



BIROn - Birkbeck Institutional Research Online

Eve, Martin Paul (2018) Professorial Inaugural Lecture: Digital Architectures: Literature, Interpretation, and Computation. In: Martin Paul Eve Professorial Inaugural Lecture, 12th September 2018, Birkbeck, University of London, UK. (Unpublished)

Downloaded from: <https://eprints.bbk.ac.uk/id/eprint/23821/>

Usage Guidelines:

Please refer to usage guidelines at <https://eprints.bbk.ac.uk/policies.html>

or alternatively

contact lib-eprints@bbk.ac.uk.

Digital Architectures: Literature, Interpretation, and Computation

Text delivered at the inaugural lecture of Professor Martin Paul Eve at Birkbeck, University of London, 12th September 2018.

We have known for quite some years that the medium through which the text of literature is delivered is both central to its reception but also of academic interest. The study of the history of the book and of material texts has an excellent pedigree, although I'm going to try to avoid citing overly trite formulations from Marshall McLuhan.

It is also the case now that for at least the past thirty years, almost every commercially published work of literature has been a digital object before it has been a print volume. The materiality of the book, in the past half-century, has moved from manuscripts, to carbon copies and typewriter-ribbon typescripts, to magnetic storage arrays, backlit CRT screens, solid state drives, through to network transmission. Certainly, the codex, as perhaps the best technology ever developed for humans to seek both sequential and random access, persists. But the field of the future history of the book, as Matt Kirschenbaum has called it, will also have to envelope XML markup, XSL transformations, Adobe InDesign workflows, digital delivery platforms, network infrastructures, word processing software, engineering patterns, machine learning, and natural language generation within its field. Future historians of the book must be multimedia and technological historians if they are truly to understand the mediations through which literature passes in the twenty-first century.

The seeds of such work are already being sown. To take the last of my list, natural language generation, it is already the case that sports columns are being written by machine. Indeed, feeding statistical sports data to recurrent character-based neural networks is a common practice when creating summarised standard journalistic writing in this area. Indeed, the desired formulaic and simplistic nature of this writing is perfect for this kind of process. But lest we think ourselves too superior, I should point out that we academics are also not exempt from the rise of the machines. **[SLIDE]** Take for instance one of my recent weekend hobbyist games in which I used my home

computing rig to train a recurrent character-based neural network on the literary studies journal *Textual Practice*. A character-based neural network essentially samples input data to build a model of the source texts on which it operates. The model, of course, is reductive and aims to statistically replicate – with a certain degree of loss – the underlying data, while never fully representing them. After just twenty-four hours of training, though, my network was able to produce some alarmingly good parodies of academic writing. For instance, the network told me that “The series of temporal inventions of the object is intelligible only afterwards”, which certainly sounds a probable proposition. Yet the network also produced statements that seem worryingly self-aware. For ‘The problem’, as the network aptly phrased it, ‘is that the poem is a construction of the self as a strategy of self-consciousness and context’. So what might it mean that there is no self-consciousness or context behind the generation of such a statement? What does it mean when a neural network writes, as though writing of itself, that it quote ‘provides the fraud of the epistemological practices of knowledge’? The network writes also that it “shall find our intellectual values, by rewriting their very ties’. It seems clear, also, that this particular network inherited an academic sense of humour in its generation of footnotes, telling me of a fictional work, previously unknown to me, entitled ‘The privatized climax’., by John Spottisley, from 1929, p. 4, emphasis in original’.

As humorous as such an exercise may be – and I do find it funny – there is a serious point from which I want here to kick off. Namely, that the ways in which we read media, content, and form of literature are now conditioned by digital technologies. For if, as Donna Haraway once put it, ‘Writing is pre-eminently the technology of cyborgs’, the balance of activity between our inert selves and our lively technologies appears ever more to be weighted in favour of the latter. Indeed, I believe that it makes little sense to speak now of contemporary fiction as somehow opposed, in a rarefied way, to computational environments – as novelists such as Jonathan Franzen are wont to do. It is with this in mind that I titled this lecture *Digital Architectures*, playing on Gérard Genette’s notion of the architext – that is, with a focus on the placement of text within categories as the object of the study of poetics. While, as Genette puts it, “Jokes on the word text form a genre that seems to me indeed overworked”, I do want to posit that another useful category for this type of taxonomic

exercise might be digital infrastructures, and the relationship of literatures to their enmeshment within digital architectures. For the digital is the underlying architecture on which all contemporary fiction, in one way or another, is now hung.

