Personalised medicine: the dream and the reality

Marc Beishon

Recent years have seen a scramble to identify the genetic variance that predisposes or protects against certain cancers and the tumour gene signatures that could indicate which therapies will work and which won't. A picture is now emerging of an infinitely complex field that is unlikely ever to live up to the high hopes of some scientists, but is certainly confounding the sceptics.

The lay public must be a bit confused about the term 'personalised medicine', which has become one of the hottest health topics at present, judging by the sheer number of mentions it is getting in the mass media. On face value, it seems to imply that we have arrived at a promised land of individual treatment, certainly where the genetic make-up of people and disease is concerned. After all, the one concept that everyone knows is that we are all – bar identical twins – genetically different from each other. But in fact we are far away from applying many different, individual treatment regimens based on genetic characteristics.

What personalised medicine currently means in practice is treatments or procedures that apply to groups of people, although those groups could be quite small parts of the population. In general medicine, a transfusion of a certain blood type is a 'personalised' approach. In cancer there are many that can be mentioned, such as the Herceptin (trastuzumab) monoclonal antibody for Her2-positive tumours and screening for the BRCA1 and 2 breast cancer genes. A recent technique showing promise allows personalised levels of chemotherapy dosing for colorectal cancer patients based on a blood test, rather than the 'gold standard' of calculating drug dose by body surface area.

It is genomics – and all the other molecular sciences such as proteomics – that are making the running in the personalised medicine debate. One person's tumour is as different from another's as a fingerprint or iris, so it is no wonder that cancer is a prime target for personalisation. There is now a massive body of published papers – more than 50,000 alone on molecular signatures for cancer, for example.

Interest and publications are also growing fast in complex genetic variance that may confer risk (or protection) – hardly surprising as the human genome has now been sequenced. Luminaries such as Craig Venter and James Watson have had their own DNA analysed (in Venter's case this showed he had no known indication for skin cancer – and yet he did develop melanoma).

There has been a lot of hype about the value of both genetic variance and tumour profiling. Helping to bring us down to Earth is a professor who is expert in research methodologies with a special interest in genomics and cancer, and is helping to navigate what is becoming an exponentially complicated field. John Ioannidis, head of the department of hygiene and epidemiology at the University of Ioannina, Greece, made waves with a paper, 'Why most published research findings are false' (*PLoS Medicine*, 2005). He suggested that the hottest scientific fields – with many research teams involved – are less likely to come up with true findings, and there is 'ping pong' between extreme positive and negative results – and no topic illustrates this better than molecular genetics.

Markers of risk

Bearing this 'reality check' in mind, Ioannidis is by no means pessimistic about progress, and has an overview of where we currently are in both genetic

