

Ethical issues in computational pathology

Tom Sorell,¹ Nasir Rajpoot,² Clare Verrill³

¹PAIS, University of Warwick, Coventry, UK

²Computer Science, University of Warwick, Coventry, UK

³Nuffield Department of Surgical Sciences, Oxford University, Oxford, UK

Correspondence to

Prof. Tom Sorell, PAIS, University of Warwick, Coventry, UK; t.e.sorell@warwick.ac.uk

Received 24 October 2020
Revised 2 February 2021
Accepted 5 February 2021
Published Online First
3 March 2021

ABSTRACT

This paper explores ethical issues raised by whole slide image-based computational pathology. After briefly giving examples drawn from some recent literature of advances in this field, we consider some ethical problems it might be thought to pose. These arise from (1) the tension between artificial intelligence (AI) research—with its hunger for more and more data—and the default preference in data ethics and data protection law for the minimisation of personal data collection and processing; (2) the fact that computational pathology lends itself to kinds of data fusion that go against data ethics norms and some norms of biobanking; (3) the fact that AI methods are esoteric and produce results that are sometimes unexplainable (the so-called ‘black box’ problem) and (4) the fact that computational pathology is particularly dependent on scanning technology manufacturers with interests of their own in profit-making from data collection. We shall suggest that most of these issues are resolvable.

Digital pathology is cellular pathology conducted with digital whole slide images (WSIs) rather than tissue sections and light microscopes. The use of WSIs obviates transport and physical sharing of tissue samples, with cost savings and reductions of damage to, or loss of, glass slides. WSIs are generally clear and detailed, even at low levels of magnification and allow rotation, panning and zooming.¹ WSIs can improve clinical workflow. They can aid collaboration: a few experts in different places can work at the same time on analyses of the same slide images for diagnostic and prognostic purposes. In principle, WSIs could even permit expert crowd-sourcing of morphological analyses. They have clear uses for teaching pathology^{2,3} and for routine quality assurance (eg, by UK National External Quality Assurance Service) of pathology practice, without the need to send slides to each of several hundred pathologists during the year. Some of these latter advantages of using WSIs are shared with ‘telepathology’, the longer-established practice of transmitting images from a remote-controlled light microscope for example, for obtaining second opinions.

Digital pathology is not entirely free of ethical issues. For one thing, if it is collaborative, it can involve sharing sensitive personal data, which is subject to distinctive ethical and legal norms. There is also the fact that the scanners used to make WSIs are a new technology only recently permitted for use by regulators in the USA and the UK following large-scale validation studies.^{4,5} The Royal College of Pathologists in the UK found that, by the beginning of 2018, very few sufficiently large studies of the reliability of interpretation using WSIs had been completed, and some of these were not independent

enough of scanner manufacturers to justify full confidence.⁶

Pathology carried out with WSIs is an enabler of computational pathology: artificial intelligence (AI)-aided modelling, analysis and discovery of patterns in large sets of high-resolution and information-rich WSIs. Computational pathology, rather than digital pathology, is the concern of this paper. Crudely, we focus on machine learning applied to WSIs—as opposed to the use by pathologists of WSIs in preference to slides and light microscopes. After briefly giving examples drawn from recent literature of advances in computational pathology, we consider some ethical problems it might be thought to pose. These arise from (1) the tension between AI research—with its hunger for more and more data—and the default preference in data ethics and data protection law (in this paper European Union (EU) and UK data protection law are considered) for the minimisation of personal data collection and processing; (2) the fact that computational pathology lends itself to kinds of data fusion that *prima facie* go against some data ethics norms and some norms of biobanking; (3) the fact that AI methods are esoteric and produce results that are sometimes unexplainable even to experts (the so-called ‘black box’ problem) and (4) the fact that computational pathology is particularly dependent on scanning technology manufacturers with interests of their own in profit-making from data collection.

COMPUTATIONAL PATHOLOGY: SOME EXAMPLES

Computational pathology can quickly classify malignancy (or normalcy) in WSIs. It can be used to predict patient outcome and life-expectancies for different cancer types. It can also identify patterns from the fusion of heterogeneous data, for example, test results of biobanked samples, clinical notes in a natural language, and WSIs of tissue resection or tissue microarrays. Practitioners of digital pathology can often assist the development of computational pathology—for example, with annotations of WSIs and validation of computational algorithms—but computational pathology involves more computational techniques than pathology *per se*.

To enlarge on possibilities of diagnosis and prognosis afforded by computational pathology, Madabhushi and Lee¹³ describe quantitative histomorphometry (QH) analysis, ‘which can now enable a detailed spatial interrogation (eg, capturing nuclear orientation, texture, shape, architecture) of the entire tumour morphologic landscape and its

¹op. cit. note 2.



© Author(s) (or their employer(s)) 2022. No commercial re-use. See rights and permissions. Published by BMJ.