In some ways, I feel that contemporary literary studies has already misrecognised the consequences of this shift. The recedence of textual scholarship – that is, the study of the conditions and sociology of publication for contemporary fiction – seems to me evidence that somehow we believe that digital methods of book production have eradicated human error from the process. This, seems to me, clearly not to be the case.

For just one example, in 1992, Jennifer Egan (a Pulitzer prizewinning author) won the UK Cosmopolitan short story award with the piece “Sacred Heart”. [SLIDE] As a result – the prize – Picador published her first collection, *Emerald City: The Collected Works of Jennifer Egan* (hereafter referred to as Emerald A). At that point, however, Egan had written neither “Why China?” nor “Sisters of the Moon”, the two stories that she considers the strongest in the eventual collection from 1996 that would receive the same title, “*Emerald City*” (hereafter, Emerald B).

It is, therefore, of interest that the “prior UK edition” of *Emerald City* is, in Egan’s own words “missing material and full of mistakes and hopefully consigned to oblivion by now”. It also, though, contains short-story material that has never appeared elsewhere and of which critics are currently unaware. The original version of the collection is nearly impossible to buy at the time of writing, but it is available for consultation in national deposit libraries in the UK, such as the British Library. If future work is to cite this material, I argue, scholars need to be aware of the potential version variants and alternative implications of the different texts by this prominent and important American author. Yet, very few people in the space of contemporary fiction studies are conducting such bibliographic work, or using digital techniques to map such differences.

At this point, it is probably necessary to say a few words on the “digital humanities” and the controversies that enmesh this particular field of study. For, I have been a computer programmer far longer than I have been a literary-studies academic. [SLIDE] And certainly, reading works of literature with the aid of statistical techniques is controversial and pointing out literature’s inescapable enmeshment within technological fields is, to some, heretical. In the minds of some commentators, it’s not just that quantification is the opposite of literary art, it’s not merely that, apparently, counting the word “whale” in *Moby-Dick* can only ever tell us one thing: how many times the word whale is used in *Moby-Dick*, as Timothy Brennan put it. It is, instead, that, for many people, the digital humanities are increasingly seen as “neoliberal”; a critique voiced most forcefully by Daniel Allington, Sarah Brouillette, and David Golumbia. I want to pull apart this argument a little more.

In order to assert that the digital humanities are neoliberal, one must first define both the digital humanities and neoliberalism. Despite the mind-boggling quantity of nonsense written on the topic, it turns out to be neoliberalism that is the *easier* of the two on which to be precise. Put bluntly, neoliberalism is a mode of political economy that emerged most prominently from the 1980s onwards in which politics is replaced by economics as the dominant societal logic, to use William Davies’s pithy formulation. Under such a logic, economics is the most powerful of the academic disciplines and must form the basis for all State decisions, with the state itself contracted merely to ensure that the conditions for market exchange can be legally enforced, even while the state itself must work on a market logic – the state under the supervision of the market, as Michel Foucault had it.

Within such a definition, one can begin to see how the digital humanities as a field might look a bit, well, neoliberal. If, for some, neoliberalism is the disenchantment of politics by economics, then, for others, digital methods for studying literature become the disenchantment of art by mathematics and computers, an old debate by now in literary theory that goes back at least to the early twentieth-century Frankfurt School thinkers. Yet, to be blunt, the histories don’t quite match up. Literary

studies has used quantitative – and even computational – methods for quite some time. For instance, as Erik Ketzan, one of my Ph.D. students, pointed out to me, Thomas Schaub claimed in 1981 that “the word ‘bloom’ is one of the most oft-repeated words in” Thomas Pynchon’s *The Crying of Lot 49* (although he’s actually wrong about this). Dartmouth College offered a course entitled “Literary Analysis by Computer” as far back as 1969. Further, as Nicholas Dames has noted, Vernon Lee proposed a “statistical experiment” – a quantitative analysis – on literature in her 1923 *The Handling of Words*, itself prompted by a letter to *The Times* from Emil Reich, several years earlier. The term “code” has been used extensively in the history of concrete poetry and various types of machine formalism have emerged in that space, as another of my students, Bronac Ferran, has charted. Quantitative and even computational approaches to the study of fiction are really not that new.

Yet it is the development of digital, computational technologies, rather than just applied mathematics, and their application to the study of artforms that really gets the neoliberal accusations flowing. The libertarian political cultures of Silicon Valley, for instance, feed into a perceived economic amorality of technology to which literature can be posited by scholars as a correcting humanist and even democratic force, notwithstanding the troubling moralistic value framework that sits below such an assertion.

