Pharmaceutical patents and data exclusivity in an age of AI-driven drug discovery and development

Medicines Law & Policy

Christopher Garrison

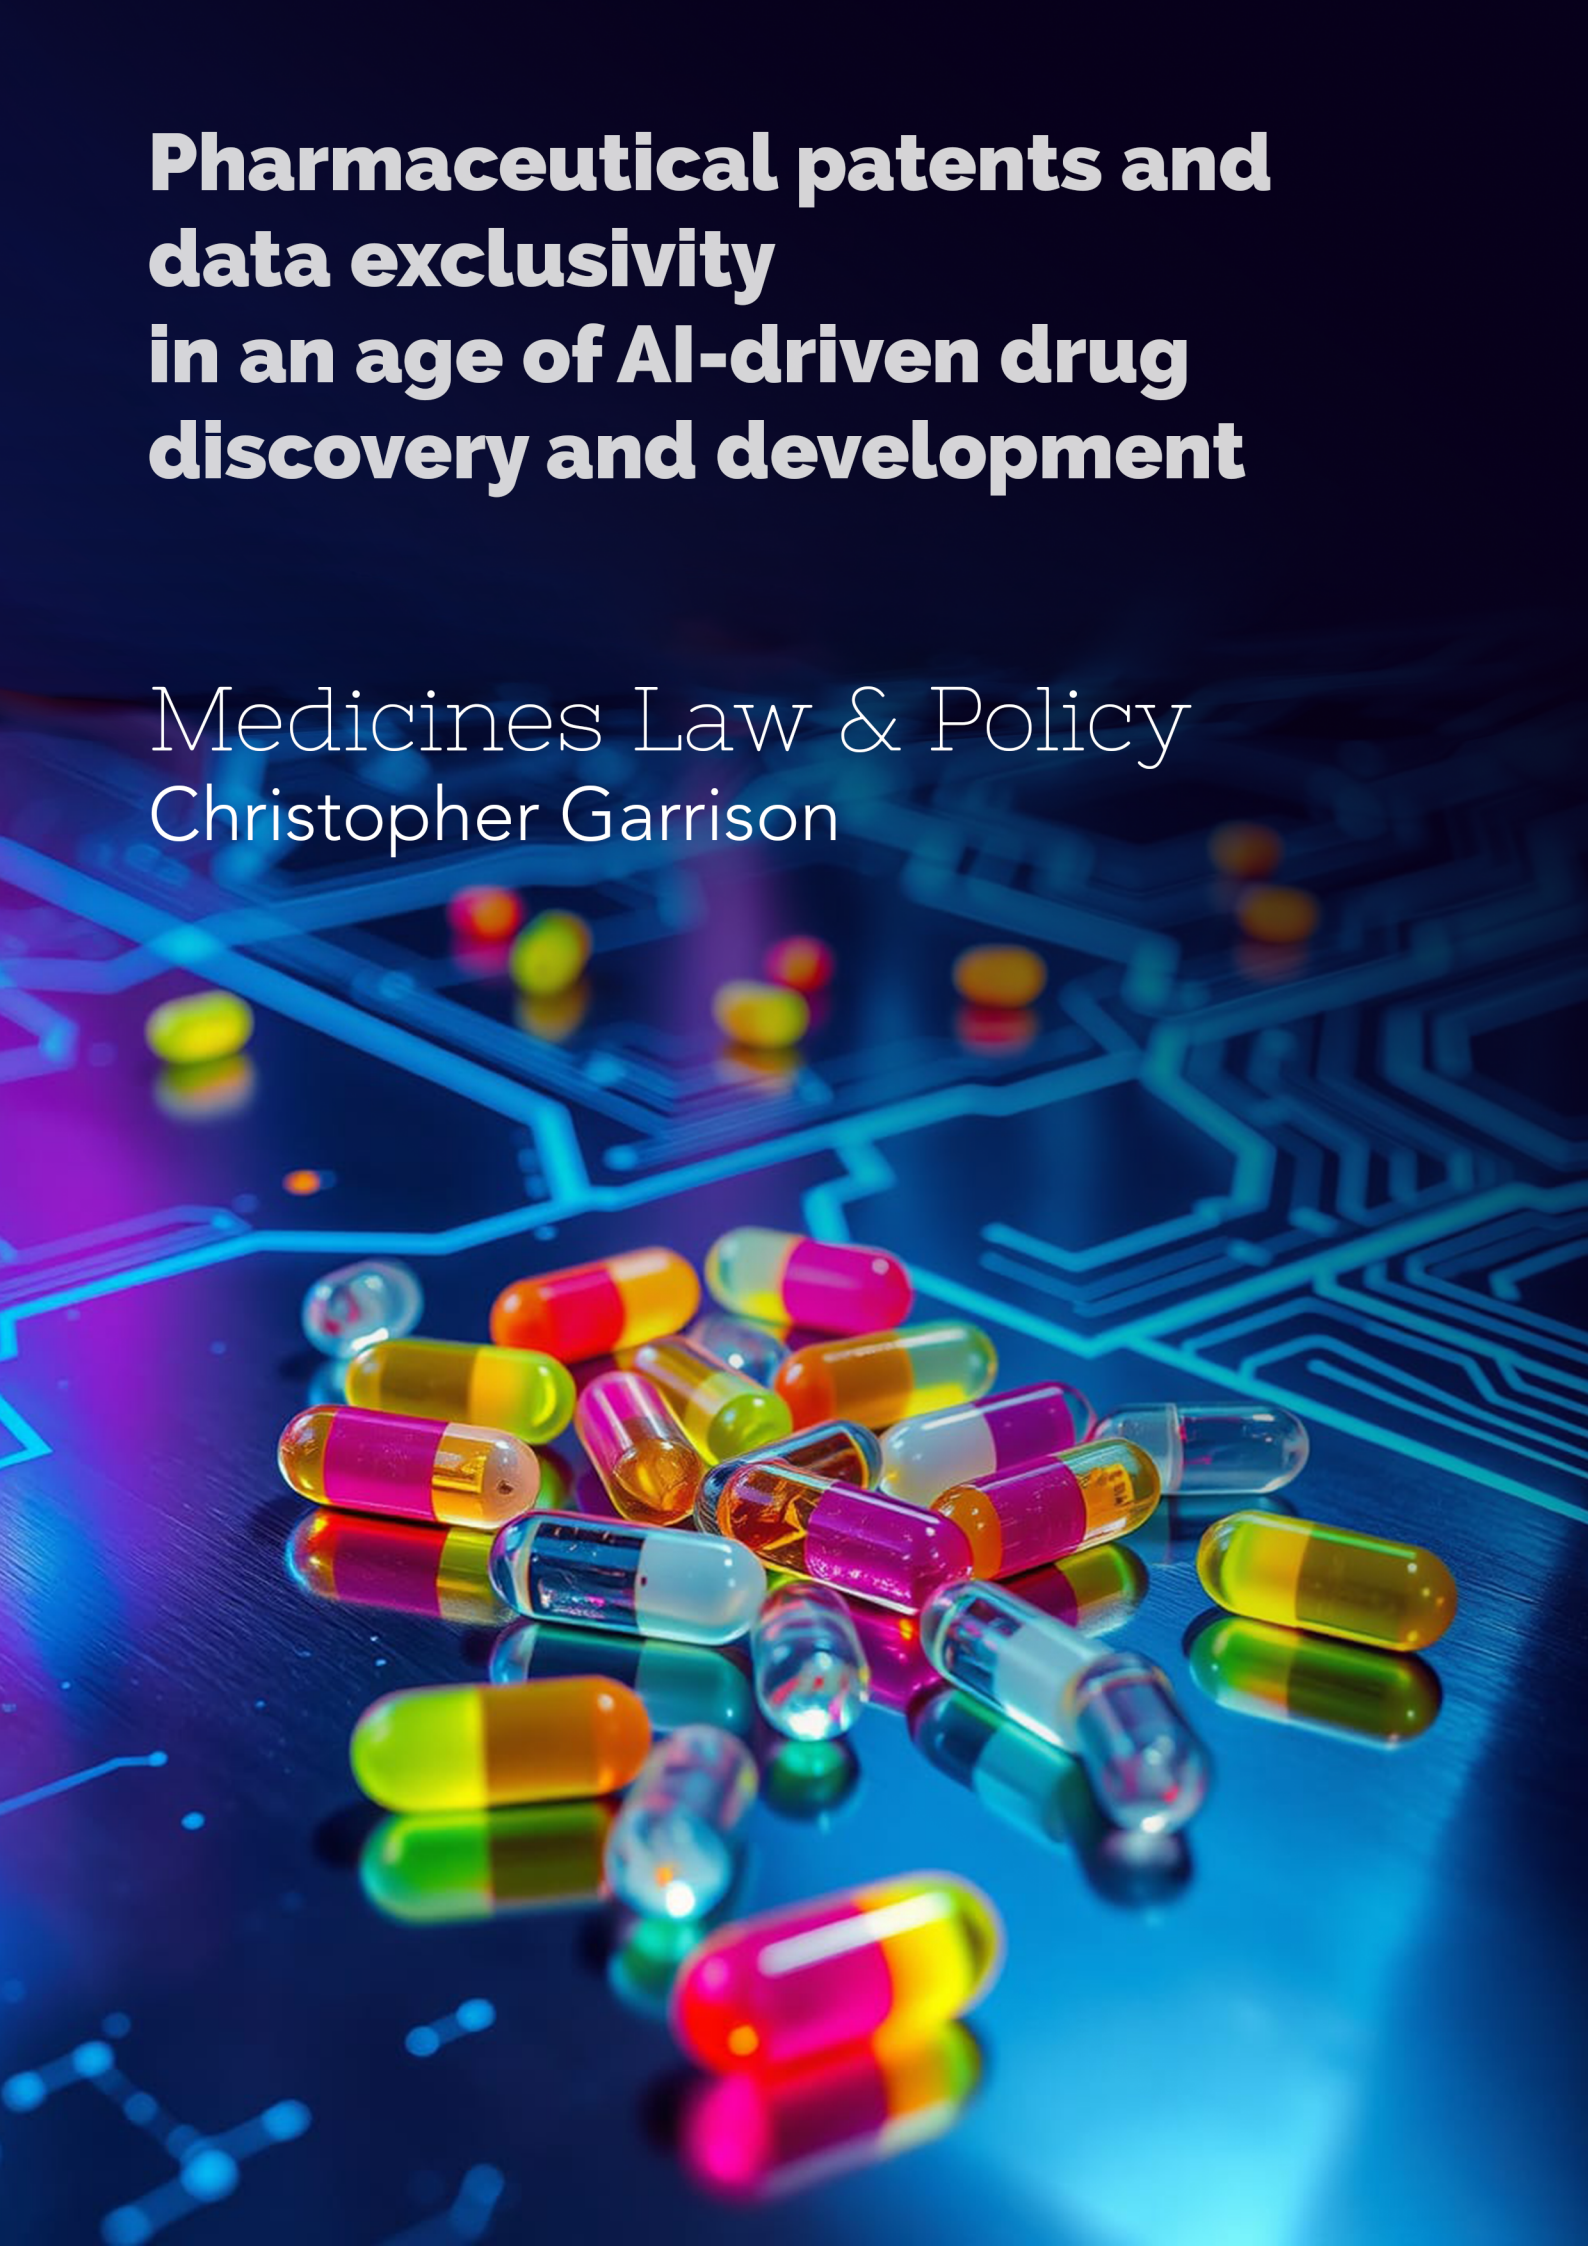

Pharmaceutical patents and data exclusivity in an age of AI-driven drug discovery and development