To cite: Sorell T, Rajpoot N, Verrill C. *J Med Ethics* 2022;**48**:278–284.

most invasive elements from a standard H&E slide.’ QH analysis depends on the detection and segmentation of nuclei and glands in images.ⁱⁱ

Madabhushi and Lee go on to mention (1) algorithms for identifying stromal features in images that have been found relevant to prognosis; (2) algorithms and feature approaches for automated tissue classification and disease grading; and (3) histological image-based companion diagnostic tests for predicting disease outcome. There is a distinction between ‘domain inspired’ approaches, that is, approaches geared to specific disease (eg, cancer) types, and ‘domain agnostic’ approaches that cut across several types.ⁱⁱⁱ Certain domain-agnostic approaches use gland shape and size, tissue texture and architecture in prognosis and grading, for example, wavelet and tissue texture features for automated Gleason grading of prostate pathology images. By contrast, ‘[d]omain inspired features ... are typically specific to a particular domain or in some cases to a particular disease or organ site. An example of this class of feature is the co-occurring gland angularity feature presented by Lee *et al*^{iv} which involved computing the entropy of gland directions within local neighbourhoods on tissue sections’.^v

Niazi *et al* mention the possibilities of exploiting not regions but the total area of a tissue section from a WSI. For example,

a whole slide is partitioned into superpixels on the basis of similarity at some magnification. Superpixels are grouped into anatomical regions (specifically epithelium) on the basis of graph clustering... Finally, each cluster is classified as ductal carcinoma in situ or benign or normal on the basis of features extracted by deep learning...^{vi}

And this is not the only instance of work using the whole of WSI content.^{7 8} Finally, the literature sometimes points to the power of AI to unify data drawn from, on the one hand, patient histories, and, on the other hand, heterogeneous types of tissue-archived and biobanked samples.^{vii 9 10}

The benefits of computational pathology, then, can be organised under four headings: improved classification of regions and objects of interest in WSIs; facilitated discovery of patterns correlating tissue architecture with patient outcomes in specific cancers; facilitated discovery of patterns correlating tissue architecture with cancer in general; and the detection of patterns involving tissue architecture and further data not derived from WSIs to predict, for example, life expectancies.

COMPUTATIONAL PATHOLOGY AND PERSONAL DATA

The previous section suggests that, among other things, computational pathology can improve diagnosis and prognosis for patients suffering from cancer and other diseases. Since earlier, more accurate, diagnosis can lead to more timely and more effective treatments, and increase the number of cancer survivors and the length of their lives after diagnosis, computational pathology has clear moral benefits: life-saving and life-lengthening are

among the clearest examples of moral benefits there are—other things being equal.^{viii}

What, if anything, counterbalances these benefits? Research ethics and local law often restrict what can be done with human tissue, and data ethics and law constrain the processing of personal data. WSIs and the pixels that make them up are personal data not in the sense that they always carry explicitly identifying information about whose tissue is imaged, but in the sense—embedded in the European General Data Protection Regulation (GDPR)^{ix 11}—that these identities can be inferred (perhaps using computational techniques), for example, when databases are fused. Again, data ethics and law—we work with the GDPR in this paper—operate with a principle of minimising the collection and processing of personal data¹² and discourage the repurposing of personal data sets. If AI-assisted analytics on WSIs are to identify reliable biomarkers, however, large amounts of data from images—probably the images of tissue derived from a large number of patients over long periods of time—may need to be used to train computational models. Here is where the data hunger of AI runs up against the norm of personal data minimisation from data ethics and data processing law.¹³ Indeed, the hunger in computation pathology for large amounts of personal data is just a special case of the data hunger typical of modern machine learning (particularly, deep learning) algorithms.¹⁴ Again, adding clinical information or longitudinal outcome data to WSI data, though it adds greatly to the potential for clinically useful pattern discovery with the help of AI, sometimes involves repurposing.

At first sight, then, practice in computational pathology seems to flout all or many of the norms governing the use of tissue and personal data derived from it. As in other instances where data ethics seems to tell against an apparently beneficial practice, the key to resolving the tension may lie in distinguishing the case at hand from cases raising a stereotypical risk of a privacy violation or harm on the basis of storage of larger than necessary amounts of personal data. Stereotypical risks of privacy violation occur where data enables inferences about identifiable people’s current health, wealth, sexual practices, political affiliations and friendships. These inferences may allow individuals or organisations to manipulate data subjects or make an economic gain from information about them. If an online gambling site extends credit to people who, according to the data collected, frequently stake large sums and lose, it may be feeding a gambling addiction and further harming the vulnerable. If credit scoring companies with oversimple algorithms unjustifiably count a conscientious but poorly paid saver as a probable defaulter, then, again, the access of the credit scoring company to detailed information about the low-paid person’s income level is morally questionable. These are among the kinds of risk that data ethics and data protection law typically cater for.

Data minimisation requirements also make sense where accumulations of personal data in a single huge data set would significantly add to its attraction as a target of hacking, or as a ‘hostage’ in a ransomware attack, as in the ‘Wannacry’ exploit

ⁱⁱIbid: 171

ⁱⁱⁱIbid: 172

^{iv}op. cit. note 2

^vop. cit. note 7: 172

^{vi}op. cit. note 2: e257

^{vii}The claim that AI can unify heterogeneous data is compatible with saying that the patient benefits from AI-assisted health research are in need of critical evaluation. See 11

^{viii}It is true that life-saving and life-lengthening are benefits relative to the quality of the life saved. Lengthening lives that are already very long or that are oppressive to those leading them are the uncommon cases.