At the same time, though, there are many contradictions and hypocrisies in the critiques of digital humanities as neoliberal as they stand today. The most pervasive of these, drawn to my attention by Ted Underwood, is that while many defend the humanities disciplines for their uselessness in the current age of economic utility (that is, the fact that the study of English literature is “merely” an end in itself), it is demanded that the digital humanities be *useful*. That is, the most prominent critique of DH is that it “doesn’t tell us anything useful” about literature (remember the slight on counting words, above), even while the broader study of literature is defended on the grounds that it doesn’t have to tell us anything useful.

The second hypocrisy is the lack of cultural critique around digital self-mediation. By this, I mean that while digital editions and digital approaches to the study of literature are criticized on the grounds of coercing their material into formats that are amenable to the digital, the fact that those very same people who voice such critiques will work, daily, on a word processor, writing books that will go to print, but that began digitally, seems tricky. They do not seem to feel that this compromises their own work, or somehow makes *it* neoliberal. The neoliberal critique, it seems, is only applicable to those who might *build* tools that others can use. So much for the famous pronouncement that “the point is to change it”, if you build digital things in the humanities, you will be tarred with the neoliberal brush.

[SLIDE] Finally, and most importantly, though – and this is where I will bring my thinking on open access into the picture – I think that digital technologies contain within them a seriously radical potential in their creation of non-rivalrous objects. As Peter Suber has extensively outlined, until the advent of the digital, all forms of writing were bound within rivalrous objects. That is, if I took your book object (a codex), you would no longer have access to the text within it. This is why the UK definition of “theft” includes the intent to deprive the owner of the original and why it is a poor metaphor for the digital age, for copying does not deprive the owner of the original artefact of anything. The advent of Xerox photocopying changed this relationship of text to object, although it still remained cheaper to buy the original book than a machine that could copy and/or transmit it. In the digital realm, the cost-per-copy is radically reduced, to nearly zero. The ability to copy “objects” – if one can call digital things, objects – at an infinitesimal cost causes profound changes in textual possibility. The first is that all costs of production are front-loaded; it is the cost to first copy that matters, and the labour there invested. Once that labour has been covered, the ongoing cost is extremely low. The second is the potential, within certain realms of not-for-profit education, to distribute research (and teaching) material to anyone who wishes to read it, with no charge.

There are many nuances and caveats to this, into which I don't have time fully to delve today, but I want to focus on two aspects here: rivalry and knowledge; and why research is different. On the first of these, consider that, in Peter Suber's words:

“Knowledge is non-rivalrous. Your knowledge of a fact or idea does not block mine, and mine does not block yours. Thomas Jefferson described this situation beautifully in an 1813 letter to Isaac McPherson: ‘If nature has made any one thing less susceptible than all others of exclusive property, it is the action of the thinking power called an idea....Its peculiar character...is that no one possesses the less, because every other possesses the whole of it. He who receives an idea from me, receives instruction himself without lessening mine; as he who lights his taper at mine, receives light without darkening mine.’”

The point that Suber, I think mostly accurately, deduces from this is that the digital realm is actually uniquely brilliant for conveying ideas; a non-rivalrous form. For the first time in human history, the underlying dissemination vessel can match the form of the ideas being transmitted.

The second point to which I will turn, though, is the way in which research production differs under our current economic conditions from other forms of cultural production. **[SLIDE]** Research in the university remains one of the last remaining systems of payment through patronage. Researchers are relatively free in their selection of research area (although one could argue that there are many soft-power mechanisms that steer such choice), and then are free to keep the copyright on the resultant output, to publish where they see fit. There seems to me no other area in society, apart from arts funding, that really operates on such a basis.

Yet, criticisms of open access, and of the digital space in general, continue. There are Marxist critiques – from people like David Golumbia – and Marxist defences and arguments – from figures such as Joss Winn. There are arguments that open access is neoliberal and there are arguments that open access might allow us to regain the commons, as work by my student Stuart

Lawson has gestured towards. **[SLIDE]** There are arguments that open access has only increased the profits of the mega-publishers and arguments that open access contains within it the seeds of a new world of social justice and equitable access to knowledge. This is the battlefield within which my scholarly communications work is fought. As Cameron Neylon put it, to work on open access is often to find oneself “described as 'neo liberal' (alongside uber capitalist) and 'Marxist'” all within just a few breaths of one another. My personal belief is that the radical gift of knowledge that we can bestow to society using digital technologies contains within it a socialist principle at its core and I work towards open access, in the humanities disciplines, for that reason.

But let us return from open access to the study of literature using computational methods and the study of how computational methods are integral to the architecture of contemporary literature. I have written various tools – pieces of software – in this space that undoubtedly fall under Allington, Golumbia, and Brouillette’s definitions of “neoliberal tools”. That is, these are tools that can be used by others that quantify various aspects of literature, which I see it as part of my research work to produce. Because this chimes with software production paradigms elsewhere in Silicon Valley, the critics feel threatened. They believe that “This view”, that there might be a definition of the humanities that includes computation, in Allington, Golumbia, and Brouillette’s critique, “reaches its apotheosis in the repeated suggestion that building computational tools should qualify as a replacement for scholarly writing”.