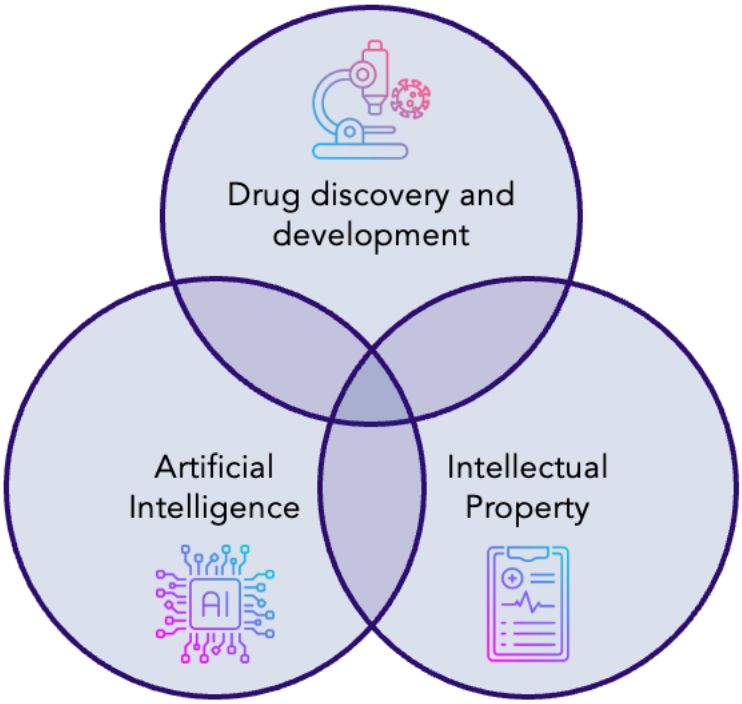

1. Introduction.

Recent advances in Artificial Intelligence (AI) are beginning to have an impact on drug discovery and development. Although it is perhaps too soon to say how transformative that impact will be, there are credible reasons to believe that it could lead to a reduction in the cost, risk and time taken for drug discovery and development. This is important since the pharmaceutical industry has argued for many decades that countries must provide increasing levels of patent and other intellectual property (IP) protection to keep pace with the increasing costs, risks and time taken to discover and develop new drugs. Their success in persuading policy makers of this argument has been reflected in the ever-increasing levels of pharmaceutical IP protection around the world, part of a process which has memorably been described by Peter Drahos in terms of an 'IP-ratchet'. What, then, if the introduction of AI technology is able to significantly reduce the cost, risk and time taken for drug discovery and development? By the same argument, should levels of pharmaceutical IP protection begin to be scaled back? Might the IP-ratchet have a reverse direction after all?

Alternatively, AI technology might take us in a rather different direction. What if it proves to be so powerful that it upends our thinking about the patent system which, after all, stretches back in recognisable form to at least the English Statute of Monopolies (1624), if

not before. What if AI systems become essentially able to make any number of inventions on demand? The founder of a new AI-firm involved in drug discovery recently said:

> "In a perfect world, the dream of purely computational drug discovery comes to life," he described. "You tell me, this is the sort of disease you're going after. Maybe there's a target protein that you have in mind. And at a push of a button, we can generate candidates that meet all of the criteria you care about, which de-risks the steps that come after, shaves time off the process, and at the end of the day, will yield better drugs, faster." ("Ex-DeepMind scientist launches AI drug discovery venture", Financial Times, 13 February 2025)

Such a dream does hold out the hope of designing better drugs that could pass through the drug development process at lower cost, risk and time taken than ever before, which would be wonderful news. However, will the patent system regard an invention made by an AI system, without any human input, as a patentable one? It currently does not. Without a change in this human-centric approach, what will therefore likely happen to the pharmaceutical industry business model in this brave new world of AI-driven drug development and discovery?

This discussion paper aims to explore these and related questions. It focusses on a subset of pharmaceutical IP: patents, regulatory patent extensions and data exclusivity. Other IP rights, including copyright and trade secrets are not therefore addressed except, in passing, where relevant. Sections 2, 3 and 4 of this paper review recent developments regarding the use of AI in drug discovery (and ancillary developments in patent offices) and discuss their consequences for the patentability of pharmaceutical inventions. Sections 5 and 6 of this paper review recent developments regarding the use of AI in drug development and discuss their consequences for regulatory patent extensions and data exclusivity.

Section 7 frames the issues raised in the preceding sections in the context of the overall incentive environment. The paper aims to show that the use of AI in drug discovery and development not only raises 'nuts and bolts' issues for the patent system, as well as for data exclusivity, but also arguably provides a rare opportunity to go back to basics in thinking about them as incentives: are they likely to remain 'fit for purpose' in encouraging pharmaceutical innovation as the twenty-first century progresses, or might the time have come to think about alternatives? Section 8 concludes with a look to the horizon and the possibility that, if Artificial General Intelligence (AGI) or Artificial Super Intelligence (ASI) is created in the next few years or decades, for better or worse, we might not have to worry about the patent system or any other innovation incentive systems anymore at all.

It is clearly impossible to be confident about the precise direction that AI is likely to take, even in just the next few years. AI has become a technology of vital strategic significance for the US, China, Russia, Europe and others, and balancing speeding the development of new capabilities against the danger of doing so – for example, whether the use of AI to develop new biological or other weapons or the development of super-human AI which is not aligned with human interests – may well prove very difficult. Assuming, nevertheless, a

manageable policy environment, this paper suggests the following hypothetical future scenarios for consideration as a starting point for further reflection and discussion:

- The development of AI systems by AI firms and/or pharmaceutical companies could bring about a significant reduction in the cost, risk and time taken for drug discovery and drug development. The most immediate 'low hanging fruit' could lie in speeding up clinical trials. Regulatory patent extensions would reduce in term although the term of the base patent and data exclusivity would remain the same. The profitability of the pharmaceutical industry business model could well increase.

- Although an increase in profitability would further stimulate the pharmaceutical industry business model, many dissatisfied with what that model delivers would likely argue that the time had come to scale-back the patent and data exclusivity incentives – perhaps through shortening them and/or converting them from an exclusivity regime to a licence of right (liability) regime – as they would have become overly favourable to industry at the expense of the public.

- In a rather different direction, the development of AI systems by AI firms and/or pharmaceutical companies could instead lead to new drug-related inventions in several important areas becoming unpatentable, for example, if the AI system is itself deemed to be the proper inventor, contrary to the currently widely held belief that only humans can be inventors for the purposes of the patent system, or if proprietary or open-source AI systems become so widespread that the new drug-related inventions become typically regarded as obvious. If excluded to one extent or another from the patent system, the pharmaceutical industry business model in those areas could become threatened.

- In response to the latter threat, amendments to national or regional patent laws could be made, for example, to permit AI systems to be properly named as inventors or to amend the test of obviousness to include broader secondary conceptions of inventive step.

- However, instead of artificially straining patent laws so far from historical understandings, it may be better to look to a new framework to encourage innovation in these areas, more suited to the realities of AI systems. Such frameworks can be thought of as part of a broader debate about moving from thinking about protection based on a difficult-to-judge level of invention toward protection more transparently based on levels of investment. Opportunities could also flourish for not-for-profit drug discovery and development with access to AI systems and contract-based drug research and development firms.

Given these hypothetical future scenarios, it seems sensible to begin a closer monitoring of relevant developments regarding the use of AI in drug discovery and development as well as its consequences for the patent and data exclusivity systems.

N.B. Although the term 'Machine Learning' (ML) is in many ways preferable to that of 'Artificial Intelligence' (AI), this paper will use AI instead of ML due to its more widespread current usage. Although the pharmaceutical industry sector includes a variety of different actors with different interests, for example, single product start-ups, 'Big Pharma' and generic drug manufacturers, this paper will use the term 'pharmaceutical industry' to describe those firms undertaking R&D activity which make use of the typical for-profit patent and data exclusivity model. The author is grateful for insightful comments on an earlier draft of this paper to Sean Flynn, Edward Griffen, Olga Gurgula, Mina Hosseini, James Love, Rohit Malpani and Anthony Taubman, as well as to his colleagues at Medicines Law & Policy, Pascale Boulet, Montgomery Dunn, Ellen 't Hoen, Kaitlin Mara and Katrina Perehudoff. Any errors are nevertheless the authors' own.

2. The use of AI systems in drug discovery

2.1 AI systems developed by AI firms and/or pharmaceutical firms

This section reviews some recent developments in the use of AI systems in the life sciences which impact drug discovery, focussing on the prominent example of Google DeepMind ('DeepMind') and its AlphaFold AI system. Drug discovery is the process of identifying new molecules (or new uses for known molecules) that could be used to treat a given human disease-related condition. The review tries to reflect some of the sense of surprise at the speed of change that is occurring at the frontier of this field, which is relevant to thinking about whether the policy environment will be able to keep up.

DeepMind was founded by Demis Hassabis, Mustafa Suleyman and Shane Legg in London in 2010 and was acquired by Google (Alphabet group) for approximately $0.5 bn in 2014. The initial focus of the researchers at DeepMind was to develop novel AI systems which could play a variety of games including video games and complex board games. These AI systems differed from earlier approaches which relied on explicitly human-coded rules, for example, the 'Deep Blue' system created by IBM which defeated the world champion Chess player, Garry Kasparov, in 1997. Instead, the first system DeepMind built, 'AlphaGo', learned how to play the Chinese board game Go by being provided with a large set of previous played human games. More impressively, its successor, 'AlphaZero', was provided with only the rules of Go and learned how to play by competing against itself millions of times over, in only hours or days of elapsed human time. The game playing complexity of Go far exceeds that of Chess, and it had been confidently expected in the Go community that no AI system would beat a human Go champion any time soon. In fact, AlphaGo defeated Lee Se-dol, the leading human Go champion, by four games to one in 2016, causing him to retire shortly afterwards. AlphaZero thereafter defeated AlphaGo by one hundred games to zero.