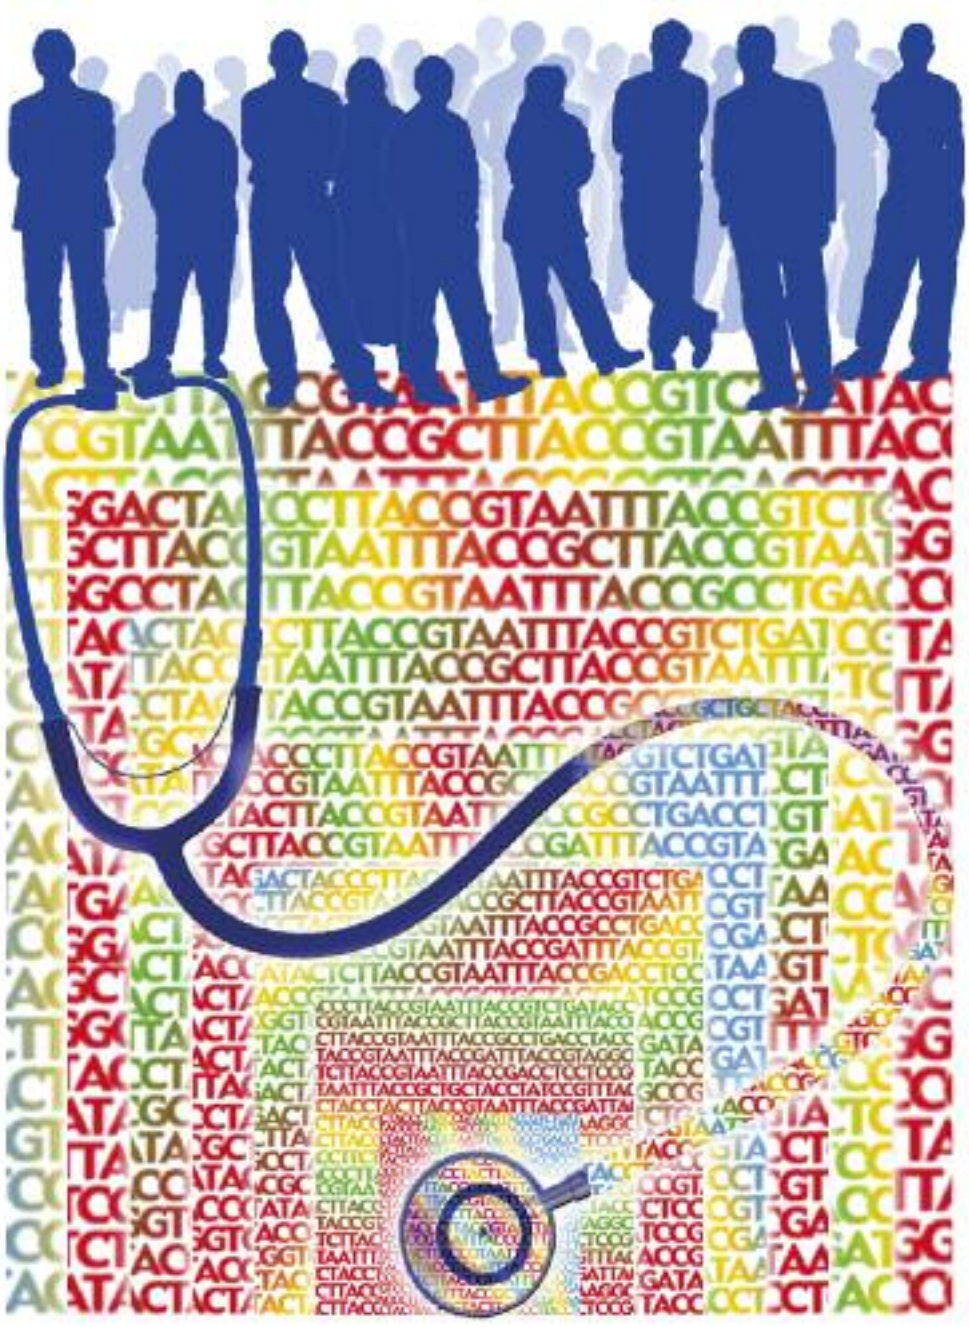

ILLUSTRATION: JANE ADES, NHGRI

variance and tumour profiling. "For genetic variance that confers risk or protects people from cancer, we have now seen identifiable progress for the first time in long while," he says. "Until recently, we had only found syndromes with a high penetrance – meaning there are few genetic factors that could contribute an enormous level of risk, but were very uncommon – but they do not explain why most people get cancer." Well-known syndromes include hereditary non-polyposis colorectal cancer, Li-Fraumeni syndrome, the BRCA1 and BRCA2 genes and so on.

Now, he adds, there are more reproducible findings about genetic variance carried by 10%–30% of people that increase or decrease risk of specific cancers. "We have had most luck with breast and prostate cancer, and found half a dozen common variants, each of which increase or decrease the relative risk of getting these cancers by about a fifth, e.g. from 5% to 6% or 4% in terms of absolute life-time risk." The variants, he says, are single nucleotide polymorphisms (SNPs, also known as 'snips'). "Someone has, say, an A instead of a T in a sequence – but we don't know whether these are the real culprits and are directly increasing risk. They could be mirroring some other genetic site they are linked to. All we really know is we have found markers for some genes – and that's about it." SNPs are found in both coding regions for genes (and so could alter proteins) and non-coding regions (where proximity to a gene can act as a marker). There can be millions of SNPs in each human genome.

"We have found half a dozen common variants, each of which could alter your relative risk by about one fifth"

"We may never be able to answer how bad exactly it is to have two bad variants that interact"

With this level of information, not much individualisation can be done – and the human genome has been screened 'pretty efficiently' now. "At best, we can explain maybe 5%–10% of the variance of the genetic risk of people getting these cancers – just a small amount of the total variability. But other tests – such as the PSA test for prostate cancer carried out on healthy men – are crude measures with low predictive ability, so we may be doing as well and a bit better than before."

Ioannidis says that going beyond identified variants that contribute to single individual risk means looking at many common variants, some of which may contribute an amazingly small increase in risk. "And that is extremely difficult to pick up, even with very large epidemiological studies." Identifying combinations is a daunting task – testing pairs of say a million each is 10^{12} interactions – testing for three is 10^{18}, and pushes even our fastest computers beyond their capability. "We may never be able to answer well the question of how bad exactly it is to have two bad variants that interact."

What the results do offer are opportunities to pursue biological studies. "If we have gene variants with *P*-values of 10^{-20} – and sometimes we are getting to that high level of probability now – we can search for the genes related to the variant and try and understand what goes wrong. There is the possibility they reflect biological pathways and machinery we do not really know about, and maybe we can develop new treatments based on the research – but it is just a promise."

Turning to molecular studies on tumours, it is striking, says Ioannidis, that despite the immense literature built up over the last 10 years, which is more advanced than that on genetic variance, the findings are modest.

For a start, an objective of improving survival by applying molecular signatures has been tempered with more modest aims such as learning how to use chemotherapy more effectively and to minimise unnecessary treatment for those who would not benefit. "We are looking to improve the accuracy of classification for predicting what would happen to patients – where of course the classic application is in breast cancer."