I would not say that I’ve gone down this road; I have plenty of scholarly writing to back me up. But there is a longer precedent for tools in the humanities that the critiquing triad are careful to dodge: literary theory. Foucault himself called his work a “toolkit” (although I’m also aware of the recent debates on Foucault’s association with neoliberalism). The worst “applications” of literary theory have simply taken a concept and applied it as a blanket “tool” for quite some years. But I want now to dive a little into what I have been doing and how the idea that the digital humanities develops “tools” that can just be uncritically applied rings false.

I am particularly interested in ways in which repetitious, brute-force digital practices can confirm or refute existing commonplace statements in literary studies. For instance, Andrew Piper’s recent work on machine-learning classification of fiction and non-fiction is a case in point. Since

the poststructuralist era of the 1970s, it has been a commonplace of literary philosophy to state that there is no fundamental distinction between the literary form and the non-literary. That is, there is nothing that a work of fiction can do to persuade you of its fictionality or factuality and vice versa in the case of non-fiction. Searle and Derrida both convincingly claimed this. For Searle, “The utterance acts of fiction are indistinguishable from the utterance acts of serious discourse” (in *Logical Status*), while for Derrida, “No exposition, no discursive form is intrinsically or essentially literary before or outside of the function it is assigned” (in *Demeure: Fiction and Testimony*). The only problem is that Piper has shown that machine classification can distinguish between fact and fiction with over 95% accuracy using just a 1,250-word stretch of text. To be clear: this machine is not determining whether a text is true or not – it doesn’t verify a relationship to the outside world. It verifies the work’s linguistic relationship to intention. It can distinguish when authors are writing fiction from when they are writing fact, but not whether the facts are actually true. The number of exclamation marks, word-percentages pertaining to families, and other features such as personal pronouns all play a role in detecting this linguistic relation to the expression of intention. Certainly, one could dispute whether this refutes the earlier assertions, but it says something extremely interesting about the relationship of large-scale empiricism against the apparently logical reasoning of literary-linguistic philosophy. Perhaps, literature is less singular, as my Ph.D. examiner, Derek Attridge might put it, than we might like to think.

Yet it is our proximity to texts that interests me and while the predominant argument about so-called “distant reading” is that it might allow us to ingest a large number of texts – more texts than one could read in a lifetime – I am interested in what digital approaches bring us in the space of close reading. One of the earliest pieces of digital work that I conducted was an attempt to quantitatively understand interconnectedness in one of the most interconnected novels of the postmodern age: Thomas Pynchon’s *Gravity’s Rainbow*. **[SLIDE]** Using modifications that I made to a tool written by David McClure, I created a force-directed graph of the top n terms in a text after computing a probability density function (using kernel density estimation) filtered by Bray-Curtis dissimilarity. This results in a map of clustered and connected terms within a specific distance; which terms occur “together” within a specified distance and are, therefore, most closely connected

within a work. My first attempt at a form of computational formalism through spatial network mapping.

There are several features of this network map that I found worthy of comment: We can note that character names are all peripheral while abstract terms, things and actions are usually more strongly linked. This seems to tie in with most assessments of Pynchon's character models. However, this is a common linguistic model for novels that is to be expected, except in the case of central characters in narratives; absent in *Gravity's Rainbow*, which has no single set of central characters but instead spreads the action over about 400 characters.

My critical assertions in my first book that a nostalgic? looking "back" as key to Pynchon's works (in my study focused on Lot 49) seem here to have some validity but one that is less centrally remarked upon in GR; the term "back" is one of the central and most connected nodes in the novel.

There's an interesting metahistorical clump around the 7pm mark; history, power, political, system, structure and possibilities are all clustered. This might indicate that the metahistorical elements of/observations made by the text are not well integrated with the narrative as a whole.

The most isolated narrative is Byron the Bulb, seen at 9pm on the far left. This makes sense as the text breaks off to relate this parable in complete disconnection from anything else in the novel. Lastly, the term "Rocket" is nowhere near so central to the novel as one would suppose, linguistically.

Gravity's Rainbow has been assessed by many readers as a text that works to generate a feeling that everything may be connected (as a form of conspiratorial plot) and that, therefore, it might equally be the case that nothing is connected. Pynchon terms these paranoia (total connectedness) and anti-paranoia (utter disconnect).