Learning to play complex board games was just a prelude to approaching more difficult real-world problems. Hassabis and Legg had prior experience in computational biology and were interested in one of science's 'grand challenges': the protein folding problem. Proteins are key subjects of study in biology and medicine due to their vitally important role in cell operation and, hence, in human metabolism. They are involved in a wide variety of

structural, energetic, recycling, regulatory (gene expression), transport and signalling processes. Proteins are formed from sequences of the order of one hundred to one thousand amino acids. These sequences fold themselves in milliseconds into complex three-dimensional shapes which, in turn, determine their functionality (Figure 1). Many drug molecules function through interaction with this three-dimensional protein shape by 'docking' with a binding site on the protein (a process often compared to fitting a key in a lock), thereby either enhancing or blocking its functionality. Predicting the specific protein structure into which a given sequence of amino acids will fold is therefore an extremely important task for biology and medicine.

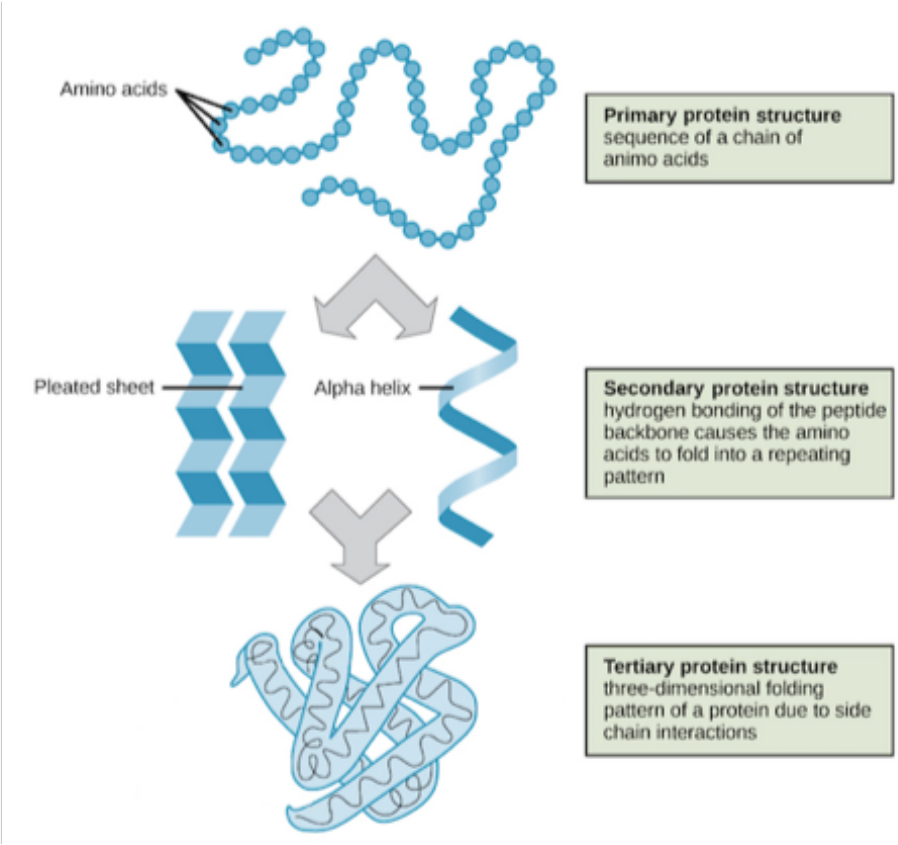

Figure 1. *Illustration of the process of protein folding. Folding of multiple amino acid chains to form a 'Quaternary' protein structure is not illustrated. Image adapted from original produced by the National Human Genome Research Institute.*

A key computational problem is the vast number of possible shapes that a typical protein could theoretically fold itself into: Cyrus Levinthal's widely cited 1969 estimate is 10^{300} possibilities, that is to say, the number one followed by three hundred zeros. This is a staggeringly large number, about 10^{220} times more than the estimated number of atoms in the known universe. To stimulate attempts to improve the prediction of folded protein shapes, a scientific competition called the Critical Assessment of Structure Prediction (CASP) has been taking place bi-annually since 1994. This allows international research teams drawn from academia, the pharmaceutical industry and elsewhere to 'blind test' the accuracy of their predictive models against folded protein structures which have been

painstakingly experimentally determined through techniques such as X-Ray crystallography but not made public.

Given that a body of suitable training data representing a few hundred thousand known amino acid sequences and protein structures was already available in public databases, mainly the Protein Data Bank (PDB), the researchers at DeepMind considered that an equivalent AI system to those they had developed for playing games could conceivably teach itself to predict how a given sequence of amino acids would fold into a specific protein structure.

As 'AlphaGo' and 'AlphaZero' astonished the Go community, so 'AlphaFold1' and 'AlphaFold2' astonished the CASP community, leading their competitive categories by a significant margin in CASP13 in 2018 and CASP14 in 2020. In an impassioned piece titled 'What just happened?", leading molecular biologist Mohammed AlQuraishi said that (here):

> "…there was, in many ways, a broad sense of existential angst felt by most academic researchers at CASP13, including myself. In a delicious twist of irony, we the people who have bet their careers on trying to obsolete crystallographers are now worried about getting obsoleted ourselves…

He had some harsh criticism for the pharmaceutical industry (here):

> "What is worse than academic groups getting scooped by DeepMind? The fact that the collective powers of Novartis, Pfizer, etc, with their hundreds of thousands of employees, let an industrial lab that is a complete outsider to the field, with virtually no prior molecular sciences experience, come in and thoroughly beat them on a problem that is, quite frankly, of far greater importance to pharmaceuticals than it is to Alphabet. It is an indictment of the laughable "basic research" groups of these companies, which pay lip service to fundamental science but focus myopically on target-driven research that they managed to so badly embarrass themselves in this episode… And if you counter with the argument that machine learning is not pharma's core expertise, then you only prove my point: why isn't it?"

Another leading molecular biologist, Andrei Lupas, who was a judge at CASP14, was baffled (here):

> "All of the groups in this year's competition improved…But with AlphaFold, Lupas says, "The game has changed." The organizers even worried DeepMind may have been cheating somehow. So Lupas set a special challenge: a membrane protein from a species of archaea, an ancient group of microbes. For 10 years, his research team tried every trick in the book to get an x-ray crystal structure of the protein. "We couldn't solve it." But AlphaFold had no trouble. It returned a detailed image of a three-part protein with two long helical arms in the middle. The model enabled Lupas and his colleagues to make sense of their x-ray data; within half an hour, they had fit their experimental results to AlphaFold's predicted structure. "It's almost

> perfect," Lupas says. "They could not possibly have cheated on this. I don't know how they do it."

DeepMind had now emerged from the status of a complete outsider to solve the protein folding problem. Along with David Baker, who had independently made significant advances in computational protein design, Demis Hassabis and another DeepMind researcher, John Jumper, were jointly awarded the 2024 Nobel Prize in Chemistry "for protein structure prediction" (here).

It is true, though, that solving the protein folding problem represents only one step toward more commercially relevant 'rational' drug design. The next necessary step is an accurate prediction of how the folded protein form dynamically interacts with other molecules. As mentioned above, many drug molecules ('ligands') function through interaction with the folded protein shape by docking with a binding site on the protein and either enhancing its functionality (for example, mimicking the binding of a natural molecule) or blocking its functionality (for example, disrupting the binding of a natural molecule).

The latest version of AlphaFold, AlphaFold3, was released in early 2024. It is not only capable of making even more accurate predictions about the structure of proteins but can indeed predict how they interact with ligands and a range of other molecules. Its increased relevance to commercial drug discovery is reflected in the fact that its development was undertaken jointly between DeepMind and Isomorphic Labs, a company which DeepMind spun off within the Alphabet group in 2021 to pursue explicitly commercial applications of its work. Isomorphic Labs has announced strategic partnerships with two Big Pharma firms, Novartis and Eli Lilly, to undertake joint AI-driven drug discovery programs.

Given this strategic change of direction, it is interesting to review how DeepMind and Isomorphic Labs have thought about intellectual property. Rather than taking a proprietary approach to their radically new technology and, for example, keeping secret the details of the AlphaFold system, DeepMind published scientific papers describing in detail both AlphaFold1 and AlphaFold2, as well as open sourcing the model code and trained model weights. The scientific community was able to review each of these versions of AlphaFold in detail and thus be assured of the quality and reproducibility of its predictions which, in turn, reinforced DeepMind's credibility. Mohammed AlQuraishi observed of the AlphaFold2 paper that (here):

> "The result is both a *tour de force* of technical innovation and a beautifully designed learning machine, easily containing the equivalent of six or seven solid ML [Machine Learning] papers but somehow functioning as a single force of nature… Soon after CASP14, I had a discussion with a friend in which the topic of how far the academic community was behind DeepMind came up…Now that I have read the paper, I think it would have likely taken at least 5-6 years before the academic community's effort could have added up to AlphaFold2."

Open-sourcing the code and trained model weights also permitted third parties to begin developing their own versions of AlphaFold.

Moving from more pure science to commerce and from sole development to joint development, however, DeepMind and Isomorphic Labs have adopted a different approach for AlphaFold3. To the dismay of the scientific community, the paper describing this latest version was accompanied only by descriptive pseudocode and for several months researchers could only access AlphaFold3 through a web-based interface and only for purely non-commercial purposes. By way of some justification, DeepMind and Isomorphic Labs indicated that they were consulting with biosecurity experts to weigh the threat that the improved capabilities of AlphaFold3 could be used by those with malign intent to design or improve bioweapons.

In fact, the AlphaFold3 source code has now been released under an open-source licence. However, in this case, requests for access to the trained model weights are only permitted from academic researchers on a case-by-case basis and again only for purely non-commercial purposes. Further evidence of a more commercial approach is provided by Isomorphic Labs having developed additional proprietary predictive and generative models which explore not only molecular binding but also, for example, "…their inherent solubility, permeability and metabolic properties which are essential for developing a successful drug" (here). It remains to be seen what the intentions of DeepMind are regarding the patent applications which have been filed to protect their AI / machine learning methodologies since at least 2018.