^{ix}“[D]ata subjects are identifiable if they can be directly or indirectly identified, especially by reference to an identifier such as a name, an identification number, location data, an online identifier or one of several special characteristics, which expresses the physical, physiological, genetic, mental, commercial, cultural or social identity of these natural persons.” See

carried out against National Health Service (NHS) computers, among others, in 2017.¹⁵ Differently, a large data set sometimes exposes a pattern of behaviour in individually identifiable data subjects that they consider private and would not want exposed. The well-known case of the US retailer Target, which correctly inferred from the search data of a particular visitor to its website that she was pregnant, is relevant here. This visitor turned out to be a teenager whose father complained about receiving invitations to purchase pregnancy products and had no idea his daughter was pregnant.¹⁶

Could computational pathology pose comparable risks in relation to its data? Here it is important to distinguish between (1) the use of a WSI related to a single patient, that is, a non-computational digital pathology exercise; and (2) patterns disclosed as part of an AI-assisted big data exercise, where possibly thousands of WSIs and other data are used to discover biomarkers for different cancer types. For example, the huge amount of data contained in a WSI might be more revealing than a tissue sample under a microscope. Suppose examination of a WSI revealed an early-stage tumour that would have been missed by an ordinary microscopic examination. In this (a)-type case, could extra data provided by the WSI disadvantage a patient wanting to be paid medical insurance (in countries where commercial medical insurance is the norm) if the insurer claims it is an undeclared condition that predated the policy?

An insurer might dispute liability, but what is normally at issue in such a dispute is not whether a cancer predates a policy, but whether the patient knows of the cancer when taking out a policy, and fails to declare as much. The introduction of WSI-based diagnosis may result in patients knowing earlier about cancers, but it is unclear why it should lead to their feigning ignorance of a cancer when applying for insurance, or being suspected of feigning ignorance by insurers. It is true that insurers might in the future require information regarding an insurance applicant's previous WSI-based diagnoses and prognoses for cancer; but it is unclear that this will lead to more disputed claims or more refusals on the part of insurance companies to offer policies to individuals in the first place. On the contrary, earlier, WSI-assisted diagnosis may make cancers more readily treatable and reduce costs for insurers.

In any case, it is (b)-type cases we are concerned with in this paper. Computational pathology directed at biomarker discovery is a big data exercise often involving very large numbers of WSIs of tissue from many patients. The more aggregated the data, the less personalised and potentially intrusive it is. Again, WSIs for such exercises are often deidentified. Deidentification is a matter of removing explicit links between pathological data and patient identities. This means that that outputs from a big exercise cannot typically be used to identify the data subjects concerned or disadvantage them. It is true that deidentified data is sometimes not anonymised in the strict sense of making all inferences—including computationally assisted inferences—to the identities of data subjects absolutely impossible. Absolutely irreversible de-identification is, if possible at all, very difficult, and might be clinically undesirable, since some of the results of big data analytics might be relevant to ongoing treatment of some of the patients whose tissue was imaged and deidentified. In any case, the techniques that would be needed to turn standardly deidentified data into identifiable data are often extremely sophisticated and expensive to apply, and it is unclear what would motivate the use of such sophistication or large amounts of money to get to the identities associated with a pathology data base, still less one particular identity, by a hacker. So while deidentification may not amount to out and out anonymisation

in the sense of GDPR, it may amount to anonymisation for most practical purposes. Once data are deidentified, then, there are fewer objections to collecting and processing more and more of it for testing and validating algorithms. On the contrary, the larger the data sets used for training and validation, the lower false positive and negative rates are likely to be, other things being equal, with corresponding clinical advantages.

LINKING PATHOLOGICAL WITH BIOBANKED SAMPLES: REPURPOSING AND CONSENT REGIMES

The conclusion of the last section is that the typical rationale for data minimisation does not straightforwardly call into question the big data requirements of computational pathology. We have not, however, seen the last of the tension between data ethics and big data requirements. Consider the following claim from Lewis *et al*:

Repositories containing high quality biospecimens linked with robust and relevant clinical and pathological information are required for the discovery and validation of biomarkers for disease diagnosis, progression and response to treatment. Ready access to such material is fundamental for meaningful translational research. In the case of cancer research, tumour banks have been established to procure fresh as well as formalin-fixed, paraffin-embedded (FFPE) tumour tissues and non-tumour control samples. These tissue collections are increasingly complemented by matched samples of blood, urine, saliva and other bodily fluids where appropriate....

While prospectively targeted collections of appropriately consented human samples are the ideal for translational research programmes, realistically the systematic accumulation of large numbers of samples linked to clinical follow-up, apart from being costly, may take many years to become established. Yet readily available resources for translational research currently exist within many pathology laboratories; indeed, in the surgical pathology archives across the United Kingdom (UK)'s National Healthcare Service (NHS), vast numbers of FFPE tumour and non-tumour control samples are currently stored often untouched for a minimum of thirty years before disposal.¹⁷

Lewis is suggesting a repurposing of samples in pathology archives for whole slide imaging, with a view to connecting analysed WSI data with matched data from biobanked samples (blood, urine) and other related data (including radiological data). Why, if at all, is this problematic morally?

