MCMC: A clever way to run the numbers

Don van Ravenzwaaij
University of Groningen

I must have first heard about Markov Chain Monte Carlo (MCMC) sampling in my undergraduate years. For the longest time, I struggled with understanding, *really* understanding, what the point was of MCMC sampling. The technique was usually explained something along the lines of "We would like to know what some unknown underlying probability distribution is, but we cannot observe it: we do not have an analytical expression to obtain said distribution. MCMC allows us to *approximate* this distribution by sampling... [and so on and so on: the rest did not much to contribute to my understanding]". Sampling how? Surely, if we have a means to calculate probabilities based on the data, what do we need MCMC for? And if we do not have a means to calculate probabilities based on the data, what exactly does MCMC add?

My breakthrough came in one of the yearly Erasmus-IP (Intensive Program) Seminars on Mathematical Psychology, the 2009 edition in Tübingen. I had one year as a PhD–student behind me, and had started to learn more about MCMC, but had still not tackled my original issue: what is this magic that MCMC contributes? One of the seminars in 2009's edition of this program was given by Francis Tuerlinckx on MCMC sampling. There was a lot of excellent content in the seminar, but it was a single sentence Francis said that would have a profound influence on the rest of my career (and, incidentally, my understanding of MCMC sampling). Francis had been talking for a long time and my brain was bulging with equations, and here he goes: "Well, time to wrap up for me, it is almost lunch time. You could always, you know, for fun, try and build one of these Gibbs samplers yourself." Credits, the end, everyone off to the lunch hall.

Except I was there first. I front–loaded as much food as I could carry and went off to my dorm room with my trusty laptop and a precariously balanced pile of chow. I had an hour and a half to make this work and I'd be damned if I would not have a serviceable version of a Gibbs sampler before the seminar would resume. Lots of cursing could be heard in the adjacent rooms in the next hour and a half and I shall not claim my final product was particularly elegant, but work it did. After the lunch break I walked up to Francis and showed him my work. He smiled at me (a clear indication of how nice a person he is, as my code really was a mess), and said "Good job".

So what exactly did I learn from building my own MCMC sampler? Well, for one, it is not a magic calculator. Armed with a likelihood, you can calculate probabilities for every conceivable combination of parameter values without needing an MCMC sampler (well, rounded to a certain decimal). The trick, as it turned out, is that doing all these calculations takes an incredible amount of time. The more parameters you are working with, the more combinations of parameter values there are, and the time it takes for all the calculations that are involved increases exponentially. What MCMC did was provide a clever way of figuring out what probabilities to calculate and approximate the underlying distribution based on a few strategically chosen calculated probabilities, rather than slogging through the whole grid of possible numbers.

The overarching thing I learned from this of course was: staring at a bunch of equations (though definitely useful) only gets you so far. By implementing "mathsy" stuff yourself, such as MCMC samplers, you get a much more thorough understanding of what is really going on under the hood. Over the course of the years, my understanding of MCMC samplers deepened, enough to write a tutorial about how they work and hopefully contribute to someone else's understanding of MCMC samplers (van Ravenzwaaij, Cassey, & Brown, in press).

The take–home message should be: "Do you want to know how something works? Build it yourself!" The examples provided in this chapter and in van Ravenzwaaij et al. (in press) should get you started, but don't shy away from experimenting with more complicated versions of MCMC samplers! Even if you never end up using them (I certainly did not touch my first Gibbs sampler ever again), the learning experience it provides could prove much more valuable.

References

van Ravenzwaaij, D., Cassey, P., & Brown, S. D. (in press). A simple introduction to Markov Chain Monte–Carlo. *Psychonomic Bulletin & Review*.