An initial plot of the text in this way allows us to start to consider whether the text constructs particular linguistic and semantic fields around particular parts of the text. Specific terms clearly occur in isolated contexts. The octopus rarely returns; most of Roger's narrative is centred around his pairing and unpairing with Jessica; Byron the Bulb is far out in his own diegetic layer with distinct terms that rarely recur.

Other terms, though, seem scattered across the Zone of the text. As an initial hypothesis that I need to explore much more: it could be that many of the isolated action segments of the text, to which critics have turned their thematic and historical attentions, may share common linguistic cores with many other parts of the text. This might begin to contribute to their ultimate connectedness within the novel.

Following on from this early experiment, **[SLIDE]** I also did some work on authorship attribution in the novels of Thomas Pynchon. For, in mid-2015, Art Winslow caused something of an online furore when he suggested that the pseudonymously-authored novel by “Adrian Jones Pearson”, *Cow Country*, was, in fact, a work by Thomas Pynchon. A full-blown argument then erupted when this was countered by Nate Jones and Pynchon’s own publisher. Indeed, Penguin thundered: “[w]e are Thomas Pynchon’s publisher and this is not a book by Thomas Pynchon”.

To be frank, Winslow’s evidence was slight and bound to irritate fans and critics alike. He argued that the author biography of *Cow Country* pointed to a recluse; someone who didn’t like to be in the limelight, a J.D. Salinger. From this, he leapt to the conclusion that Thomas Pynchon was a likely candidate. The second piece of evidence that Winslow furnished was that the novel is a work of metafiction; another potential Pynchonian connection. Winslow also evaluates the dismantling of binaries within *Cow Country*, the humorous character names that abound throughout (we have “Dr Felch”), apparent running Pynchon-like gags and more. For Winslow, “[t]he off-kilter sensibility one sees in the work of both [authors] would not be [...] easily ‘replicable’ by another”.

Cow Country itself is an amusing-enough read. It’s a campus novel set in a backwater Community College and that documents the trials of the newly arrived educational administrator, Charlie. This hapless individual is set to head the College’s re-accreditation drive, which, to be frank seems a big ask. For instance, the College’s new staff orientation event consists of the castration of a calf as an exercise in team building. You get the idea.

In order to work a little on this problem, I turned to a set of computational authorship attribution methods to examine the stylistic properties of Pynchon’s novels in comparison to *Cow Country*. These methods consist of statistical analyses of the most frequent words used by Pynchon as opposed to Pearson and the order in which they occur. This may sound extremely dry and of little

worth – counting words does not always go down well in literary circles – but, in fact, these techniques have been shown to be highly accurate under specific circumstances. Indeed, at the time of writing, according to Ariel Stolerman, computational forensic stylometry “can identify individuals in sets of 50 authors with better than 90% accuracy, and [can] even [be] scaled to more than 100,000 authors”.

So, when using such a method (in this case Burrows’s Delta), do Thomas Pynchon and Adrian Jones Pearson write in similar ways? I am afraid to say that the answer is a fairly conclusive “no”, as you might have already concluded yourself. I first explored this using the aforementioned method using single words, sets of two-words (bigrams), and sets of three words (trigrams) inside the R programming language’s “Stylo” package. I then clustered these using 30 to 100 words (and all numbers of words in between: e.g. 31 words, 32 words etc.) and taking the most common result in 80% and 90% of cases.

Admittedly, when using only single words, there is a similarity between Pynchon and Pearson that seems alarming. Note that I have also added Don DeLillo’s corpus here in order to provide a foil for the system and so that we can gauge the accuracy of its profiling techniques. At an 80% consensus between all word frequencies, this algorithm believes that *Cow Country* is most similar, linguistically, to *Bleeding Edge*, *Gravity’s Rainbow* and *V*. At 90% there is very little consensus and the algorithm does not correctly cluster any of the Pynchon works as significantly close to one another.

[SLIDE]

Things do, though, get a bit more interesting when we move on to bigrams and trigrams. In the bigram and trigram models – that’s sets of two words in a row that authors share with each other along with the ways in which these tropes are ordered – most of Pynchon’s novels are tightly clustered together, indicating a distinct writing style. The same goes for DeLillo. However, *Cow Country* does not cluster with these other authors. This points towards the likely probability of a different author being responsible for this novel.

[SLIDE]

Indeed, others have speculated, based on hints dropped by Steven Moore, that the actual author of *Cow Country* is one A.J. Perry. Perry is the author of two other books, but lists himself as the author of three on his website. I took one of these books and dropped it into my authorship profiling tool.