Repurposing of personal data in big data research is in itself morally questionable from the standpoint of research ethics¹⁸ and data law.¹⁹ To enlarge on the research ethics issues, digital data is subject to fusion or analytics often without the knowledge of people who have given consent to its collection. Fusion and analytics are not always subject to formal oversight. Once data are in digital form, there is often confusion as to who should decide about its reuse and which reuses are legitimate. Consents to the use of tissue for research, for example, may reasonably be understood to extend to the use of WSIs of tissue for pathological investigations. The further research possibilities of WSIs, however, may lie in data science rather than medicine. It is unclear whether these indirect uses of tissue are grasped by those giving consent. Again, it is unclear whether the secondary uses of WSI data for algorithm development, including algorithm development for profit, are always a use of data for 'research' envisaged or understood by the cancer patient.

The issues are further complicated, in the UK at least, by the ethics and law of tissue retention. There are two cases, corresponding to the difference between a diagnostic archive and

a postmortem archive. A diagnostic archive is composed of samples taken from the living for a consented medical procedure. These can be used for that procedure and research so long as there is no objection to research from the patient. Samples taken from deceased patients are treated differently. They are covered by the Human Tissue Act 2004, which was in part a reaction to scandals about the retention without permission of organs and tissue, sometimes extracted from children, in at least two English hospitals in the 1980s and 1990s.²⁰ The Act resulted from a public protest not only against the retention of tissue, but the failure to put it to any use in research. Tissue samples in post-mortem archives accordingly require consent for both storage and research.

Do restrictions in the Human Tissue Act inspired by loose hospital practice decades ago still fit public opinion about the use of tissue? There is some evidence that, in the UK at least, attitudes to tissue retention for research have changed. In 2017, a significant consultation exercise took place on future-proofing consent to the use of tissue and health data, sponsored by the Human Tissue Authority (HTA), the regulator associated with the Human Tissue Act.²¹ Participants in general strongly supported a relatively relaxed consent regime to minimise obstacles to health research. They were against the ‘waste’ of already extracted material—its not being used for research—through lack of clarity on consent. They recognised a tension between, on the one hand, giving genuinely informed consent to collection of data on tissue and other biological samples, and, on the other hand, experiencing information overload. At times they worried that inferences might be made from tissue about individual identities and identified people’s lifestyles, but they were reassured by the fact that studies typically used aggregated data and methods of deidentification for tissue or data linked to patient records. Finally, when given the choice between, on the one hand, broad, one-off consent to future research on tissue^x and, on the other hand, dynamic (or periodically renewed and possibly overloading) informed consent, they preferred broad consent. The only ‘red lines’ were the use of tissue by profit-making commercial firms. We come back to this in the final section.

Public support for a broad consent regime as opposed to a ‘dynamic’ one is not by itself a moral justification for that regimen, even when that support is informed by the purposes and methods of biobank-based research or by research involving repositories of science. But if there is a clear health gain, potentially to a population, from tissue research in general, and if broad consent enables more of a gain more quickly without countervailing harm, then that is already something of a moral argument for a broad consent regimen. There may be further arguments for broad consent rather than dynamic consent based, for example, on the way that repeated consenting of biobank donors may create false expectations of personalised gains for donors.²²

Biobanks collect samples from donors to be used longitudinally for research that benefits a wider population from which donors are drawn. These samples are collected with consent from the donors to storage and research. Diagnostic archives of pathology, as already noted, are different: they are part of the medical record. ‘Repositories of science’ in the sense of Lewis

et al bring together not only biobanked samples but those *and* tissue samples in diagnostic and postmortem archives. Granted that the Human Tissue Act strongly discourages the repurposing of pathology archive samples from the deceased, is there a good moral argument against scanning the tissue as a part of a WSI-analytics exercise?

We cannot see that there is. If the scan contributes to training pathologists in tumour recognition or produces images for training an algorithm with the power to make improved diagnosis or grading of tumours, then it makes a contribution to saving lives. It is hard to see how the now dead patient is disrespected or exploited by scanning a donated tissue sample, since scanning is not contrary to a stated preference, or out of keeping with a previously collected consent. Nor is scanning a case of breaking faith with the motivation of the Human Tissue Act. What scandalised people was the storage of tissue and organs, especially the organs of children, without permission and to no clinical or research purpose. The contents of pathological archives are kept with permission and were once put to a clinical purpose when their donors were living. Scanning tissue from these archives to make digital images is not obviously a misuse of tissue, and had the (deidentified) disused organs and tissue samples at the Alder Hey or Bristol Royal Infirmary simply been photographed for research purposes rather than stored, it is not clear than anyone would have been scandalised.

There is a corollary for the assembly of data sets that are different from full scale repositories of science but that promote complementary purposes. A pathology data lake assembles WSIs made from tissue samples of many research centres into a single digital repository suitable for the training of algorithms for diagnosis, prognosis and general biomarker discovery. In the PATHLAKE project,²³ the digital repository brings together deidentified samples of various cancer types from various UK centres, the original tissue having been gathered under a variety of consent regimes. This repository will be open to commercial algorithm development by commercial partners who belong to the PATHLAKE consortium, as well as to others who can be granted access to the data under certain conditions, including payment conditions for commercial applicants. Although a number of moral issues are raised by public–private partnerships in computational pathology (see final section), the repurposing of tissue samples for not-for-profit clinically useful algorithm development seems permissible if the repurposing of the otherwise unused contents of pathology archives in general is permissible. And we have argued that it is.

