

# David MacKay

David MacKay is a scientist at Cambridge University and an enthusiastic user of  $\text{\LaTeX}$ .

[Interview completed 10 February 2009.]



*Dave Walden, interviewer:* I already know something of your background, since you and I met in person two or three times in the 1990s; and, in the years since, I've periodically reviewed your web site (<http://www.inference.phy.cam.ac.uk/mackay>), including your personal (<http://www.inference.phy.cam.ac.uk/mackay/bio>) and scientific (<http://www.inference.phy.cam.ac.uk/mackay/bio2.html#scibio>) bios. Nonetheless, for readers of this interview, please speak about your personal history.

**David MacKay, interviewee:** I studied Natural Sciences at Cambridge University (1985–1988), specializing in Physics; then did a PhD at Caltech in Computation and Neural Systems. I've always enjoyed working in a wide range of sciences, and computing, engineering, and mathematics too. At Caltech I considered working in experimental neuroscience but decided that theoretical work was more my thing. I did a PhD on applying Bayesian probability theory to neural networks, so as to make them work better.

Back in Cambridge I took advantage of my postdoctoral research fellowship to pursue some new research interests: via my interest in approximate inference methods, I stumbled into coding theory, developing sparse-graph codes for communication systems. Radford Neal and I managed to reinvent low-density parity-check codes (first invented in 1962 by Bob Gallager, but forgotten by the coding community for 30 years), and the codes that I developed have influenced some new standards for satellite broadcast and for hard drives. What the coding theory community realised at the end of the 90s was that most of Shannon's communication problems are best solved by codes based on relatively simple sparse random graphs—rather than the elaborate algebraic coding theory of the previous 30 years. These sparse-graph codes are decoded using simple local message-passing algorithms. I happened to be writing a textbook on information theory and machine learning (*Information Theory, Inference, and Learning Algorithms*, Cambridge 2003) while this revolution in coding theory happened, so this book managed to become (among other things) the first textbook on modern coding theory. I like to make and exploit connections between diverse topics, and this textbook also features a chapter on evolution, natural selection, and sex.

Another interest that developed during my postdoc was the idea of making a new communication system by taking arithmetic coding (the state of the art in text compres-

sion) and turning it on its head to define an information-theoretically-optimal gesture-to-text system. I wrote a first prototype of Dasher, which is what we call the software based on this principle, and recruited a PhD student to develop the idea. Dasher is now distributed as free software and is part of the GNOME desktop.

After three and a bit years of postdoc I was lucky enough to get a faculty position in the Physics department in Cambridge, teaching Physics, and continuing my research in information theory, approximate inference methods, and computational neuroscience. Over a period of four years, I did several months of voluntary teaching in Africa. Then I decided there was a need for a numerate, factual, unemotional book about energy, to guide constructive conversations on energy policy. During the last three years I have devoted more and more of my time to energy matters, especially the public understanding of 2+2.

*DW:* In fact I downloaded Dasher from its web site (<http://www.inference.phy.cam.ac.uk/dasher/>) and tried it while preparing for this interview, and I have been reading your writings on energy.

But let's come back to your 2+2 and other work later, and first ask a  $\text{\TeX}$  question. When and how did you first encounter  $\text{\TeX}$  or  $\text{\LaTeX}$  and begin to learn to use it?

*DM:* My second summer job, as a student (1987), was at a Ministry of Defence establishment, working in a pattern-recognition/machine-learning group. The group had a cluster of Sun workstations running Unix on which I ran neural networks as background jobs. During my ten-week placement I picked up  $\text{\LaTeX}$ , which was the recommended way of writing reports in the group. I wrote two reports—one on my research project, and a second one reviewing the field of neural networks, which was having a loud re-birth at the time. (I'd attended the first International Conference on Neural Nets earlier that summer.) A year later when I arrived at Caltech, Sun workstations and  $\text{\LaTeX}$  were standard there too.

*DW:* I believe you have written two big books (the abovementioned *Information Theory, Inference, and Learning Algorithms* (<http://www.inference.phy.cam.ac.uk/mackay/itila/>) and *Sustainable Energy—without the hot air* (<http://www.withouthotair.com/>)), and presumably many papers, using  $\text{\LaTeX}$ . Please tell me about the  $\text{\TeX}$  distribution, editor, and other tools you use for this work, and how your approach has developed over time.

*DM:* Initially on Sun workstations I was using the Suntools text editor, and perhaps I wrote my PhD thesis in that; but for ages (probably since 1992) I have been using Emacs. I think I picked up  $\text{\BibTeX}$  to automate my bbls about 1992 too. I stubbornly used  $\text{\LaTeX}$  2.09 as long as I could, because if things are not broke I don't want to fix them. So I was still using  $\text{\LaTeX}$  2.09 while writing the first book on Information Theory. It probably would have been good to switch, because (doing my own document design) there were a lot of things that probably would have been easier in  $\text{\LaTeX}$  2<sub>ε</sub>. Eventually, towards the end of writing that book (which took me eight years) I did switch to  $\text{\LaTeX}$  2<sub>ε</sub>.