DeepMind and IsoMorphic Labs are only two of a rapidly expanding number of firms using AI systems in the drug discovery field and new initiatives and advances are continuously being announced. EvolutionaryScale, a start-up funded earlier last year with $142 million raised from investors including Amazon and Nvidia, has not only released its own AI system with which to study proteins ('ESM3') but has already used it to design a wholly artificial green fluorescent protein ('esmGFP'), which should prove useful in biotechnology. The researchers at EvolutionaryScale suggest that the design of esmGFP represents the equivalent of five hundred million years of natural evolution. (See the press release, here.) Another example that is further down the drug development pipeline is provided by Insilico, a company which has used AI techniques to identify the protein most likely responsible for causing Idiopathic Pulmonary Fibrosis, a stiffening and scarring of the lung walls, and to design a drug candidate to disrupt the operation of this protein. In 2023, this drug candidate became the first designed using AI techniques to enter phase II clinical trials (see, for example, Insilico blog here.) Finally, the scientist quoted in the introduction is Simon Kohl, who has recently left working on AlphaFold at DeepMind to found another AI-firm, Latent Labs, also with the intention of supporting pharmaceutical industry work on designing new proteins.

There is also a flourishing open-source community in drug discovery. Examples include those building on the open-sourcing of AlphaFold3, mentioned above, and Rosetta (here), originally created by the above-mentioned Nobel laureate David Baker.

How likely is it that such AI systems, whether proprietary or open-source, will revolutionise drug-discovery, reducing its cost, risk and time taken? Some positive evidence is already being reported. For example, Insilico have stated their use of AI techniques (here):

> "…was a milestone in drug discovery – from novel target discovery to Phase 1 in under 30 months, about half the time the process would take with traditional drug discovery, thanks to the speed and efficiency provided by AI, and for a fraction of the cost."

A recent review looking across a much larger data set ("Unlocking the Potential of AI in Drug Discovery", Boston Consulting Group / Wellcome Trust, available: here), came to a similarly positive conclusion:

> "Modelling based on extrapolation of publicly available data from early AI programmes suggests AI-driven R&D efforts from discovery up to preclinical could deliver time and cost savings of at least 25-50%."

and:

> "Whilst time and cost savings are helpful, modelling shows it will be improvements in probability of success in the clinic that delivers the biggest impact from AI in drug discovery and a step change in the economics of R&D."

Being able to terminate the study of drug candidates which are unlikely to be successful at an earlier stage than before could have a huge impact on the cost and risk of a large pharmaceutical company's 'portfolio' of drug discovery and development. In this sense, the use of AI systems in drug discovery has memorably been described as a helpful 'grim reaper', see here.

However, it is still perhaps too soon to tell whether the impact of AI in drug discovery will truly be as transformative in the long run as this initial evidence might suggest. There are still many sceptics, not least due to the steps in commercially relevant rational drug design that remain between modelling protein / ligand interactions in the abstract and modelling them in the much more complex and messy environments of cells and human bodies. Derek Lowe, a well-known commentator in the drug discovery field, recently observed of AlphaFold3 (here) that:

> "In almost every area it's a significant improvement over anything that we've had before - including previous AlphaFold versions - and in some of them (protein-antibody and protein-RNA) it appears to be (for now!) the only game in town, even though it's not an infallible oracle in those cases by any means…I'm particularly interested in the protein-small molecule ligand accuracy, though, as a guy who's spent the last few decades working on those things. The numbers presented show that AF3 is indeed better than any of the competition for predicting such structures, based on the competition sets, anyway…[but]… *Structure is not everything.* It's very useful, very good to have, and it will accelerate a lot of really useful research. But it does not take you directly to a drug, nor to a better idea about a target for a drug, nor to a better chance of passing toxicity tests, nor to a better chance of surviving oral dosing and the bloodstream and the liver. Better structure predictions are tools

> that we can use to attack those crucial problems, but they don't answer any of them. Drug discovery has not been solved by software, no matter what you might read."

A recent Nature editorial likewise expressed caution (here): "AI's potential to accelerate drug discovery needs a reality check. Companies say the technology will contribute to faster drug development. Independent verification and clinical trials will determine whether this claim holds up." A recent article in STAT news reported that two firms claiming to have used their proprietary AI systems to generate new drug-related antibodies from scratch are better understood as having optimised known antibodies (here). Interestingly, the latter article also notes a call from industry and academia to develop a competition similar to the CASP protein folding competition, in which teams could compete on their ability to generate new (*de novo*) antibodies against test targets.

One key limit to the ability of AI systems to improve in the field of drug discovery, both in terms of proprietary and open-source systems, may be lack of access to suitable training data. The explosive increase in performance of AI systems in Natural Language Processing, such as ChatGPT, has been supported by access to enormous quantities of high-quality training data in digitised libraries and on the internet. (The resulting disputes over copyright infringement and misappropriation of authors' and artists' work is an important issue but not one studied in this paper.) By contrast, far less high-quality training data exists to train AI systems in the field of drug discovery and much of that appears to be held by pharmaceutical companies, suggesting another reason for tie-ups between AI firms and pharmaceutical companies. (The issue of access to publicly held health data, which could also be very valuable in training AI systems, is discussed in section 7 below).

Another challenge is perhaps the attitude of drug regulatory agencies towards candidate drugs developed using AI systems. How much confidence are these agencies likely to be able to have in data generated *in silico* from modelling? In the future, might it substitute for traditionally generated pre-clinical data? Anticipating the following section on drug development, might it even in the further future substitute for traditionally generated clinical trial data? Regulatory agencies have now opened dialogues with the pharmaceutical industry and the public on the acceptable use of AI systems in drug discovery (for example, see discussion papers from EMA here, and FDA here.)

3. The use of AI systems in patent offices

The use of AI systems in drug discovery is not the only direction from which the patent system might be impacted. Patent offices have also begun to use AI systems to improve the quality of the patent granting process. For example, the European Patent Office (EPO) has been training AI systems using the millions of documents in its own internal patent and prior art (literature) database to assist in patent search and classification, as well as machine translation. A new EPO search tool for patent examiners, AI-PreSearch, was accordingly publicly announced in late 2023. A publicly available AI-powered legal search tool has also just been released by the EPO (here).

4. Drug discovery, AI and patentability.

4.1 Introduction

Pharmaceutical companies tend to file patent applications identifying promising candidate drugs once drug discovery has concluded but before drug development begins; that is to say, after pre-clinical testing has been concluded but before clinical trials in humans have begun. With the preceding sections in mind, we can now therefore consider what might happen when pharmaceutical companies make use of AI systems to invent new drugs and file patent applications for them and, equally, when patent offices make use of AI systems to examine their patentability. Two of the most fundamental requirements for the patentability of an invention are that is it 'novel' and that it involves an 'inventive step', as now expressed in Art. 27 (1) of the World Trade Organization (WTO)'s Trade-Related Aspects of Intellectual Property Rights (TRIPS) Agreement. The following sections therefore focus on these tests of novelty and inventive step before also casting an eye over the additional question of inventorship.

4.2 Novelty

The rationale for the novelty test is that the grant of a patent is supposed to be a reward to tempt the inventor to disclose an invention to the world which was not known before. If they, or someone else who had made the same invention, had already made it public then there would be no need to reward the inventor as the world would already have it. It would also be wrong to deprive the public of the free use of an invention that is already known by later granting a patent over it. Different jurisdictions adopt different definitions in applying this test. Under the European Patent Convention (EPC), for example, novelty is regarded in absolute terms (Art. 54 EPC): the publication (disclosure) of the invention anywhere in the world in any form before the patent application is filed 'destroys' the novelty of the invention such that that no valid patent can thereafter be granted.

Bearing in mind the brief review of the use of AI systems undertaken in the previous sections, it is not difficult to imagine several different possibilities for how the test of novelty might be impacted:

- Patent examiners undertake prior art (literature) searches to determine whether the inventions described in patent applications are novel or not, according to the relevant national or regional patent law. Searches using AI systems are likely to be more effective than those currently undertaken using present day computer assisted methodologies. Such searches may therefore raise the likelihood of a given invention, such as a drug molecule, being found in the prior art and, hence, being found to lack novelty.

- At the frontier of the use of AI systems in drug discovery, it may be that new inventions, such as new drug molecules or new uses for known drug molecules, are more likely to be more novel than before. This could be especially so, for example, if

the AI systems were able to explore previously unknown families of drug molecules with radically different structures.

- Extensive textbooks, handbooks, catalogues and databases have long been compiled and published for public use. DeepMind has arguably turbo-charged this type of disclosure, however, by releasing the AlphaFold Protein Structure Database, describing over 200 million new protein sequences and predicted folded structures. Although this decision appears to have been motivated by a desire to strengthen the research resources publicly available in this field, perhaps recalling that AlphaFold was originally trained using data from publicly available databases too, these sequences and structures will therefore no longer be strictly novel for the purposes of patent law. (This is not to say that inventions related to, or derived from, or using those sequences and structures could not be patentable, just that the sequences and structures would not themselves be patentable *per se*.) In a parallel project, DeepMind has also made public a database containing over two million new crystal structures for investigation by material researchers with applications in, for example, new batteries and superconductors.

- Companies can choose to behave in a similar way when they have decided not to pursue patent applications for the result of some research project but want to ensure that no-one else will be able to patent those results either ('defensive publication'). The use of AI systems in a similar but more deliberately defensive way to that of DeepMind in publishing the database could likewise turbo-charge this type of strategy.

It is important to keep this latter prospect in perspective. The space of possible chemical compounds is absolutely vast – one recent project enumerated more than one hundred billion possible organic molecules even only containing up to seventeen atoms (here) – so it is perhaps more likely that it would be relevant in much more limited contexts. It is nevertheless already causing some concern among patent professionals. For example, Matthew Chun (2023) has suggested that:

> "…in the current environment, companies are incentivised to publish massive libraries of AI-generated compounds as early as possible, with little confirmatory experimental testing, in the hopes of carving out "untouchable" chemical spaces within which companies are prevented from obtaining patent protection."