‘WHOLLY AUTOMATED’ PROCESSES AND EXPLAINABILITY

So far, we have identified a range of tensions between data ethics, research ethics, tissue-use ethics and AI ethics when applied to digital pathology. In this section, AI ethics, data ethics and medical ethics move into the foreground and research ethics slips into the background. The issue to be discussed is the difference made by digital pathology to the reliability and intelligibility of cancer diagnosis and grading. This issue can be sharpened by reference to a norm of AI ethics on the one hand, and a norm of data processing law on the other. The norm of AI ethics is that algorithms ought to be as transparent or as explainable as possible²⁴; the norm of data processing law (GDPR) is that no decision making with significant effects on an individual should be wholly automated.²⁵

At first sight, both of these norms tell against computational pathology in some form. An important output of computational pathology is automated classification of tissue into normal and

^xBroad consent is in the spirit of the new NHS opt out for research with (de-identified) patient data, in that the presumed default position for a patient is broad agreement to research conducted under the protections of data law and Research Ethics Committees.

malignant. Patients, general practitioners, some pathologists and some oncologists will have no idea how a diagnosis generated by an algorithm has been derived, and, if they look at the Best Practice Recommendations of the Royal College, they may think that the College's own attitude towards digital pathology is at best cautiously supportive. These facts are certainly consistent with, and may even support, a norm to the effect that AI-driven diagnosis on its own should not trigger treatment, such as chemotherapy or radiotherapy, with major effects. Of course, it is highly implausible that such treatment would flow automatically from a diagnosis—human or AI-generated—in any case. The treatment would need informed consent. But since even the communication of a cancer diagnosis is often traumatic, and since an AI-generated diagnosis can sometimes be wrong, there may be support for a norm to the effect that an AI-generated diagnosis should be communicated directly to the patient by a doctor who is well informed about the relevant AI and in a position to explain it to some extent. The norm of circumspect communication by people who understand relevant AI may have the odd exception, as when an AI scientist is the subject of a diagnosis, and his or her medical team is familiar with digital pathology, but does it not hold for large groups of patients who know nothing about computers, and for large numbers of medically trained personnel who know little or nothing about AI?

The answer to this question, it seems to us, is 'it depends'. To begin with, what is it for a decision or diagnosis to be 'wholly automated'? WSIs are humanly annotated before algorithms are trained to produce diagnoses and estimates of life expectancy. Machine-generated diagnoses are trained to agree with a set of expert human ones which establish 'ground truth' for the relevant algorithm. Admittedly, deep-learning after supervised machine-learning introduces inscrutability. Furthermore, various machine-learning approaches permit the discovery of patterns between diagnoses and for example, the deep architecture of tissues, that human pathologists do not recognise and are perhaps incapable of recognising, patterns that unexpectedly track the presence of tumours. That does not mean that the machine is wholly unconstrained by judgements of human pathologists.²⁶

The cases that most strongly support the GDPR norm against automated decision-making are ones in which administrative decisions—decisions to investigate, distribute benefits or assign penalties—are subject to requirements of impartiality, consistency and proportionality. Here, the norm against automatic decision-making operates against the injection into an algorithm of personal bias or disproportionality. But making diagnoses in pathology is not like this. Although it is possible for algorithms to be biased by an insufficiently varied training set, this is not likely to be the result of the influence of the stereotyping that can blight administrative decisions.

Let us turn now to the demand for 'explainability' in algorithms. This, too, makes most sense in relation to automated administrative decisions, for example, lending decisions based on automated credit scoring; automated sentencing decisions, and some decisions about prioritising the deployment of police in response to calls for help. Digital pathology does not lend itself to the norm of explainability in algorithms if explainability is a safeguard against arbitrariness and unfairness. Arbitrariness and unfairness are a matter of how the human will is directed and what considerations are given weight in decisions that are made primarily by humans. On the other hand, whether someone gets a cancer diagnosis, or a higher or lower grading of cancer, based on WSI data, is not a matter of arbitrary human decision. It is a matter of what conditions of tissue are, independently of

anyone's will, markers of tumour development. How algorithms track those markers is a separate issue, and even if it is unexplainable in cases of deep learning, its being unexplainable is compatible with a low error rate in diagnosis and grading, and, in particular, an error rate lower than that of human diagnosis and grading. The relatively low error rate is morally important when the alternative to AI-derived diagnosis and grading would be human diagnosis and grading.

Another morally important consideration is speed. The speedier diagnosis and grading are with better or comparable error rates, the more they encourage earlier and effective treatment. This is a large part of the moral argument for relying on computational pathology. It is true that the difficulties for doctors not well versed in AI of communicating the basis for the accuracy of AI-driven diagnosis, and the difficulty for patients of understanding AI, hamper informed consent. But this is not a decisive consideration against relying on computational pathology; it is an argument instead for training doctors in AI so that are not mere mouthpieces for algorithms.²⁷⁻²⁹

Citizen jury work on explainability of algorithms shows that, for patients, it is the effectiveness, rather than the intelligibility, of the algorithm that matters in medical contexts.³⁰ This means that even if the AI behind an algorithm is inscrutable to specialists—a case of the 'black box' problem—that too may be regarded as secondary under citizen jury conditions, that is, under conditions of relatively full information about the significance of black box problems in practice.³¹

Of course, computational pathology *does* lend itself to the medical ethics norm of informed consent for biopsies that will lead to algorithm-based diagnosis. And the medical ethics norm of informed consent does require that, as far as reasonably possible, the patient know what the biopsy is for and how reliable a diagnosis will be. But meeting this norm when the diagnosis is AI-driven does not translate into a demand that the patient be given an introductory course on AI. Neither does informed consent to a procedure involving an X-ray involve an introduction to either radiography or radiology. What matters is to be given an accurate and understandable account of the ill effects of a small dose of radiation compared with the benefits to a choice of treatment after access to X-ray imagery. The comparable information about automated diagnosis might be information about 'concordance' and 'discordance' rates between digital pathology and pathology with light microscopy,³² and general information about pattern discovery in WSIs.