I used xfig and gnuplot to make PostScript figures (to include in the  $\text{\LaTeX}$  documents) for a long time, but always found the results a bit disappointing—especially those of xfig. A friend introduced me to MetaPost in about 2002, and I have become quite keen on MetaPost as a replacement for xfig. I started using MetaPost in the first book. MetaPost is a product of the  $\text{\TeX}$  world, so (a) you can get things to look perfectly matched with the  $\text{\TeX}$  text; and (b) you have to learn a bit of classic  $\text{\TeX}$  to be happy using MetaPost. I still use gnuplot for most graphs, but I am thinking of migrating graphs to MetaPost eventually. I love the appearance of MetaPost output, but working

with MetaPost is always a bit cumbersome, and the error messages are often worse than useless. During the first book I used a nice Emacs mode called `reftex` to automatically traverse the hundreds of TeX files that made up my book.

At some point a system upgrade broke my `reftex` mode, however, so I haven't been using it in the second book. I just open all the TeX files manually with Emacs, and switch buffers by hand instead of by `reftex-magic`.

While writing my own document designs for my books, I have rewritten the definitions of many of the objects in L<sup>A</sup>T<sub>E</sub>X (parts, chapters, sections, figures), but I don't really know what I am doing. I'm sure it would be a good idea to read *The TeXbook* properly some time. Some L<sup>A</sup>T<sub>E</sub>X packages that I always use are `booktabs` (to make L<sup>A</sup>T<sub>E</sub>X tables look great instead of crappy); and `colordvi` (to get colourful text). I currently use `natbib` to handle my in-text citations.

I don't use L<sup>A</sup>T<sub>E</sub>X only for my papers, for my two books, and for examining and teaching.

I also use L<sup>A</sup>T<sub>E</sub>X in presentations, but in a rather non-standard way. I don't like the standard L<sup>A</sup>T<sub>E</sub>X presentation style — I only tried using it once or twice. It's a bit too slow and cumbersome to use and the results tend to be dull and stereotyped. (I especially loath those fake three-dimensional blue m+m bullet points.) I've become addicted to a little known super-lightweight presentation tool called `magicpoint`. `Magicpoint` has the lightest imaginable slide-writing syntax, and it's very easy to include colourful L<sup>A</sup>T<sub>E</sub>X in it, using a simple script to filter L<sup>A</sup>T<sub>E</sub>X source embedded in the `magicpoint` file (<http://itila.blogspot.com/2008/08/how-to-use-latex2epssh-to-make.html>). So, when I'm out giving talks, the main way in which I am using L<sup>A</sup>T<sub>E</sub>X is for little one-liners to make nice-looking simple presentations.

**DW:** The L<sup>A</sup>T<sub>E</sub>X part of your web site (<http://www.inference.phy.cam.ac.uk/mackay/SourceTeX.html>), has a six-page introduction to L<sup>A</sup>T<sub>E</sub>X called "Please write your report using L<sup>A</sup>T<sub>E</sub>X", stating "L<sup>A</sup>T<sub>E</sub>X is the typesetting program to use in scientific publishing." In your view, is suggesting L<sup>A</sup>T<sub>E</sub>X simply part of teaching students how to get along in the scientific community, or is there more to it than that? Also, is there explicit support for learning and using TeX in your college or within the whole university, or is each department or research group more or less on its own?

**DM:** The last time I looked at alternative ways of writing reports, I felt there was no alternative. I've occasionally been forced to use Microsoft Word, and have hated every moment. I think the logical structure of L<sup>A</sup>T<sub>E</sub>X is good for scientists, emphasizing content and leaving the details of document design to Leslie Lamport. If people want a WYSIWYG environment, they can always use a clever L<sup>A</sup>T<sub>E</sub>X front-end such as LyX.

One of the main reasons I recommend L<sup>A</sup>T<sub>E</sub>X is I often meet grad students who are groaning about the pain of dealing with their bibliographies. A user of L<sup>A</sup>T<sub>E</sub>X/B<sup>I</sup>B<sub>T</sub>E<sub>X</sub> who has a sensible makefile never has any such pain. The bibliography just happens, and is perfectly formatted in whatever house style is desired.

I'm not saying there are no alternatives out there, of course. But I think I would recommend all students to use L<sup>A</sup>T<sub>E</sub>X for its bibliography-handling alone.