Whether using single words, bigrams, or trigrams, I received exactly the same result in this case. Perry's *Twelve Stories of Russia, A Novel I Guess* published in the year 2012, markedly clusters with *Cow Country*. Note also that, once Perry is thrown in, Pynchon and DeLillo are clustered almost entirely accurately and distinctly by this algorithm.

From these hobbyist beginnings, I then began work on my latest book, forthcoming with Stanford University Press next year, which is titled *Close Reading with Computers*. One of the areas that I explore therein is the features of historically imagined style in the contemporary novel. **[SLIDE]** Specifically, I have been working on David Mitchell's *Cloud Atlas*, a novel that spans six different generic styles, ranging from a seafaring journal of the 1850s through to a far-future SF dystopia modelled on *Soylent Green*. When you want to write fiction that sounds as though it was written in the 1850s, though, how do you actually do it? Is mimetic accuracy of the language the crucial point? Mitchell himself thinks not, stating in an interview that "Historical fiction isn't easy; it's not just another genre. How are they going to speak? If you get that too right, it sounds like a pastiche comedy—people are saying 'thou' and 'prithee' and 'gadzooks,' which they did say, but to an early 21st-century audience, it's laughable, even though it's accurate. So you have to design a kind of 'bygone-ese'—it's modern enough for readers not to stumble over it, but it's not so modern that the reader kind of thinks this could be out of *House* or *Friends* or something made for TV".

Given this statement and the linguistic features of the text, I was interested to discover how far mimetic accuracy goes in the Pacific Journal of Adam Ewing in *Cloud Atlas*. The first thing to note is that the novel gives its own internal dating for the Ewing narrative. We are told, by Frobisher, that "[m]ention is made of the gold rush, so I suppose we are in 1849 or 1850". If we take the diary at face value, then Frobisher is almost right. In fact, the year must be precisely 1850, since this is the only year in the 1850s range that has the 7th November (the first dated entry in the diary) falling on

a Thursday. Hence, also, by the internal chronology, when Ewing notes that “[t]oday is [his] thirty-fourth birthday” on Sunday the 12th January 1851, Ewing’s precise birthday is the 12th January 1817. In its tight internal chronology that *does* match the historical record, the text even manages here to parody the (or, my) act of literary interpretation; Frobisher is akin to the paranoid critic who would seek out such information.

However, the first thing to note is that we cannot take the date of the diary at face value. As Frobisher again notes, in a perhaps defensive authorial move for Mitchell, there is “[s]omething shifty about the journal’s authenticity—seems too structured for a genuine diary, and its language doesn’t ring quite true”. Frobisher clearly suspects the entire thing to be a literary forgery; which, of course, it is. Mitchell is the ultimate forger here (although it is by license of the reader), but in the intra-diegetic setup of the text, Jackson Ewing looks likely to have doctored the diary.

The reader knows, from the final pages of the diary, that Jackson Ewing, the son who has “edited” this published diary, was born before Ewing set sail in 1850. However, we are also told that Jackson Ewing is the same age as the first hazing victim aboard the ship: “Rafael was Jackson’s age”. Assuming, then, an approximate earliest birthdate of 1st January 1835 for the late-teenager Jackson Ewing, it seems likely that the furthest date within the text’s internal chronology for editing and publication of the diary, taking an optimistic human lifespan average of 60 years for the time, might reasonably be 1895. The diary would also have to have been edited after Ewing’s return at a later date. If one wanted to be generous, one could extend this by 15 years to 1910, so as to also chime approximately with the date of the diary’s “discovery” by Frobisher in the Letters from Zedelghem section of the novel, a few years later. The date range that this yields for Mitchell’s Ewing chapter is, then, 1851-1910.

By most accounts, Mitchell’s novel is successful at imitating the linguistic style of the period in which it purports to be set. However, the questions that I ask here are: how does Mitchell achieve this? What are the limits of linguistic mimesis in *Cloud Atlas*? And what kind of historical

imaginary could function as a model against which we could measure Mitchell's prose? Specifically, I was interested in the extent to which Mitchell's chapter is accurate in the use of language from its claimed period. That is, I wanted to know whether, of the 13,246 words in *The Pacific Journal of Adam Ewing*, any would have been inaccessible to a writer living between 1851 and 1910.