COMPUTATIONAL PATHOLOGY AND COMMERCIAL INTERESTS

We come finally to the ethical issues arising from the role of commercial firms in digital and computational pathology. These issues have a distinctive character when commercial firms use data from a public health service, such as the UK NHS, with unique and valuable data sets, and where the consequences of misuse might be particularly damaging to an institution at the heart of a national welfare state. (We concentrate on the UK context, and will ignore the contrasting issues that arise in jurisdictions dominated by private healthcare, such as the USA.)

Scanner manufacturers are a leading type of commercial participant in digital pathology. Their equipment produces WSIs,

²⁸Schiff and Borenstein make the good point that many people in different roles share responsibility for making algorithmic results in medicine intelligible. Not only doctors but coders and others are or ought to be involved. For general discussion about AI and medicine, see ^{28 29}

and the more widely distributed it is in hospitals or laboratories, the more money they make. Scanner manufacturers, then, have an interest in the growth of digital and computational pathology quite apart from the gains to patients, and they have paid for some of the studies that compare the accuracy of diagnoses based on WSIs to diagnoses using microscopes and slides.^{xii}

Patients, for their part, are sometimes suspicious of for-profit uses of tissue and data. According to the HTA report of the public dialogue on biobanking data that we referred to earlier,

Participants' most common red lines were no access for commercial companies like insurance companies or marketing companies using data to sell a product.^{xiii}

The same red lines are reflected in the UK NHS Code of Conduct for data-driven research.³³ Principle 10 is directed specifically to for-profit technology developers and researchers. To this audience the code says, 'Define the commercial strategy'. And the code spells out what this means. Among other things, commercial activity has to conform to a Framework introduced in July 2019.³⁴ This restricts purposes that can be pursued with patient data to those that benefit health, and it asks that authorities in charge of health data be aware of the commercial value of data sets, prohibit exclusivity in access to it for commercial partners, audit use of it, and communicate arrangements and practices to data subjects and the wider public.

In keeping with a Framework that insists on openness with the public in general and patients in particular, consultations or 'dialogues' involving a wide range of stakeholders in data-driven research will probably be relied on to inform its application. Such a dialogue was recently conducted by the UK Academy of Medical Sciences.³⁵ The Academy study found 'universal support'

for data-driven technologies which are based on scans and imaging automation for diagnosis. Data collected and used in this way for direct clinical care was accepted by all participants; these new technologies were enthusiastically welcomed especially by healthcare professionals. There was support for outcomes from machine learning being used to support shared decision making.^{xiv}

This endorsement clearly includes digital and computational pathology. Patients, healthcare workers and non-affiliated members of the public seem to support this technology whether or not it depends on commercial manufacturers of scanners. Participants in the Academy study appear not to have been asked whether the use of images for developing commercially saleable algorithms was supported where this also advanced diagnosis. But in both the Academy and HTA studies it was non-health uses of data—for marketing or insurance—that seemed to be most disapproved of.

Does the strong public support for imaging research and data use help to legitimise commercial activity in this area, when that activity also conforms to the exacting 2019 Framework? We believe the answer is 'Yes, for practical purposes'. After all, the public dialogue approach fits in with the democratic principle that those affected by policy and practice should have a say in it, and public dialogues insure informed support by using the techniques of citizen juries. Experts are able to communicate to the

public relevant facts about the relevant science and the groups involved in research.

The question left open by the dialogue approach is whether the interests of, for example, scanner manufacturers are adequately represented. Are the manufacturers and other commercial interests able to participate, or to be heard in deliberations leading to a Framework, and, if so, how? In some jurisdictions, the UK, the USA and Canada included, academic research is sometimes geared to partnerships between commercial actors, academics and public sector bodies, with commercial partners being expected to make 'in-kind' contributions, sometimes through the donation of equipment or staff time to joint projects, alongside grants from government. Digital and computational pathology are being pursued this way in Britain, under research organised and funded by Innovate UK. The terms of that co-operation are as much in need of a multiparty dialogue as the use of NHS data, for at least two reasons. First, not only data subjects but academic researchers are liable to be at a disadvantage in contract negotiations over rights to the proceeds of joint research. Second, the estimate of in-kind contributions in dollars and cents or pounds and pence is deeply contentious. In this case the multiparty dialogue cannot just involve the public and healthcare professionals. Commercial partners are clearly stakeholders.

