the semester was a bit rushed. I would feel more confident in my paper if I was given more time to complete the statistical methods in homework 9. However, we were given more time to review our own rough drafts.

I think the office hours and the homework really help my understanding of the course materials. I hope we'll have more time to work on the final project.

I do appreciate the demos we have, I just wish we had more time to cover the topics at the end of the semester. Maybe the beginning few lecture could be condensed more so we have more time for the end of semester content? It would also have been nicer to start on the final project several weeks before it was due, as it seemed incredibly interesting to fully compare classical against popular.

I think Professor Junker does a great job expressing interest in students' learning. I think the office hours session are extremely helpful and collaborative, making this course even better. I also love how the homework solutions are routinely released after the assignment is due so that we can analyze our work with the solutions. I think this class is not trying to trick you, which I like. It seems as if as long as you put the work in a go the extra mile to do what is asked of you in the course, you can succeed. Professor Junker is also always flexible with due dates and truly wants us to limit stress as much as possible. Overall, great course. I learned a lot.

the powerpoint is a little bit make me confuse, more explanation of those code will be better

(continued)

Comments

I really like how the course is constructed. A few comments about improvement - I think this course should have more than one TA, perhaps it'll help the grading process much smoother (since every hw is quite long). Also, the deadline for hw9 (analysis before final IMRAD project) was way too tight, I think there should be improvement in scheduling (and I do NOT think it is a good idea to have 2 day grace period - it did not really feel like a grace period and frustrated me quite a bit).

I like all the material but I need more time to do the reading before the class, or I would feel kind of hard to follow the instructor.

The instructor is pretty good, I really like the way he taught, and his Office Hour is the place where new ideas and solutions are most stimulated, I love it!

Zoom recordings is useful.

Homeworks in the last half semester become more reasonable in terms of time spent.

The IDMRAD project is great, It would be better to bring forward or push back the draft dupe has caused me some problems around Thanksgiving, but most of them are my own.

The level of the material is a little bit difficult for me, but I almost understand, the piazza using is good for me. The zoom recording is useful to review the course content.

I really like the recorded videos that we can help to review classes after lecture.

I think the major issue I had overall was the pace of the course as it covered a lot of different concepts very quick, however the second part of the semester relatively better in terms of pace. Definitely, the week of thanksgiving and the weekend before was quite stressful due to the load of the assignments, but otherwise I believe that I learned a lot of very interesting and new concepts that I would be expected to use in real world. The assignments and quizzes did help a lot in understanding the concepts, even if the assignments were quite long. I would definitely say that the course would need an additional hour for office hours for the week as the homeworks are long and conceptually challenging and it might harder to accommodate every person's every question. The files for submission online were very easy and straightforward. Overall, I am very thankful for the course.

It was sometimes hard to process the material in class the first time hearing it usually. I often had to reengage with the material to understand it. I think this is normal, but it felt a little discouraging in class. Going to office hours helps with understanding the material and I'm glad I went so frequently. I was able to get on top of the difficult material this way, but I don't think there was always enough time to fully understand all of the homework. I also wish we had more time to do the final paper and that the assignments were either longer with more time to do them, or that there were less of them. It always took more hours to do the homeworks than I anticipated so I wouldn't suggest making them longer. The textbook wasn't always helpful for specific questions but it did sometimes help to look at it for general material understanding. A lot of times I forgot about the quizzes until the last minute because of the time of their release and due dates. I also really didn't like that the final paper rough draft was due over Thanksgiving break because I traveled and had to work on it in my hotel room with spotty wifi. I liked the idea of the peer reviews but I think having more time to create a rough draft/having longer between the solutions of HW9 and the rough draft due date would've been better because there was a lot of variability in the quality/amount people were able to write for it. I also really appreciate there being late due dates because of other life things that get in the way. Usually it was easy to submit online, but the longer the assignment pdf the harder it was to upload, especially because it felt like Gradescope and Canvas could sense fear. In general I think having more time to review the solutions of the previous assignment before the next assignment would be helpful. Overall I think the class improved and I enjoyed the material even though it was hard.

The lecture recording and piazza are really useful. I really like this course even though it is a little hard for me to understand all of the contents. Homeworks helps me to get a better understanding of the material and also catch the points that I might have missed. Professor Junker is patient and knowledgeable. I learned a lot from him.

(continued)

Comments

The materials of the second half of the semester were a bit hard for me to understand since I did not have much experience with causal inference, mixed effects model, and MCMC. I hope we could have more time to talk about these contents in details.

Office hours and Zoom recording helps me a lot in understanding and reviewing the class materials.

If there are more diagnosis for models or the methodology in the slides, rather than just theory, that'll be great. I really like that so far many slides have summarized the benefits and concerns for some new stuffs, and given us some practical suggestions.

The material is a little bit difficult for me.

The material is kind of hard for me, but piazza is a good thing, we can access the TA and instructor directly and immediately. Also, every class has recordings, so if students is ill, they can access the lectures through the recordings. The homework, take-home midterm and quizzes are very useful for us to review the knowledge we learnt in class. We applied everything we learnt in this course in the IDMRAD project.

Overall. the course is very well structured and I like the general content. I think we kind of go through some of the materials at the end really quickly however. From my understanding, most people in this program already know most of the basics of linear regression and the first half of the course, so perhaps that could be shortened a little bit and the other sections get a little bit more time spent on them.

I like how everything is submitted online, I just think in future iterations of the course, quizzes can also be done through gradescope to keep all of the assignments in one place.

It is pretty good for everything. One thing I really concerned is due dates in the last two weeks, like we have two assignments due in the same week which is pretty stressed.