To obviate this problem, Chun suggests that:

> "…Congress could consider a sui generis carveout for AI-generated prophetic drug molecules, granting them anticipatory effect only if they are publicly disclosed with an established utility…AI-generated compounds disclosed without further testing would no longer qualify as prior art, restoring the incentive to test and develop these compounds into marketable drugs." (*ibid*)

As discussed further below, patent law has certainly been amended in the past to permit carve-outs from the usual novelty rule for some pharmaceutical inventions. However, Chun's proposal would represent the blending of two different patentability concepts. In addition to novelty, patentable inventions must also be capable of 'industrial application', or equivalently must be 'useful', as now expressed in Art. 27 (1) TRIPS. This test became important, for example, in the context of the debate over patenting DNA sequences. However, Chun's proposal seemingly inverts this example – instead of requiring that an invention possess an industrial application to be patentable, he suggests that molecules generated using AI systems should only prevent those molecules from later being patentable (on the ground of lack of novelty) if an industrial application is already known for them at the time.

4.3 Inventive Step

The rationale for the inventive step test is that there is no reason to reward an inventor with a patent for something which, while strictly speaking is new, only involves a straightforward and predictable step beyond what the public knew before. To justify the grant of a patent, the invention should instead involve an unexpected or 'inventive' step forward. Again, different jurisdictions adopt different definitions in applying this test. Under the European Patent Convention, for example, the test is conceived in terms of considering whether the invention, as described in the patent application, would be obvious to a fictitious 'Person Skilled in the Art' (PSA) (Art. 56 EPC). This PSA is assumed to be of average technical ability and to be familiar with all the standard knowledge in the field of the invention available through publications, oral descriptions, by use or in any other way. As appropriate, a PSA may be regarded as a team of researchers rather than an individual.

Bearing in mind the brief review of the use of AI systems undertaken in the previous sections, it is not difficult to imagine several different possibilities for how the test of inventive step might be impacted:

- Searches undertaken by patent examiners using AI systems are likely to be more effective than those currently undertaken using present day computer assisted methodologies. Even if a given invention, such as a drug molecule, cannot explicitly be found in the prior art, it may be that sufficiently close and suggestive references are more likely to be found which, in combination, make it more likely that the examiner (with the PSA in mind) will judge that the invention is obvious. (A step further would perhaps permit AI systems to help patent examiners in reaching that judgment by suggesting how references could best be combined; a step perhaps too far for patent offices might even see human patent examiners replaced by AI systems?)

- At the frontier of the use of AI systems in drug discovery, if equivalent systems – including equivalent training data - are not widely available to other pharmaceutical companies and, hence, notionally PSAs, then it may be that new inventions, such as new drug molecules or new uses for known drug molecules, are more likely to be judged inventive than before. Again, this could be especially so, for example, if the

AI systems were able to explore previously unknown families of drug molecules with radically different structures and modes of operation.

- However, the situation would be different if equivalent proprietary or open-source AI systems – including equivalent training data – had become widely available to other pharmaceutical companies and, hence, notionally PSAs. In this case, it may be that a patent examiner (with the PSA in mind) would be more likely to determine that the invention lacks inventive step. To make a simple analogy, it was mentioned above that the protein-drug molecule docking problem is often compared to the problem of fitting a key to a lock. If, when given the publicly known design of a lock, an easily available AI system can automatically generate the optimal design for a key for any user, no valid patent could be granted to protect that key design; it would be obvious to anyone equipped with that AI system 'tool'.

The prospect that many or all pharmaceutical inventions could eventually become regarded as obvious by PSAs using AI systems is already provoking a degree of concern amongst some patent professionals.

Looking at a broader context, Ryan Abbott published a striking paper entitled 'Everything is Obvious' (Abbott 2019), stating that "As inventive machines continue to improve, this will increasingly raise the bar to patentability, eventually rendering innovative activities obvious. The end of obviousness means the end of patents, at least as they are now". Looking more specifically at drug discovery, Olga Gurgula (2020) suggests that:

> "AI should be incorporated into the standard of the skilled person as a tool that such a person uses to achieve the invention. Importantly, to establish an appropriate level of skills and technology for the skilled person, such a person must be equipped with an equivalent AI that was used in the creation of the invention or the best available AI in the relevant field."

Noam Shemtov and Garry Gabinson (2022) provide a similar perspective. Speculating on the replacement of human PSAs with a 'Machine of Ordinary Skill In The Art' ['MOSITA'], Daniele Fabris (2020) suggests that:

> "In the most extreme scenario, hence, one could even argue that any new invention, at least in the pharmaceutical field, might be declared obvious from the point of view of a MOSITA. So that the pharmaceutical industry might soon be completely excluded from the patent system."

In order to avoid this 'unthinkable' outcome, Fabris suggests that:

> "I believe, however, that a solution to avoid such a drastic consequence could be found in current patent law and jurisprudence, and namely in the "secondary considerations" of non-obviousness." (*ibid*)

where such 'secondary considerations' could look, for example, to factors such as how long it has taken for a solution to a particular problem be presented: "a long-felt but unsolved need".

Such secondary considerations can already be taken into account by examiners at the European Patent Office. They are usually used by patent applicants to try to strengthen what would otherwise be a weak case to demonstrate that an invention is sufficiently inventive. Arguing on the basis of satisfying a 'long-felt want or need' would broadly mean showing that previous attempts to solve an important identified problem have been unsuccessful but that the proposed invention has now successfully solved it. Commercial success can be another secondary consideration, although it is important to discriminate between the invention itself being responsible for that success and, for example, better marketing.

4.4 Inventorship

This latter discussion of inventive step raises a related problem. Although the TRIPS Agreement does not state explicitly that only humans are able to be inventors for the purposes of obtaining the grant of a valid patent, recent decided cases in, for example, the United Kingdom, Europe, Australia, New Zealand and the United States (the 'DABUS' / 'Thaler' cases) have all reached the conclusion that an AI system cannot be regarded as an inventor *per se* for the purposes of patent law. If the properly identified inventor in a patent application is an AI system then, however inventive, the resulting invention may be unpatentable.

Discussion is therefore taking place in the patent profession regarding how to assess the respective contributions of humans and AI systems: at what point is the contribution of the AI system to making an invention so significant that the human cannot be regarded as the inventor? Patent offices have begun to discuss and release preliminary guidance on this question (see, for example, the USPTO here and the EPO here). Antonia Sequeira, Carl Morales and Fredrick Tsang have recently authored a note entitled "Emerging Legal Terrain: IP Risks from AI's Role in Drug Discovery" (here) in which, drawing on the USPTO treatment, they state that:

> "…if AI were merely used as an acceleration tool in drug discovery cycles, it should not be treated differently from other computer or laboratory tools that have existed for decades. Hence, the researchers will still be the rightful inventors. Yet, if the researchers cross an unidentified threshold by allowing AI to replace the human role of inventing, IP rights may be threatened."

Under US patent law, a distinction is made between the 'conception' of an invention and its 'reduction to practice'. It is the former which counts for inventorship, not the latter. The authors indicate that this could have the following consequence regarding the use of AI systems in drug discovery:

> "A conception is completed when the inventing entity has formed the idea in a sufficiently final form and has a reasonable belief that the invention is operative. This sounds similar to a situation where a machine learning model (inventive entity) outputs a compound formula in digital form (sufficiently final form) with a predicted affinity score representing the affinity for a target enzyme (appreciation that the compound is operative in interacting with the target enzyme). As such, under a superficial application of inventorship law, it seems that companies are at high risk of losing patent rights because, in some situations, conception could be deemed completed at the moment AI digitally outputs a chemical structure or formula. In such situations, humans have no opportunity to contribute to the conception." (*ibid*)

However, the authors point out that US case law supports the view that, where the synthesis of such a compound is difficult and unpredictable, conception and reduction to practice may occur simultaneously with the successful isolation of that compound. In which case, the relevant human researcher whose work enabled reduction to practice may yet have made a sufficient contribution for the purposes of inventorship. The authors indicate that:

> "Logically, this leads to an observation that the risk level of losing patent rights with AI involvement is different for different classes of drugs. For example, antibodies and polypeptides may become the highest risk class for a company to lose patent rights if AI outputs the sequence and there is no subsequent human alteration of the sequence. The reason for this is that there are well-established synthetic methods for preparing antibodies and polypeptides of almost practical sequence…at the moment that AI generates an amino acid sequence, it is almost certain that the sequence can be synthesized in the lab. In this type of setting, a court in the future may hold that conception of the sequence is completed by AI alone…The risk for a small molecule drug is lower for using AI tools in screening. A small organic molecule is typically three-dimensional in its overall structure (instead of a linear sequence), and the manner of synthesis is typically considered by courts as unpredictable and can vary molecule by molecule. Hence, a synthetic organic chemist will likely need to be involved in determining how the digital formula outputted by the AI can be made into a physical form, thereby giving the organic chemist a contribution to conception under the doctrine of simultaneous conception and reduction to practice. From the IP perspective, employing another AI model to replace the organic chemist's role in determining how the organic molecule should be synthesized would likely increase the IP risk." (*ibid*)

This seemingly produces a paradoxical situation. The introduction of AI systems into drug discovery may hugely improve that process but, at some point, it may become so automated that the inventions made become unpatentable. To preserve that patentability, it may be necessary to 'shoehorn' a human back into the inventive process, even at the risk of reducing its efficiency. This would surely look like a case of the 'tail wagging the dog'. At the limit, it must instead be sensible to consider designing a replacement system which would reward drug discovery which was as efficient as possible, whether employing humans or not.

A difficult ancillary question these considerations raise is how often human applicants might simply conceal the fact that an invention was made by an AI system. Art. 29.1 of the TRIPS Agreement requires that "…an applicant for a patent shall disclose the invention in a manner sufficiently clear and complete for the invention to be carried out by a person skilled in the art…". What if, for example, a human is improperly listed as the inventor on a patent application which adequately discloses a new drug molecule, as well as a satisfactory synthetic pathway for making it, when an AI system is wholly responsible for inventing it? How will a patent office know? For that matter, might this already be happening?

4.5 Discussion

The review undertaken in the previous sections suggests that the use of AI systems in drug discovery could have several different outcomes depending on which parties have access to them and when:

- Baseline case. It is possible that the use of AI systems in drug discovery will prove useful but not transformative. It is likely that more efficient AI-assisted search and examination in patent offices will nevertheless raise the bar for the novelty and inventive step tests, making it more difficult for pharmaceutical companies to obtain patents.

- Asymmetric transformative case. If a pharmaceutical company is the only one to possess an AI system which is capable of reliably supporting the human-led discovery of unusually novel and inventive drugs then it will likely be easier for that company to obtain patents. However, if the AI system becomes too capable and discovers unusually novel and inventive drugs without any human intervention at all then it will likely become more difficult again for that company to obtain patents since, at least at present, an AI system cannot validly be listed as the inventor.