For mathematical students I think L<sup>A</sup>T<sub>E</sub>X is absolutely mandatory. In L<sup>A</sup>T<sub>E</sub>X you can easily make math expressions look professional. It's very rare that people manage to make good-looking math without L<sup>A</sup>T<sub>E</sub>X. (Though of course I know it is possible.)

**DW:** Do you participate in or draw on the resources of the worldwide TeX community?

**DM:** Yes, I've often used CTAN when I was adjusting my book design and trying to use extra features of L<sup>A</sup>T<sub>E</sub>X. I've sometimes emailed other members of the open source TeX

community to ask for their insights, and I think I may have once or twice contributed patches.

*DW:* As a major user of L<sup>A</sup>T<sub>E</sub>X but self-professed as not being an expert, no doubt there are things about L<sup>A</sup>T<sub>E</sub>X that frustrate you.

*DM:* Yes. Because L<sup>A</sup>T<sub>E</sub>X is beautifully content-oriented, and doesn't give you control over layout on the page, there is inevitable frustration when you are actually trying to lay stuff out on the page :-).

The way that L<sup>A</sup>T<sub>E</sub>X places floats (figures and tables) and handles page breaks in text is what I'm talking about. Friends have given me some helpful hacks which seem to help, sometimes. For example:

```
{% begin troublesomepage hack
% this should be *before* the start of troublesome page
\renewcommand{\floatpagefraction}{0.8}

blah blah

% troublesome page
...
}% end troublesome page hack
```

but when these problems arise, it feels a lot like herding cats. Or playing a marble/maze/holes game.

*DW:* I'm in the same boat with you — major-user-non-expert. When it's time to prepare the final manuscript, I find myself making a good bit of use of the position option in figure and table environments, some use of `\enlargethispage` plus or minus one or two lines, and occasional use of changes to the `\floatpagefraction`, etc., commands. Maybe a reader of this interview will have suggestions for dealing with this issue more automatically.

*DM:* Another minor defect of basic L<sup>A</sup>T<sub>E</sub>X 2.09 and L<sup>A</sup>T<sub>E</sub>X 2<sub>ε</sub> is the way that it is willing to let marginpar objects overhang the bottom of the page. I make heavy use of marginpar for marginal figures, and it's a shame to have to manually check all of them.

*DW:* Both your books have the entire book available as a PDF for free download in addition to being available in paper copies. The first book was published by Cambridge University Press, and the publisher of the second book is listed (at Amazon.co.uk) as UIT. Please tell me what your thinking is about making these books available to the world, and what were the publishers' points of view.

*DM:* I am enthusiastic about open-source and the fact that, in the digital age, information can be copied almost for free. I think it makes sense to exploit this virtue of electronics.

UIT Cambridge is a traditional but relatively small publisher, who was willing to go along with the "free online" approach. Cambridge University Press are also now happy with "free online", but when I initially suggested the idea to them, they were not so keen. My view on free-online books is that the *only* way for written materials to have a significant impact is for them to be available for free online. If work is not available free, people will simply move on and read something else that is free.

I think that giving away a book for free online is also a good way to enhance sales of the paper book, as long as the book is a good book. If people download my information

theory textbook and print it out without paying, I am delighted. They have chosen to spend their own printing resources on creating a piece of advertising for the book! Publishers spend lots of money trying to get advertising onto peoples' desks, and most of that advertising just turns into junk mail that goes straight in the bin. What a waste! If Fred prints out my textbook in order to read it, maybe ten of his friends will see the printout, and maybe two of them will end up buying the textbook on paper. Last time I checked, the information theory text had sold 10,000 copies, so I don't think that CUP is disappointed. I think they view it as a success.

I believe that *Sustainable Energy—without the hot air* has a good chance of being commercially successful too, not “in spite of being free online”, but “because of being free online”.

One benefit of the open free-online attitude is that the author can put half-written materials on the web before publication, and benefit from bug-fixes from early readers. Both my books have certainly benefitted in this way.

I am passionately enthusiastic about open source software. I depend on it. Everything I have done for the last 15 years has been done using open source software. It's not a perfect system—two examples of imperfections with open source:

1. I worry that the open source community will keep on creating more and more new apps (some of which may be flashy but flakey replacements for existing high-quality old apps) and will not have enough enthusiasts to look after the old apps.

2. My group's software, Dasher, which is free software mainly intended for people with disabilities, is probably not reaching its target community very effectively because that community is served by middlemen salesmen with glossy catalogs and commissions to earn; free software doesn't fit into the glossy catalog system so it doesn't get promoted.

My group is a partner in a new open source research project led by Sun. The project is called AEGIS and the aim is to develop a new open source programming framework such that all apps developed in that framework will have excellent accessibility support, i.e., support for users with any disabilities).