The course content is pretty challenging to me. I have learned several topics during the undergrad school, but it was shallow. Professor and TA are both very kind and helpful when I have questions about the course material and homework. The IDMRAD project is the first time I do the full research process and finish a report by myself, which I think will have great help to my future study and career. I skipped several in-class lectures due to my bad mental condition (really sorry for that), but the recording is really helpful and explicit for me to understanding the knowledge. Overall, I love and enjoy this class!

I like this course and I think it's one of the most useful courses for me this term. Although the course material is a little bit difficult for me, the professor answers question during office hours and on the piazza, which helps a lot. The homework helps me understand the course material better.

Questions
I learned that there can be multiple level-2 equations in a multi-level model. One question that I have is whether you can assess the normality of a model from a conditional residual plot for lmer models or if you need qqnorm plot of the conditional residual plot to assess the normality.
I learnt how to transfer a mathematical formula of multilevel model as suitable R code. But I still have a question, if sometimes whether there is a "1" in the lmer formula or not doesn't change the meaning of the formula, why would we add it to the formula?
One thing that I learned from the past few lectures is the reasoning behind using Markov-Chain Monte Carlo simulations, and sampling high dimensional data of densities. One question I did have is how implementable and transferable are the tools and methods we used in this class with R into other languages such as Python.
One thing I've learned a lot through lmer and now stan is the differentiating of subjects and distributions that can be done to really 'personalize' models and predictions. One question I still have is determining what priors should be in general. For example: For a given problem, how do we decide what the first round of priors should be when we don't have an expert to consult?
Why the abstract of an IDMRAD paper doesn't have to be long.
I learned multilevel model which I never learned in previous courses. one question I have is that I don't quite understand how to connect bayes rule with multi-level model towards the end of this semester.
I like the Bayesian and hierarchical topics covering the late courses. Wish we could have more covered but, I guess see you next semester!
I learned how to modify and improve Bayesian models by interpreting the posterior predictive distributions.
I would love to know more about the different kinds of distributions used in Bayesian models to estimate parameter.
Coming into the class, I had a vague understanding of mixed effects model. It made sense to weight groups in regression analysis (in some scenarios) rather than treating each observation completely independently. After these past 4 weeks, I have a solid understanding of the mechanics behind different mixed effects models but more importantly the motivation why we design them in the first place. One lingering question I have is why do MCMC converge in the first place? I understand that they converge and what the burn-in is and what the steady transient state is. But I don't completely understand why/how the convergence happens in the first place.
I learned multilevel model, Markov Chan Monte Carlo stimulation. sometimes I find it's hard to interpret the multi level model.
Before this course I had never been exposed to hierarchical models. I see the real-world applicability of this modeling technique and I am glad I learned about them in this course. I still have questions about exactly how markov chain monte carlo works. I understand from a high level but there are some details that are unclear.
I learned multilevel models. My question is can we apply other types of Bayesian algorithms, such as DRAM, to run the models we have learned in class?
I've never used R stan and I've never heard of Bayesian modeling. It's been interesting learning about them!
I learned about hierarchical models and implementing them in R using lmer and other libraries. One question I do have, maybe you cannot answer, but where do you think is a good place look to find current / recent research being done. I find it interesting to see these implementations in practice. learned: multi-level regression, how to analyze database, how to write paper question: how do we check or improve our paper? (Except read it several times)

(continued)

Questions

One thing I have learned - I think I now have a good sense of how to approach with the data analysis and modeling when I face a raw data set from the wild. For instance, before this course I did not know about such variable selection methods as AIC, BIC, and DIC - now I have learned them, I think they are quite nice ways to derive our final model.

One question - I am still not so sure when we are using the Bayesian modeling (in real-world context). Why are we using Bayesian modeling instead of other methods we learned earlier in the class? It would be something to be clarified.

Multi-level model is pretty useful and new to me. I think this is a way which approach the real situation. lmer() and glmer() are awesome!

I'm still not quite sure about the difference between a Bayesian-based model and a model using lmer(). For example, which model can we use in what situations.

I was still a little bit about the bayes usage, sometimes I can't create a good model. I have a good understanding of transformation to improve the model accuracy .

Have more understanding about linear model and interpret models with R.

I think I really enjoyed the conceptual foundation for multilevel models and the paper on it really opened my perspectives on multilevel model and how they can be applicable in real life. I think non-parametric regression was interesting but I would have liked to see more lectures on it as it is a concept that can be discussed very much more in detail.

One thing I learned (which is mostly the same as what I answered for the mid semester evaluation) is that the interpretability of results is the most important part of any analysis. This is especially true when working with clients because they want clear answers and complicated models can be difficult to apply in the end. We also have had this fact reiterated in several of our other classes. One question I have overall is "How long does it take to feel confident giving statistical advice in real world client situations?" I know that I can do the analysis, but I still don't feel fully confident in applying what I know to an unfamiliar situation where it would matter if I was wrong. I'm assuming that the more experience you have, the more confident you feel, but maybe there's some way to check that accuracy of your results in a systematic, consistent way.

I learned how to fit the multilevel models, how to improve the model and also how to use Stan.

I learned the process of fitting mixed effects model and analyses of its diagnostic plots.

I don't have any question for now.

Multi-level model is the model I didn't know before. Good to learn it.

I think multi-level model is the useful thing I have learned, and I think I am still confused about the Bayes Rule.

Multi-level models. I have never had experience with them.

What other types of models would you recommend learning about that are outside of the scope of this course?

I am quite confused about the Bayesian stuff and how does it work in the model. I could construct the model but cannot clearly understand the mechanism of it.

One thing I learned is the multilevel model, I never used it before.

The multilevel model analysis is what I learned most. Previously, I would only consider the simple or multiple regression model during the model building process, while now, I have learned to consider more due to the dataset variables.

For the material mentioned in slide 26, I find it might make me feel confused about lm, glm, lmer, and glmer.