- Symmetric transformative case. If equivalent proprietary or open-source AI systems capable of discovering unusually novel and inventive drugs eventually enter widespread use such that they can be regarded as a standard tool for the notional PSA then it will likely become more difficult for pharmaceutical companies to obtain patents as inventions are more likely to be determined to be obvious. Again, if these AI systems become too capable and discover unusually novel and inventive drugs without any human intervention at all, then it will also likely become more difficult for companies to obtain patents since, at least at present, an AI system cannot validly be listed as the inventor.

As discussed above and as will be discussed further below, it is possible that amendments to patent law could be considered to bring AI-assisted inventions better back within the scope of the patent system.

5. The use of AI systems in drug development

This section reviews some recent developments in the use of AI systems in drug development, focussing on clinical trials. Drug development is the process of testing the new molecules (or new uses for known molecules) identified during drug discovery to determine whether they are indeed safe and effective in treating the intended human disease-related condition.

Clinical trials are regarded as the most reliable methodology with which to test new drugs in humans. However, that reliability currently involves spending a great deal of time and money – usually the largest single element of the cost of discovering and developing new drugs – and it is recognised that there may be a great deal of room for improvement in how they are conceptualised and run. A recent article in Nature (Hutson 2024) reviewed several areas in which researchers are investigating how AI systems could be used to improve the efficiency and speed of clinical trials including:

- Trial design. A robust trial design is important to maximise the likelihood of generating reliable and useful data. AI systems are being developed which more accurately summarise the design and outcome of previous clinical trials than previously possible, to help those designing future clinical trials with design decisions such as drug dosages, number of patients and data to be collected.

- Trial patient recruitment. Nearly all clinical trials reportedly exceed the expected time for recruiting patients. AI systems are being developed which better define suitable participation criteria, to speed up recruitment without sacrificing trial safety. Other teams are developing AI systems to help match those participation criteria with patients more quickly and efficiently, as well as the reverse, matching interested patients with trials having suitable participation criteria.

- Trial patient maintenance. AI systems are being developed to help with monitoring patients in clinical trials and directing extra clinical support to those in most danger of dropping out.

- Trial data management. The burden of data management from clinical trials is enormous. AI systems are being developed to help with aggregating trial data in many different forms as well as generating reports for submission to drug regulatory authorities.

As with the use of AI systems in drug discovery, there are challenges to overcome such as obtaining access to suitable training data. However, given the safety critical aspect of running clinical trials, there are additional challenges such as ethical concerns relating to the transparency of AI system outputs which regulatory authorities may need to take into account. If these challenges can be overcome, given that clinical trials are usually the largest single element of the cost of discovering and developing new drugs, the use of AI systems in drug development could have a more pronounced and more rapid business impact than in drug discovery.

6. AI, drug development and pharmaceutical patents

Considering that the use of AI systems might be able to speed up drug development (clinical trials), the following sections consider what impact this could likely have on pharmaceutical patents, as well as on another type of relevant pharmaceutical IP, so-called data exclusivity.

6.1 Patents and patent extensions

Pharmaceutical patents *per se* will be unaffected by a speeding-up of drug development. However, many jurisdictions do now provide for regulatory patent extensions, which are designed to compensate patent holders for some or all of the patent life-time taken in drug regulatory review during which their product cannot yet be marketed. The compensatory patent extension term is accordingly shorter or longer depending on how long that review takes. Such patent extensions are not required under the TRIPS Agreement. The details of the extension mechanisms used vary between, for example, the European Union, the United States and Japan. Strictly speaking, in Europe, the mechanism used is not an extension of the patent itself but an ancillary legal instrument, a Supplementary Protection Certificate (SPC), which is designed to reproduce much of the effect of the underlying patent.

If the use of AI systems were able to reduce the time taken in drug development (the different mechanisms variously represent this in terms of the time between either the filing of the patent application or the grant of the patent and the grant of marketing approval), it would effectively reverse the effect to compensate for which the patent extension system was introduced. Instead of adding longer patent extensions to make up for longer periods of patent-based market exclusivity lost in the regulatory process, shorter patent extensions would automatically be added to make up for the shorter periods of lost patent-based market exclusivity. At some point, no patent extension would need to be added at all.

Some jurisdictions do not offer regulatory patent extension mechanisms at all. Others do offer patent extension mechanisms but do not necessarily offer a complete compensation for the loss of patent rights. For example, in the EU, the term of the SPC is designed to add to the remaining life-time of the underlying patent term to ensure a constant period of market exclusivity of fifteen years from marketing authorisation. However, it restricts the relevant 'drug development' period between filing the patent application and obtaining marketing authorisation to between five and ten years; at or below five years no SPC is added whereas above ten years, no SPC longer than five years can be added. Reducing the drug development period from above this range may therefore increase the absolute market exclusivity time enjoyed by the patent holder, as may reducing it from the bottom of the range even further. Further, in the EU, the SPC is more limited in scope than the underlying patent. It is restricted to protecting only the specific medicine which has received marketing approval. A shift back to a longer period of patent-based market exclusivity and a correspondingly shorter period of SPC-based market exclusivity would accordingly increase the scope of the patent holder's later market exclusivity.

6.2 Data exclusivity

Art. 39.3 TRIPS requires that:

> "Members, when requiring, as a condition of approving the marketing of pharmaceutical or of agricultural chemical products which utilize new chemical entities, the submission of undisclosed test or other data, the origination of which involves a considerable effort, shall protect such data against unfair commercial use."

WTO Members are not required to interpret protecting against unfair commercial use as requiring a period of exclusivity for that data but, for example, the European Union, the United States and Japan do, although the details of the exclusivity mechanisms used again vary. Such exclusivity is regarded as providing an additional incentive to recompense pharmaceutical companies for the cost of undertaking preclinical tests and clinical trials as a prelude to demonstrating to a regulatory authority that their medicines are safe and effective.

In the European Union, following the submission of such test data by a pharmaceutical company to the European Medicines Agency and the grant of marketing approval, competitor companies are prevented from making reference to that data to demonstrate that their bioequivalent or biosimilar version of the medicine is also safe and effective (data exclusivity) for a period of eight years, followed by a further two years during which competitor companies are otherwise prevented from marketing their own versions of the medicine (market exclusivity). The duration of this combined ten-year period (which can be extended by a further year for a new therapeutic indication) is not the result of any objective policy decision. Rather, as a representative of the European Commission indicated: "In the end, it [data exclusivity] is not a mathematical decision; it is a political estimation of how much incentive for innovation you need" (quoted in Adamini et al. 2009). Whether or not drug development is speeded up by the use of AI systems will have no impact on the term of these European data/market exclusivities since they are fixed from the date of marketing authorisation.

(N.B. For the sake of simplicity, the term 'data exclusivity' is used in this paper to denote both data exclusivity proper and, in the case of the European Union, the additional market exclusivity too.)

6.3 Discussion

The review undertaken in the previous sections suggests that the use of AI systems in drug development could indeed have an impact in reducing the cost, risk and time taken for clinical trials but that it would only be reflected in an automatic reduction in the term of regulatory patent extensions and not in the fixed term of base patents or data exclusivities.

7. An overall incentive perspective on pharmaceutical innovation

7.1 Quick recap of the rationale for the patent system

It has been interesting and valuable to examine the individual 'nuts and bolts' issues raised in the preceding sections regarding the patent and data exclusivity systems. However, as mentioned in the introduction, it is also interesting and valuable to bring them together and reflect on what they mean in terms of an overall incentive structure for encouraging pharmaceutical innovation. By way of background, this section therefore first quickly recaps the rationale for having a patent system in the first place.

Patents are one exception to the general rule of free competition in market economies. They permit the suppression of free competition in patented inventions, within a limited scope and for a limited period of time, with the justification that it should result in an overall innovation-related benefit for the public. The rationale is, of course, that private sector inventors and their investors will only invest in innovative research and development projects if they think the sale of the resulting invention is likely to be sufficiently profitable. That likelihood will be increased if they are able to obtain a patent for the invention and thus suppress its market competition, again albeit within a limited scope and for a limited period of time.

It is true that the public will be inconvenienced by obtaining access to the patented invention at a higher price in the shorter term, during the patent lifetime, than otherwise would be the case if there were market competition. However, without the patent, the invention might not have been made at all and they will have access to the invention at the lowest price able to be provided through market competition in the longer term, once the patent has expired. 'Inconvenienced' takes on a very different meaning, of course, in the context of access to essential medicines: for patients whose health or life could be maintained with (but only with) a newly patented drug, there is a stark difference between being able to access it at an affordable competitive price and not being able to do so due to an artificially high price maintained through a patent monopoly. This will be true whether the drug is unaffordable for the patients themselves or whether supply to the patients cannot be justified on cost grounds by health insurance companies or public health systems.

This simple recap raises a number of important questions. Perhaps the most obvious is: how long should patents therefore last? The suppression of market competition is sufficiently serious that patents should presumably last for the absolute minimum time consistent with providing an adequate incentive to inventors. Should patents therefore last different periods of time for different technical fields, depending on the difficulty of making inventions? Indeed, should those periods of time be able to be adjusted up and down depending on changes in the difficulty of making inventions? Although these questions have prompted rich and long-standing debates among economists, lawyers and others, for

a review of which see, for example, Lester and Zhu (2019), the real-world answer to these questions is a surprisingly unsophisticated one.

The TRIPS Agreement adopts a harmonized one-size-fits-all approach: Art. 33 specifies a basic patent term of twenty years from the date of filing a patent application, with no variation by field of technology and no variation with changes in the difficulty of making inventions. (Art. 27.1 TRIPS specifies that "...patent rights shall be enjoyable without discrimination as to the…field of technology…") This approach is not the result of any detailed analysis suggesting that twenty years optimises the overall public benefit flowing from the patent system averaged across all fields of technology. Rather, it represents the end of a slow evolution from, for example, a term of fourteen years being set in English law as the time taken to train two generations of apprentices (Statute of Monopolies, 1624), gradually lengthening to a widely accepted twenty years by the late twentieth century through a process of haphazard comparison and harmonisation between different countries.

The adoption of a one-size-fits-all approach suggests that patent lifetime has expanded to satisfy the needs of those technical areas with the highest difficulties in making inventions, thus likely over-rewarding other technical areas with lower difficulty. In addition, as already touched on in this paper, the pharmaceutical industry has argued for and received additional and even longer-lasting protection through regulatory patent extensions and data exclusivity, justified in terms of the ever-increasing cost, risk and time taken for drug discovery and drug development. The consistency and magnitude of the profitability of at least the larger pharmaceutical firms ('Big Pharma'), who are able to better manage risk through a portfolio approach to drug discovery and development than smaller pharmaceutical firms, suggests that they may have been over-rewarded with this additional protection too.