Assuming that Mitchell's diary object attempts an accurate depiction of language from the time of its purported authorship, an obvious first question is: are there words in the diary whose first usage falls later than the date of the Ewing section of *Cloud Atlas*? In order to gauge the "authenticity" of the diary through the appropriateness of its linguistic register, I made two initial preparatory modifications to the E edition's first section of *The Pacific Journal of Adam Ewing*, the portion to which Frobisher refers [**QUICK NOTE ON VERSION VARIANTS**]. The first pass that I made was to split all words within the text into their own lines and then to eradicate any words that appeared in the Project Gutenberg version of Herman Melville's *Moby-Dick*. Using this text as a filter enabled me to eradicate words that were clearly in use in 1851, the first publication date for Melville's novel. This greatly reduced the effort involved in sequential etymological date checking of these words. The second step was to produce a piece of software that would "scrape" sets of open-access dictionary sites for claimed "first usages" of words and to run the remaining words through that software. The idea behind this was that it should give an indication of any obvious outlier words, which I would then be able to check more thoroughly. [**SLIDE**]

In order to militate against the challenges of etymological research data, I decided to reduce further the terminologies studied (in addition to de-duplication and the *Moby-Dick* filter) to words that appear in Ewing Part I that have etymological data in both the OED and in Dictionary.com. This resulted in a final unique vocabulary of 896 words out of an original 13,246 terms for which I then had two sets of etymological first-use dates.

Taking, then, a latest in-text "publication date" for the "Pacific Journal"'s first section as 1910 yields, in my search of Dictionary.com, just six anachronistic words that would definitively not have been available to either Adam or Jackson Ewing and that occur in both editions of the text:

“home-town” [1910-1915], “spillage” [1920-1925], “lazy-eye” [1935-1940], “returnees” [1940-1945], “latinos” [1945-1950], and “A-frame” [1960-1965]. However, the Oxford English Dictionary disagrees. For “hometown”, it tells us, was coined in 1851; “returnee” in 1870; and “A-frame” as far back as 1827. The OED also yields a number of terms from the novel as being after our cut-off date that Dictionary.com does not. In the OED, through the automatic approach, we are given: “bizarreness” [1920], “spillage” [1934], “slumped” [1937], “pulsed” [1942], “colour” (P) / “color” (E) [1944], and “scuttlebutt” [1945]. There are some strange things going on here that are worth briefly unpacking.

In the case of “bizarreness”, “slumped”, and “pulsed”, the OED API simply disagrees with Dictionary.com, claiming that the specific forms of these words, deriving from older ancestors, were not used until these later points. This is probably because my software is pulling out the incorrect part-of-speech definitions for first usage within the specific contexts. Two words have more interesting stories behind them, though.

“Color” seems an unlikely candidate to have been coined, even in its American spelling, in 1944. Indeed, this is the case. What has actually happened here is that the OED API has taken “color” in the sense of “[a]ny of various musical devices or techniques used to enhance the performance of a piece, esp. a repeated melody in late-medieval isorhythmic motets”; a very specific definition of “color”, with the main entry for perceptions of electromagnetic radiation listed instead under “colour”. This usage of color in the musical sense appears later in *Cloud Atlas* but hardly applies to the initial use here: “a Bonapartist general hiding here under assumed colo[u]rs”.

“Scuttlebutt” also has two different meanings. The older, given by Dictionary.com as first occurring around 1800, means “an open cask of drinking water”. The usage in the text, though, is that “Henry shall inform the ‘scuttlebutt’ that Mr Ewing has a low fever”, meaning in this case a person who puts a rumour about. This second definition as a colloquialism, according to the OED, comes from 1945, while Dictionary.com yields 1905. Interestingly, Mitchell puts this term in quotation marks, as though the speaker is using an informal or new word. Although there is disagreement between my two etymological sources, scuttlebutt is definitely an edge-case here. It is very unlikely that it would have been used in the informal sense during the period of purported

authorship of the document. On the other hand, Dictionary.com does put such a use at 1905, so it makes sense to exclude this from the final definitive list.

This leaves, then, just three terms that, I feel, can be said with certainty to have been absolutely inaccessible either to Mitchell's historic author or the intra-diegetic editor: spillage, from ~1934; latino, from ~1946; and lazy-eye, from ~1960. These are not terms, though, that one would readily identify as anachronistic, particularly since "lazy-eye", for instance, sounds like a politically-incorrect term for the medical condition.

These questions of historical accuracy at the micro level bring to the fore a problem that has vexed historical fiction and its study for many years: to what extent is mimesis to the historical record important for historical fiction? For instance, Harry E. Shaw notes that there are two types of representational phenomenon at play in historical fiction. A work "may represent societies, modes of speech, or events that in very fact existed in the past" or it may "promote some sort of historical effect within the work".