So far, there has been a small gain in predictability, says Ioannidis. "We could classify 65 out of 100 people correctly in the past. Now it is about 70. The number depends on the participation of different subgroups of tumours you have in your sample, so this is a simplified statement, but the gain is modest." That so much emphasis has been placed on the 'hot results' of gene profiling is understandable, he adds, in such a rapidly blossoming field, but there must now be a concerted effort to develop more rigorous methodologies.

"In particular, we must start much larger studies to improve the accuracy of molecular signatures and to reduce the 'noise' level, and it would certainly help if we have more robust validation and replication plans. Until now, most are small studies on a few dozen or maybe a few hundred samples, and validation procedures have sometimes been very shaky."

At least half of papers, in his estimation, have serious methodological problems that could exaggerate their validation performance, especially for cross-validation (i.e. when not using new data). "Many are very complex studies that require a lot of effort in design, analysis and reporting, and are very rich in data. Even minor divergences can produce huge biases."

Markers of response

Two very well known studies that include large sample numbers are the European-based MINDACT trial (Microarray in Node Negative Disease may Avoid ChemoTherapy – see also *Cancer World* 7, July–August 2005), and, in the US, TAILORx (Trial Assigning IndividuaLized Options for Treatment), both of which are looking at molecular signatures in node-negative breast cancer to avoid unnecessary treatment and which are aiming for patient numbers of 6,000 and 10,000 respectively. "They are trying to validate the clinical performance of molecular signatures and I am quite happy if, as a result, we can just improve quality of life and minimise toxic side-effects," says Ioannidis.

A good review article of the progress and practical limitations of this technique, 'Enabling personalized cancer medicine through analysis of gene-expression patterns', has recently been published in *Nature* (3 April 2008) by Laura van 't Veer and René Bernards, pioneers of gene expression profiling who developed the MINDACT signature (see also *Cancer World* 12, May–June 2006).

A feature of the molecular signatures that puzzles many outside the field is that there is little or no commonality between the signatures of identified genes for a particular disease such as breast cancer. As Ioannidis explains, this is probably because there are several genes that have the same effect on the cell cycle, carcinogenesis and the spread of the tumour. Most of these genes are proliferation genes, he says. But having more commonality in the signatures and larger sample sizes should improve accuracy (and also reduce the majority of genes that are just 'noise' in most of the signatures currently used).

So with this 'proof of concept' and with large studies, Ioannidis' guess is that we will be able to increase correct patient classification to about 80–85 out of 100. The gene expression and microarray techniques used in trials such as TAILORx are not the only biological profiles that could show promise, he adds: "Microarrays and gene expression is one level, and other levels include proteomics, metabolomics, epigenetics, nutrigenomics, and instead of measuring messenger RNAs measuring micro RNAs. They can all offer complex biological pictures we can look at, but again there are only pilot studies on very small numbers at present. We need to investigate which techniques are best to take further."

As with genetic variance, his take on what we are likely to see is a small, incremental contribution from adding various techniques – again pointing to the need for large-scale studies of even greater complexity. "Each is a snapshot that may not be independent from another. They should all be pursued."

The TAILORx and MINDACT trials do show that where there is backing from funders, widespread collaboration on large sample sizes for molecular signatures has no barriers. But as Ioannidis also points out, there is irony at present in the way that most of the studies are carried out when compared with work on genetic variance. "The paradigm of molecular signature studies is for one or two centres to do research with a low number of samples, but often investigating a few dozen or maybe several hundred gene expressions. But the genetic variance studies are now screening huge numbers of people for just a single gene variant, or maybe a few."

The genetic variance research did start with relatively few participants, but now there is work on populations of 100,000 or more. "People will not be able to publish in major journals in this field unless they use massive platforms for genetic measurement, and that means massive sample sizes." Consortiums are now looking at variance and susceptibility for diseases such as type 2 diabetes, breast cancer and osteoporosis, with as many as 50,000–150,000 participants, notes Ioannidis, while smaller numbers – but still over ten thousand – are involved in the search for variants for Parkinson's disease.

The good news, he adds, is that researchers are now looking at combining data on molecular signatures, public availability of data is now greatly improved (most databases were not open to researchers only two to three years ago), and there are papers reporting convincing evidence that different labs can achieve similar results when carrying out gene expression profiling.

That may help to create greater consensus about exactly what results we are seeing from studies. Ioannidis has an amusing example of one paper that could be interpreted as very positive to rather negative, and at several points in between – and which one you prefer depends on your expectations. He concludes: "Overall, the best we can currently do is create stratification of risk for certain subgroups – full individualisation is far away."

10,000 HUMAN GENOME PROJECTS

The National Cancer Institute and National Human Genome Research Institute (NHGRI) in the US has started to build The Cancer Genome Atlas, a hugely ambitious project to map genomic changes linked to cancer. Its pilot phase is examining three tumour types, glioblastoma, lung cancer and ovarian cancer. The scale of the challenge is noted in a *Scientific American* article (18 February 2007) by Francis Collins (director of the NHGRI) and Anne Barker: for the 50 most common cancers the effort would be equivalent to 10,000 Human Genome Projects in the volume of DNA to sequence.

"Where there is backing from funders, widespread collaboration on large sample sizes has no barriers"