It might be thought that since commercial firms and industry bodies have channels of their own for making representations to government, the need for them to be included in a dialogue that gauges public opinion for data use or that determines a framework for conducting data-driven research is correspondingly slight. Again, it might be thought that when scanner manufacturers belong to corporations with global reach and resources for influencing legislation, their being included in dialogues adds to an already disproportionate influence. Our own view is that at least their reaction to dialogues that exclude them should be taken into account by governments and regulators, if not the public. But their views need not disagree with those of the public in relation to every kind or use of health data. It is perfectly possible that digital and computational pathology are unalloyed goods from quite a number of points of view, including those of scanner manufacturers and patients.

CONCLUSION

The preceding discussion has identified and outlined approaches to resolving certain ethical issues in computational pathology. In particular, certain tensions between computational pathology, data ethics and tissue ethics have been addressed. Repurposing of tissue use for scanning and research seems highly justifiable. Ethical issues in the handling and storage of tissue for research seem to be tangential to research with WSIs, since WSIs do not seem to alter or damage tissue, and since research with WSIs also avoids the 'waste' for research purposes of tissue in pathology archives. Again, neither data hunger nor fully automated decision making in computational pathology seem to carry the risks that personal data minimisation and other principles are typically intended to counteract. Demands for data minimisation and the ban on automated decision making seem to be prompted by problems of arbitrariness in the application of rules that do have clear counterparts in cancer diagnosis, prognosis and grading. The business ethics of commercial firms in computational pathology is another sort of issue, and one that has been anticipated early by codes of conduct in the UK. It is too soon to say whether Principle 10 of the NHS code is adequate for resolving this issue, but it is also too soon to say that it is not.

^{xii} *Ibid.*

^{xiii} *op. cit.* note 26, executive summary

^{xiv} *Ibid.*

Contributors TS will act as guarantor. He satisfies all of the BMJ criteria for authorship. NR satisfies all those criteria but has been particularly involved with section I. CV satisfies all the criteria, and has been particularly involved with section III.

Funding This study was funded by Innovate UK (Grant number: 18181).

Competing interests All authors are members of the Innovate-UK funded Pathlake Centre of Excellence consortium. The PathLAKE (Pathology image data Lake for Analytics, Knowledge and Education) is a cross-faculty research consortium comprising researchers from the University of Warwick, University Hospitals Coventry and Warwickshire NHS Trust, and Royal Philips to create a national centre of excellence in AI in pathology, linked to five digitised NHS pathology labs. The cutting-edge AI technologies will assist pathologists in diagnosing cancer more efficiently and selecting the optimal treatment for cancer patients.

Patient consent for publication Not required.

Provenance and peer review Not commissioned; externally peer reviewed.