This "historical effect" can be seen in many contemporary takes on fantasy. In Kazuo Ishiguro's remarkable (though not universally well-received) fantasy novel, *The Buried Giant* (2015), for instance, the inhabitants live in a hybrid ancient Saxon world of Arthurian knights, ogres, and dragons. The novel engages, clearly, in a mode of mythopoesis: sowing its other-world within a reader's own past reality, cross-fertilising between a Tolkien-esque universe and ancient Britain. Among the most interesting facets of this novel, however, is the mist of forgetting that covers Ishiguro's land. Indeed, the inhabitants of his story have only the weakest sense of history, a poor historical consciousness, and are unable to remember even recent events that were nonetheless, by all accounts, central to their lives. As the warrior Wistan puts it, "who knows what went on here in ancient days?"

The Buried Giant is clearly not simply a fantasy tale but is also a rumination on history and its mediation. In fact, it is a rather pessimistic meditation on the ability of fiction to reconstruct a past. Like Wistan the reader may hope that, "by travelling beside" Ishiguro's characters, "the memories would awaken". Like Wistan they may too find, though, that "they've not yet done so". Ishiguro's novel closes with an ambiguous warning. As the memory fog lifts, the infidelities and

betrayals of the journeying couple, Axl and Beatrice, are revealed, in parallel to a warning of ancient nationalistic grudges. At the end of the novel it remains unclear whether Axl and Beatrice are allowed to journey together to the heaven-like island or whether these past grudges cause them to fail the ferryman's potentially unpassable test. What is clear is that *The Buried Giant* warns of recovering history, unless one is prepared to forgive its wrongs in the present.

These are the types of question raised by my work on anachronism in contemporary historical fiction, using computational methodologies. Using brute-force, repetitive, machine-driven labour gives us some answers; in fact, some data. But it does not, on its own, get us any closer to understanding works of literature. Indeed, using digital tools still requires a great deal of interpretative effort in order to understand what one has found. For instance, one of the other aspects that I wanted to investigate was whether Mitchell creates a “historical effect” by simply using words that are not as “popular” as those used in a contemporary vocabulary. I quickly knocked up a script that would scrape the Merriam Webster online dictionary for popularity statistics. **[SLIDE]**

Sadly, it turned out that there was a significant problem with this methodology. The underlying question that must first be answered is: what does the Merriam Webster dictionary mean by “popular”? It turns out that the Merriam Webster score for popularity is calculated by the number of times that each word is looked up by users online. In other words, “popularity”, as defined in the Merriam Webster online dictionary, is not taken from any representative corpus of contemporary use, but is determined by how frequently users visit the definition page in question. This, in turn, raises questions as to what “popularity” might actually mean that hinges upon the reasons for which people turn to online dictionaries. By Merriam Webster's measure, “popularity” is actually constituted by a range of socio-behavioural and technological aspects.

For example, when a word scores in the “bottom 10%”, does this mean that: 1.) the word is hardly ever searched for because it occurs so rarely in contemporary usage?; 2.) the word is hardly ever searched for because it is such a common word that everybody already knows how it is used?; 3.) the word is both uncommon in usage but also very well known?; 4.) the word is easy to spell? On the other hand, when a word is scored as being very “popular”, this could mean that: 1.) the

word is extremely rare and not very well known; 2.) the word is very common and is being searched for by non-native speakers of the language; 3.) the word has a subsidiary use that is less familiar; 4.) the word has a specific grid pattern that fits with various word puzzles and is looked up disproportionately by players (“okapi”, for instance, is by far the most common crossword answer for any grid that reads: “o_a_i”); 5.) the word has a difficult spelling and is frequently looked up not because it is uncommon in usage but for its composition; 6.) the word has one or more homonyms and users are seeking to disambiguate the term(s); 7.) the word has an interesting and unapparent etymology; 8.) the word has recently featured in a popular context, giving undue exposure to the term.

In these cases, there are not only matters of literary interpretation with which to contend, but also matters of technological and social interpretation. The accusation that the approaches of the digital humanities reduce us to mindless word counting – surpassing interpretation – is, in fact, a total and pernicious falsehood. New approaches to literary empiricism, using mechanization and digitization, leave us only with thornier questions about how we interrogate texts.

[SLIDE]

The part I don’t have time to cover today concerns terms that Mitchell uses that do not appear in a wide contemporary magazine corpus. In the book, I show how racist, colonial terms form the key marker of difference between Mitchell’s historical fiction and contemporary writing, but I’m going to have to leave you in suspense for that one.

And that is where I here wish to close. I have tried, in this lecture, to paint a picture of how I came to be where I am. Of how I believe the digital realm holds remarkable potential for education and research dissemination and of how I have come to consider the use of digital tools as a promising avenue for the study of literature. I have also charted the way in which the pre-digitality of almost all print must be considered as a precondition of contemporary textuality. I have here attempted to show why I believe that our digital architectures are worth both studying and worth studying with.

[SLIDE]