7.2 Hypothetical future scenarios

With this recap in mind, let us now return to the development of AI and its use in drug discovery and development. Let us consider the following hypothetical future scenarios:

- The development of AI systems by AI firms and/or pharmaceutical companies could bring about a significant reduction in the cost, risk and time taken for drug discovery and drug development. The most immediate 'low hanging fruit' could lie in speeding up clinical trials. Regulatory patent extensions would reduce in term although the term of the base patent and data exclusivity would remain the same. The profitability of the pharmaceutical industry business model could well increase.

For the reasons discussed in the preceding sections it is plausible that the use of AI could materially reduce the cost, risk and time taken for drug discovery. It is true, however, that the necessary investments in AI systems could be very expensive in the short to medium term, including generating high-quality training data for drug discovery, which could delay cost savings. For the reasons discussed in the preceding sections, it is also plausible that the use of AI in clinical trials could materially reduce the cost, risk and time taken for drug development. Although the balance of difficulty between drug discovery and development

will vary between particular disease areas, it is therefore plausible that the overall cost, risk and time taken for drug discovery and development will fall. Let us assume that not only do these reductions occur but they are materially large, widespread and persistent across the pharmaceutical industry. (Let us also assume that the pharmaceutical industry does not switch to a personalised model of medicine, with target disease markets effectively falling to single individuals, such that patent-based models are no longer useful in any case.) What would this mean in terms of the overall incentive model?

At least one element of the overall incentive will automatically respond to the reduction in time taken: patent extensions will begin to reduce in length. However, the other elements of the overall incentive – the base patent and data exclusivity – will not respond in this way. Once the length of patent extensions has fallen to zero, it could be argued that this trajectory should ideally continue such that the overall patent incentive continues to decline. However, of course, it cannot be reduced any further at the present as it hits the 'bedrock' of the base patent, with its fixed twenty-year term. The situation is similar with data exclusivity. There is no correspondence at present between the cost of generating the necessary pre-clinical and clinical trial data for regulatory approval and the length of the data exclusivity period. As the cost, risk and time taken for drug development reduces this therefore has no impact on the length of the data exclusivity period unless, perhaps, the reduction is so pronounced that the Art. 39.3 TRIPS test is no longer passed ("…the origination of which involves a considerable effort"), in which case no data exclusivity would be granted at all, *i.e.* a binary on/off response.

Such reductions in the cost, risk and time taken for drug discovery and development would therefore suggest that the pharmaceutical industry will become progressively more rewarded, in incentive terms, at the expense of the public. In fact, as mentioned above, since it is arguable that many pharmaceutical firms have already been over-rewarded, this would suggest that they are likely to become progressively even more over-rewarded. It is difficult to be certain that this would be the case. On the one hand, it might be that competitor pharmaceutical companies are able to use AI to compete more effectively, to better design around patented drugs or to better substitute them with other drugs with similar clinical properties, weakening market monopolies; although it may seem unlikely, it might even be that pharmaceutical companies pass on the reductions in voluntarily lowering their prices. On the other hand, it might be that pharmaceutical companies are able to use AI to engage even more effectively in practices such as evergreening, strengthening and/or prolonging their market monopolies. Let us nevertheless assume, therefore, that the pharmaceutical industry could become progressively (even more) over-rewarded.

- Although an increase in profitability would further stimulate the pharmaceutical industry business model, many dissatisfied with what that model delivers would likely argue that the time had come to scale-back the patent and data exclusivity incentives – perhaps through shortening them and/or converting them from an exclusivity regime to a licence of right (liability) regime - as they would have become overly favourable to industry at the expense of the public.

Policy makers could respond in a few different ways, alone or in combination. They could do nothing and let the pharmaceutical industry become more profitable, potentially turbo-charging their present business model. The public, mostly in wealthy nations, would continue to benefit in the way that they presently do. However, given the evidently unsatisfactory performance of the model in several important respects (for example, contributing to obvious access problems regarding drug pricing and skewing R&D priorities toward profitable disease 'markets' irrespective of public health need), many would certainly object to raising the profitability of the pharmaceutical industry even further.

Policy makers could instead therefore reduce that over-reward through, for example, revisiting drug pricing and profitability models or scaling-back the overall patent and data exclusivity incentive to return to the *status quo ante*. Such a reduction would emphatically not be about penalising the pharmaceutical industry and/or their AI firm collaborators for their success in developing new technologies and improving drug discovery and development but would instead be about restoring the appropriate balance between incentivising inventors and investors and delivering benefits to the public as quickly and affordably as possible.

Regarding the patent system, it is perhaps interesting to reflect on one prior example where precisely such a scaling-back was advocated. The Sainsbury Committee was set up by the British government in 1965 to undertake a wide-ranging investigation of the operation of the pharmaceutical industry and its relationship with the British National Health Service (NHS), to see whether the public interest was being best served with the contemporary patent system and pricing arrangements. At that time, British patents lasted sixteen years. It is true that there was also a special compulsory licensing regime for pharmaceutical patents although it was so cumbersome and so little used that it had hardly any practical effect. The Committee was very much aware of the over-rewarding problem:

> "If the patent monopoly is stronger or lasts longer than is necessary to achieve its intended purpose, then prices may be higher, or high prices may persist longer, than is necessary. Unless some control or regulation of prices were introduced, the public would in such cases be paying the pharmaceutical companies more than would be required to induce them to do their job." (Sainsbury Committee 1967, para. 146)

However, as to what was the appropriate patent term, the Committee found itself in difficulty:

> "The industry produced no factual evidence, so far as the United Kingdom is concerned, to support its case for an extension of the patent period, but it is fair to point out that the debate over the optimum length of the patent period has been continuing in European countries for over 150 years and no one has yet been able to do much more than draw inferences from *a priori* propositions. We must thus sympathise with the industry's difficulty. There is, so far as we can ascertain, no objective way of measuring the effect on the amount of invention and innovation that a patent monopoly lasting 16 years would call forth as contrasted to one of 10 or of 20 years. The same difficulty is inherent in the problem of appraising other

> provisions of the patent law. The simple assumption made by the industry seems to be that the longer and more complete is the monopoly, the greater will be the flow of new pharmaceutical products." (*ibid*, para. 144)

Based on the evidence reviewed, the Committee nevertheless concluded that:

> "The majority of the Committee is in doubt whether over the extensive field of products of varying importance, the result of varying periods of research, a patent period of as long as 16 years is necessary. Most of us are inclined to think it too long, and that the position would be met by a shorter period of complete protection, the patents having a licence of right endorsement from some intermediate period within the 16 years." (*ibid*, para. 150)

In addition to considering patent term, the Committee thus picked up on another important question raised by the simple recap of the patent system above: how extensive should patent rights be? Instead of a patent providing exclusive rights throughout its whole term, the Committee recommended that a patent should provide only the right to receive a royalty income during the latter part of its term, *i.e.* transforming an exclusive rights regime to a liability regime. (Although the voluntary form of a licence of right signals that a patent owner is willing to provide a licence to any interested party, in this context, a licence of right endorsement is effectively a form of compulsory licence with a reversed burden of proof: instead of an applicant having to prove they should be granted such a licence, an applicant will be granted one unless the granting authority considers them unsuitable to work the invention.)

Rather than making this recommendation a formal one, though, the Committee suggested that these patent law questions should be handed to a parallel Committee – staffed by patent lawyers - for expert assessment. It will perhaps come as no surprise to find that this Banks Committee rejected the recommendation in favour of a significant expansion of patent rights - not on the basis of any domestic cost/benefit analysis but rather to align the UK with ever-higher internationally agreed standards – such that the term was increased to twenty years and the special compulsory licensing regime for pharmaceutical patents was abandoned; the IP-ratchet in action.

Despite this outcome, this paper suggests that an analysis along the same lines as the Sainsbury Committee could be useful in thinking about scaling-back the overall incentive provided by patent and data exclusivity where the use of AI in drug discovery and development warrants it. This would have to take into account an analysis of the degree to which different elements of the overall incentive contribute to profitability: scaling-back in the later years of patent term may have a disproportionately reduced practical impact if the material profits have already been generated in the earlier years. Some thought could also perhaps be given, in the event of collaboration between the private and public sectors, to the manner in which publicly owned health related data could be used in training AI systems. As mentioned above, meaningful training data is extremely valuable for AI firms. If public health authorities have such data and choose to share it with AI firms for training

purposes, they should be aware of that leverage and should consider *quid pro quo* commitments from the firm.

Consider, for example:

(a) although the period of patent extensions should automatically reduce, the associated exclusive rights regime could be partly or wholly converted to a compulsory licence of right (liability) regime, where the level of suppression of market competition could be graded through the payment of appropriate royalty fees;

(b) once the period of patent extensions reaches zero, thought could be given to a matching compulsory licence of right (liability) regime for the latter end of the base patent life-time or, alternatively, patent owners might be encouraged, for example as a *quid pro quo* commitment, to agree a voluntary licence of right (liability) regime for the latter end of the base patent life-time; and/or

(c) the period of data exclusivity could be reduced and/or it could be converted from an exclusive rights regime to a compulsory licence of right (liability) regime, where the level of suppression of market competition could be better graded than a simple binary on/off through the payment of appropriate royalty fees.

Since both patent extensions and data exclusivity are TRIPS-plus measures (that is, measures going above and beyond what is laid out in the TRIPS agreement), there remains comparatively more room for thinking about the design and implementation of any such amendments. It may be though, one day, that the impact of the use of AI in drug discovery and development – and, indeed, across all technical fields of invention – becomes so pronounced that even the most fundamental base patent provisions of the TRIPS Agreement have to be revisited too. It is possible that the rich debates about patent term reviewed by Lester and Zhu (2019) may be about to become much more relevant again.

- In a rather different direction, the development of AI systems by AI firms and/or pharmaceutical companies could instead lead to new drug-related inventions in several important areas becoming unpatentable, for example, if the AI system is itself deemed to be the proper inventor, contrary to the currently widely held belief that only humans can be inventors for the purposes of the patent system, or if proprietary or open-source AI systems become so widespread that the new drug-related inventions become typically regarded as obvious. If excluded to one extent or another from the patent system, the pharmaceutical industry business model in those areas could become threatened.