**DW:** Do you do your own book design, as well as typesetting, or do your publishers help you with the design?

**DM:** When I was signing up with CUP (in about 1996) I suggested to them that they should give me a book design. They showed me some options, but I didn't like them, so I decided to make my own book design, and just ask CUP's “experts” for  $\text{\TeX}$  help when I was trying to make particular features work. In the end I asked their expert for help once, and the expert was useless and gave incorrect advice. So I did it all myself, with a little bit of feedback from friends (especially Sanjoy Mahajan (<http://mit.edu/sanjoy/www/>)); a little bit of advice from a typographer, and with the help of (a) Tufte's books and (b) Robert Bringhurst's *Elements of Typographic Style*; and of course by poring over other textbooks for design ideas.

**DW:** Returning to your interest in energy and your mention of  $2+2$ , will you say a bit more about what your interests and intentions are with this work. Or is this more of a hobby—or does it somehow fit within your academic situation?

**DM:** There is a problem of wishful thinking today. People seem to believe, when it comes to energy, that  $2+2 = 100$ . But the correct answer is that  $2+2 = 4$ . People are anti-fossil fuel, anti-tidal barrage, anti-wind farm, and anti-nuclear. Yet they want the lights to come on. They seem to think that a few renewable facilities the size of a figleaf could power a modern lifestyle. In Britain, total energy consumption is 125 kWh per day per person. (In the USA, it's 250 kWh per day per person.) In Britain, renewables provide about 1 kWh

per day per person, or perhaps 2, depending how you do the thermodynamics. Of that, wind power contributes less than 0.2 kWh per day per person. If we had a really big increase in wind power (against public opposition), it might deliver 4 kWh per day per person.

Hardly anyone is talking about energy plans that add up!

I wrote the book to try to help the public have an intelligent numerate and constructive conversation about energy (instead of the normal Punch and Judy show of anti-wind and anti-nuclear).

This activity is an amateur interest for me, but it has become a near-full-time activity, and it is an activity that my department supports warmly. Energy is a popular topic with students, and I'm engaging with academics (in economics, engineering, and business) who work on energy policy.

I haven't figured out exactly where I am going next with this full-time hobby. I wrote the book to try to get all the numbers down in one place, so that now I can talk to people. The task for the next couple of months is to figure out whom I should talk to. I'm not sure whether I should focus on public education, or on talking to civil servants. We need a plan that adds up. It's not going to be easy (especially not in a high density country like Britain) — but it is possible.

*DW:* Since learning about your book, I have been recommending it to everyone I know who is interested in thinking about the energy situation based on some factual numbers rather than an a priori political or moral position, and it surely doesn't hurt that it has a completely professional and authoritative feel while being available on-line for free.

Perhaps we can close with another personal question. You write and think about the energy problem. However, it seems you also try to live a green life to some extent. Your web site mentions that part of the reason for returning to Cambridge after Caltech was getting back to a bicycle rather than car society; your web site also mentions your involvement in something called Cast Iron aimed at replacing a bus system with a train; and I loved your analysis of the efficient way to boil hot water (<http://www.inference.phy.cam.ac.uk/sustainable/hotwater/>). Can you say a few philosophical words about what you see as the right mix of just trying to get along in life (make a living, stay warm in the winter, etc.), spending time trying to understand the realities of the situation, living green at a individual level, and trying to improve the world more generally.

*DM:* When I give talks about energy, people quite often ask me what car I drive, how much I fly, and so forth; it goes down well with these people if my response indicates that I'm trying to live a green lifestyle. And indeed I don't drive cars any more — I just ride bikes and trains and buses.

But the truth is, when you honestly quantify everything, it's incredibly difficult to get close to sustainable targets, in particular to the greenhouse gas emission targets for 2050 recommended by climate scientists. So while I enjoy playing the game of reducing my energy consumption at home (in the winter, because my home is now usually below 55 F, I've cut my natural-gas consumption by 60 percent), I honestly view these personal privations as curious science experiments rather than effective actions. Even my 60-percent-reduced domestic gas consumption still contributes 2 tons per year of CO<sub>2</sub>! (And the recommended target for 2050 is 1 ton per capita per year.)

So for me the bottom line is pretty clear: the only way to live a useful green life is to influence governments' policies. Through policies, we can transform the way everything works. There will still be consumer choice; but through good legislation we can ensure

that the consumer can have any colour, as long as it's green. If it seems like I have to get on an aeroplane occasionally to try to influence government policies, I'm going to do so. Maybe I'll see you in Boston this year!

*DW:* Thank you, David, for taking the time to participate in this interview. I do hope we meet again in person at some point. I recommend to readers of this interview that they spend some time scouting around your web site; there's a lot of interesting stuff there.