REFERENCES

- Bradley A, Jacobsen M. Toxicologic pathology Forum*: opinion on considerations for the use of whole slide images in GLP pathology peer review. *Toxicol Pathol* 2019;47(2):100–7.
- Niazi MKK, Parwani AV, Gurcan MN. Digital pathology and artificial intelligence. *Lancet Oncol* 2019;20(5):e253–61.
- Madabhushi A, Lee G. Image analysis and machine learning in digital pathology: challenges and opportunities. *Med Image Anal* 2016;33:170–5.
- Snead DRJ, Tsang Y-W, Meskiri A, et al. Validation of digital pathology imaging for primary histopathological diagnosis. *Histopathology* 2016;68(7):1063–72.
- Mukhopadhyay S, Feldman MD, Abels E, et al. Whole slide imaging versus microscopy for primary diagnosis in surgical pathology: a multicenter blinded randomized noninferiority study of 1992 cases (pivotal study). *Am J Surg Pathol* 2018;42(1):39–52.
- Cross S, Furness P, Igalı L, et al. Best practice recommendations for implementing digital pathology. Royal College of pathologists. §§ 2.4 and 2.5, 2018. Available: <https://www.rcpath.org/uploads/assets/f465d1b3-797b-4297-b7fedc00b4d77e51/Best-practice-recommendations-for-implementing-digital-pathology.pdf>
- Lu W, Graham S, Bilal M. Capturing cellular topology in multi-gigapixel pathology images. Wenqi. *The IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, 2020:260–1.
- Javed S, Mahmood A, Fraz MM, et al. Cellular community detection for tissue phenotyping in colorectal cancer histology images. *Med Image Anal* 2020;63:101696.
- Lewis C, McQuaid S, Hamilton PW, et al. Building a 'Repository of Science': The importance of integrating biobanks within molecular pathology programmes. *Eur J Cancer* 2016;67:191–9.
- Vollmer S, Mateen BA, Bohner G, et al. Machine learning and artificial intelligence research for patient benefit: 20 critical questions on transparency, replicability, ethics, and effectiveness. *BMJ* 2020;368.
- Personal data. Available: [https://gdpr-info.eu/issues/personal-data/#:~:text=The%20term%20'personal%20data'%20is,Data%20Protection%20Regulation%20\(GDPR\).&text=Personal%20data%20are%20any%20information,identified%20or%20identifiable%20natural%20person](https://gdpr-info.eu/issues/personal-data/#:~:text=The%20term%20'personal%20data'%20is,Data%20Protection%20Regulation%20(GDPR).&text=Personal%20data%20are%20any%20information,identified%20or%20identifiable%20natural%20person)
- Principles (C): data minimisation. Available: [https://ico.org.uk/for-organisations/guide-to-data-protection/guide-to-the-general-data-protection-regulation-gdpr/principles/data-minimisation/#:~:text=What%20is%20the%20data%20minimisation%20principle%3F,-Article%205\(1&text=Also%20bear%20in%20mind%20that,under%20the%20right%20to%20rectification](https://ico.org.uk/for-organisations/guide-to-data-protection/guide-to-the-general-data-protection-regulation-gdpr/principles/data-minimisation/#:~:text=What%20is%20the%20data%20minimisation%20principle%3F,-Article%205(1&text=Also%20bear%20in%20mind%20that,under%20the%20right%20to%20rectification)
- Information Commissioner's Office. Available: <https://ico.org.uk/for-organisations/guide-to-data-protection/guide-to-the-general-data-protection-regulation-gdpr/principles/data-minimisation/>
- Brownlee J. How much training data is required for machine learning? 2019. Available: <https://machinelearningmastery.com/much-training-data-required-machine-learning/>
- WannaCry cyber attack cost the NHS £92m as 19,000 appointments cancelled. Available: <https://www.telegraph.co.uk/technology/2018/10/11/wannacry-cyber-attack-cost-nhs-92m-19000-appointments-cancelled/>
- How target Figured out a teen girl was pregnant before her father did. Available: <https://www.forbes.com/sites/kashmirhill/2012/02/16/how-target-figured-out-a-teen-girl-was-pregnant-before-her-father-did/>
- Lewis C, McQuaid S, Hamilton PW, et al. Building a 'Repository of Science': The importance of integrating biobanks within molecular pathology programmes. *Eur J Cancer* 2016;67:191–9.
- Clark K, Duckham M, Guillemin M, et al. Advancing the ethical use of digital data in human research: challenges and strategies to promote ethical practice. *Ethics Inf Technol* 2019;21(1):59–73.
- Steinmann M, Matei SA, Collmann J. A theoretical framework for ethical reflection in big data research. In: Collmann J, Matei S, eds. *Ethical Reasoning in big data: computational social sciences*. Switzerland: Springer, 2016.
- Burton JL, Wells M. The alder Hey affair. *Arch Dis Child* 2002;86(1):4–7.
- Human Tissue Authority. Public dialogue on consent to use human tissue and linked health data in health research: one year on, 2019. Available: <https://www.hta.gov.uk/news/public-dialogue-consent-use-human-tissue-and-linked-health-data-health-research-one-year>
- Steinsbekk KS, Kåre Myskjå B, Solberg B. Broad consent versus dynamic consent in Biobank research: is passive participation an ethical problem? *Eur J Hum Genet* 2013;21(9):897–902.
- Colling R, Pitman H, Oien K, et al. Artificial intelligence in digital pathology: a roadmap to routine use in clinical practice. *J Pathol* 2019;249(2):143–50.
- Floridi L, Cows J, Beltracchi M, et al. AI4People-An ethical framework for a good AI Society: opportunities, risks, principles, and recommendations. *Minds Mach* 2018;28(4):689–707.
- Information Commissioner's Office. Rights related to automated decision making including profiling. Available: <https://ico.org.uk/for-organisations/guide-to-data-protection/guide-to-the-general-data-protection-regulation-gdpr/individual-rights/rights-related-to-automated-decision-making-including-profiling/>
- Topol EJ. High-Performance medicine: the convergence of human and artificial intelligence. *Nat Med* 2019;25(1):44–56.
- Schiff D, Borenstein J. How should clinicians communicate with patients about the roles of artificially intelligent team members? *AMA J Ethics* 2019;21(2):E138–45.
- Char DS, Shah NH, Magnus D. Implementing Machine Learning in Health Care - Addressing Ethical Challenges. *N Engl J Med* 2018;378(11):981–3.
- Obermeyer Z, Emanuel EJ. Predicting the Future - Big Data, Machine Learning, and Clinical Medicine. *N Engl J Med* 2016;375(13):1216–9.
- National Institute of Health Research Greater Manchester Patient Safety Translational Research Centre. Citizens' Juries on Explainable Artificial Intelligence, 2019. Available: <http://www.patientsafety.manchester.ac.uk/research/themes/safety-informatics/citizens-juries/>
- Castelvecchi D. Can we open the black box of AI? *Nature* 2016;538(7623):20–3.
- Cross S, Furness P, Igalı L, et al. Best practice recommendations for implementing digital pathology. Royal College of pathologists. p.5, 2018. Available: <https://www.rcpath.org/uploads/assets/f465d1b3-797b-4297-b7fedc00b4d77e51/Best-practice-recommendations-for-implementing-digital-pathology.pdf>
- Department of Health and Social Care. Code of conduct for data-driven health and care technology, 2019. Available: <https://www.gov.uk/government/publications/code-of-conduct-for-data-driven-health-and-care-technology/initial-code-of-conduct-for-data-driven-health-and-care-technology#principle-10-define-the-commercial-strategy>
- Department of Health and Social Care. Creating the right framework to realise the benefits for patients and the NHS where data underpins innovation, 2019. Available: <https://www.gov.uk/government/publications/creating-the-right-framework-to-realise-the-benefits-of-health-data/creating-the-right-framework-to-realise-the-benefits-for-patients-and-the-nhs-where-data-underpins-innovation>
- Castell, S., Robinson, L. & Ashford, H. Future data-driven technologies and the implications for use of patient data, ipsos mori, 2018. Available: <https://acmedsci.ac.uk/file-download/6616969>