If pharmaceutical inventions in a particular area fall out of the scope of the patent system then pharmaceutical firms will not be able to rely on their typical business model and can therefore be expected to be much less likely to invest resources in R&D activities in that area. It is possible that pharmaceutical firms could fall back on making greater use of trade secret protection, but the widespread use of AI by competitors would suggest that might not afford much security. There would be important public policy consequences if the

pharmaceutical industry pulled out of R&D areas – the examples of antibodies and polypeptides were mentioned above in section 4.4 – in favour of 'artificially' concentrating on areas where patents were still available. Given the historical profitability of the pharmaceutical industry business model, though, it perhaps seems more likely that the pharmaceutical industry would instead resolve to take action to try to have those areas brought back into the scope of the patent system.

- In response to the latter threat, amendments to national or regional patent laws could be made, for example, to permit AI systems to be properly named as inventors or to amend the test of obviousness to include broader secondary conceptions of inventive step.

Such amendments would not require any high-level amendment to the TRIPS Agreement. They could be implemented as new interpretations of the existing tests, through legal or administrative amendment – perhaps even at the level of patent office examiners' practice – or through litigation.

Two examples can perhaps illustrate the potential for such changes in interpretation. As mentioned above, the present conclusion that an AI system cannot be regarded as a valid inventor for the purposes of patent law stems from the 'DABUS' / 'Thaler' cases. Such a conclusion is not therefore immutable but could be altered by policy makers at a higher level, whether national, regional or indeed international.

The acceptance of so-called 'Swiss' patent claims for second and subsequent medical uses provides another revealing example. Molecules which are already known for some non-medical use can be found to have a medical use, or a second or subsequent medical use. It is helpful to incentivise the search for such new uses of known molecules as it can bring significant therapeutic benefit but potentially at a lower cost, perhaps especially in second or subsequent medical uses where the known molecule should already have been shown to be safe for a first medical use in humans.

The question of whether a second or subsequent medical use could be patented came before judicial and administrative fora in several European countries in the 1980s. Of course, pharmaceutical firms argued that they could only invest significant R&D resources in searching for new uses if patent protection could be extended to cover such activity. The somewhat tortured form of the patent claim eventually contemplated by the Swiss patent office was: "Use of compound X for the preparation of an agent for the treatment of disease Y". Grubb and Thomsen (2010) state that:

> "The Swiss form of claim suffers from the logical objection that it lacks novelty, because it claims the use of the compound for preparation of a medicament and normally the medicament itself will be the same as that already used for the first pharmaceutical indication. Accordingly, there was a concern that patents granted with such claims in the EPO [European Patent Office] could be held invalid by national courts in countries such as the UK. These fears were allayed by an unexpected decision of the UK Patents Court, in which both patent judges sitting

together decided that a British national patent could be granted with claims of the Swiss type. The novelty problem was recognised, but it was held that, for the sake of a common approach in European patent law, it should be considered that the new end-use conferred novelty on the claim."

Art. 54 of the European Patent Convention was subsequently amended to implement this newly accepted exception from the usual novelty rule. This example demonstrates that even previously unthinkable changes in interpretation – indeed it is not unfair here to go so far as to describe the invention of new 'legal fictions' - need not occur as the result of public debate, balancing private and public interests, but can be the result of quite individual decisions. If important changes resulting from the use of AI systems were contemplated to the scope of patentability of pharmaceutical inventions, it would therefore be better to ensure that they were not taken behind closed doors, whether those of judicial fora or patent offices, but were instead made with the benefit of transparent discussion and input from other public-policy perspectives.

- However, instead of artificially straining patent laws so far from historical understandings, it may be better to look to a new framework to encourage innovation in these areas, more suited to the realities of AI systems. Such frameworks can be thought of as part of a broader debate about moving from thinking about protection based on a difficult-to-judge level of invention toward protection more transparently based on levels of investment. Opportunities could also flourish for not-for-profit drug discovery and development with access to AI systems and contract-based drug research and development firms.

Although it perhaps seems unlikely, if insurmountable problems were to derail the use of patents in the pharmaceutical industry business model, there would be an obvious opportunity to re-think and design new and potentially better frameworks to support and guide private sector drug discovery and development. The public ends up paying for drug discovery and development one way or another, whether loaded to the front end (including supportive public investment in basic science or, in some cases, with market-making advance purchase commitments) or the latter end (paying market prices, either personally, through contribution to health insurance companies or through taxpayer-funded health systems), so it is perfectly reasonable to ask whether the venerable patent system is indeed still fit for purpose or whether an alternative system or systems might now do better.

Such frameworks can be thought of as part of a broader debate about moving from thinking about protection based on a difficult-to-judge level of invention toward protection more transparently based on levels of investment. One interesting proposal – so-called 'de-linkage' (see here) – breaks the problem into two separate pieces: it suggests that the costs of drug research and development are not recovered through the patent model but rather that those costs are separately covered so that as soon as new drugs are introduced into the market, they are made available at the lowest possible prices through generic competition without having to wait for patent or other IP rights to expire. Opportunities could also flourish for not-for-profit drug discovery and development with access to AI systems and contract-based drug research and development firms. A range of 'push' and

'pull' incentives have recently been discussed in the context of encouraging the discovery and development of new antimicrobials (here), the latter including publicly funded assurances of minimum viable profits even when the use of a new antimicrobial is necessarily kept to a minimum.

Given these above-discussed hypothetical future scenarios, it seems sensible to begin a closer monitoring of relevant developments regarding the use of AI in drug discovery and development as well as its consequences for the patent and data exclusivity systems. It will be very interesting, for example, to monitor the terms of regulatory patent extensions, at least in those cases where AI systems are being used, to see if evidence of a reduction in the time taken for drug development begins to become apparent.

8. Is there a chance we might no longer need to debate and discuss IP at all?

By way of some final speculation, let us look beyond the scenarios discussed in the previous section to rather more radical possibilities. The possibility that the current trajectory of increase in AI capabilities might lead in the near to medium term to the creation of Artificial General Intelligence (AGI), at or above human-peer level, or even Artificial Super Intelligence (ASI), far above human-peer level, is now being widely discussed.

It is worth noting that in many areas in which AI systems have been made public, scepticism about their ability to successfully complete a particular task has often existed right up until that task was completed, whereupon the scepticism seems to have just been recalibrated to the next task: "I can't believe an AI system could ever write poetry" becomes "well, ok, but I can't believe that it could ever write poetry that could win a poetry prize", becomes "well, ok, maybe it just isn't that hard to write prize-winning poetry, so…". It is possible to lose sight of the incredible rate of overall progress in this way; humans are notoriously bad at understanding exponential rather than linear change.

If AI systems were to reach AGI, at or above human-peer level, it has been suggested that thousands (or millions) of these could explosively automate further AI research overnight, perhaps far outstripping the progress made by the limited number of human researchers working on AI to date and perhaps driving extremely rapidly toward ASI. Imagine that, one morning five or six years from now, an ASI emerges in public and it is not Go players or computational biologists who are astonished and fearful of obsolescence but every leading scientist and thinker in every field.

Assuming that the so-called 'alignment' problem had been solved, such that any such ASI was constrained to act in the interests of humanity, or that any such ASI was otherwise benevolent, then among other things there would be no need for a patent system at all. There would be no need to incentivise human inventors if an ASI 'oracle' could generate all imaginable inventions on a simple request. However, if the ASI was not so constrained or was not otherwise inclined to act in the interests of humanity, then we could have made ourselves evolutionarily redundant. In either case, within just a few years, it is possible that

we might not be debating and discussing the patent system or any other innovation incentive systems anymore at all.

Until that day arrives, though, let us focus on finding ways to best encourage the use of the AI and the scientific expertise that we do have to discover and develop new drugs to treat or cure burdensome diseases and deliver them as affordably as possible to all those in need.

References

Abbott, Ryan Benjamin, Everything is Obvious. 66 UCLA L. Rev. 2 (2018), Available (open-access) at: https://www.uclalawreview.org/wp-content/uploads/securepdfs/2019/01/66.1.1-Abbott.pdf

Adamini, Sandra, et al. "Policy making on data exclusivity in the European Union: from industrial interests to legal realities." *Journal of Health Politics, Policy and Law* 34.6 (2009): 979-1010. Available at: https://read.dukeupress.edu/jhppl/article-abstract/34/6/979/93620/Policy-Making-on-Data-Exclusivity-in-the-European

Chun, Matthew, Artificial Intelligence for Drug Discovery: A New Frontier for Patent Law, 104 J. Pat. & Trademark Off. Soc'y 5 (2024). Available at: https://heinonline.org/HOL/LandingPage?handle=hein.journals/jpatos104&div=6&id=&page= or https://ssrn.com/abstract=4566014

Fabris, Daniele. From the PHOSITA to the MOSITA: Will "Secondary Considerations" Save Pharmaceutical Patents from Artificial Intelligence?. *IIC* **51**, 685–708 (2020). Available at: https://doi.org/10.1007/s40319-020-00953-8

Grubb, Philip W., Thomsen, Peter R., Patents for chemicals, pharmaceuticals, and biotechnology. (5th ed.), Oxford University Press, 2010.

Gurgula, Olga, AI-Assisted Inventions in the Field of Drug Discovery: Readjusting the Inventive Step Analysis. International Journal of Social Science and Public Policy, Vol. 2, No. 8; August 2020, Available (open-access pre-print) at: https://bura.brunel.ac.uk/bitstream/2438/21690/1/FullText.pdf

Hutson, Matthew. "How AI is being used to accelerate clinical trials." *Nature* 627.8003 (2024): S2-S5. Available at: https://www.nature.com/articles/d41586-024-00753-x

Lester, Simon and Zhu, Huan, Rethinking the Length of Patent Terms. American University International Law Review, Volume 34, Issue 4, Article 4. 2019. Available (open-access) at: https://www.cato.org/sites/cato.org/files/2020-02/lester-zhu-auilr-v34n4.pdf

Sainsbury Committee, Report of the Committee of Enquiry: "Relationship of the Pharmaceutical Industry with the National Health Services, 1965-1967", 1967, Cmd 3410.

Shemtov, Noam and Gabison, Garry, The Inventive Step Requirement and the Rise of the AI Machines (January 17, 2022). Queen Mary Law Research Paper No. 375/2022 in Ryan Abbott (ed), Research Handbook on Intellectual Property and Artificial Intelligence (Edward Elgar 2022) , Available at: https://www.elgaronline.com/edcollbook/book/9781800881907/9781800881907.xml