European Review, Vol. 25, No. 1, 96–106 © 2017 Academia Europæa
doi:10.1017/S106279871600048X

Cancer, Computers and Complexity: Decision Making for the Patient

MARKUS HARZ

Fraunhofer MEVIS, Universitaetsallee 29, D-28359 Bremen, Germany.
E-mail: markus.harz@mevis.fraunhofer.de

In health care, a trend may be noted to fundamentally question some of today's assumptions about the traditional roles of medical disciplines, the doctor–patient relationship, the feasibility of medical studies, and about the role of computers as an aid or replacement of doctors. Diagnostics and therapy decision-making become more complex, and no end is in sight. Amounts of health-related data are being collected individually, and through the health care systems. On the example of breast cancer care, technological advances and societal changes can be observed as they take place concurrently, and patterns and hypotheses emerge that will be the focus of this article. In particular, three key changes are to be considered: (1) the growing appreciation of the uniqueness of diseases and the impact of this notion on the future of evidence-based medicine; (2) the acknowledgment of a 'big data' problem in today's medical practice and science, and the role of computers; and (3) the societal demand for 'P4 medicine' (predictive, preventive, participatory, personalized), and its impact on the roles of doctors and patients.

Introduction

Medicine today faces several paradigmatic changes that are either already underway or likely to occur over the next few years. One main cause is the digital acquisition and storage of patient-specific health-related data. Medical records store the clinical history, medical images can be stored electronically, and patients track their behaviour using mobile devices. Central data storage systems begin to emerge, which are under administrative control of governments, or provided by companies. Legislation is being put into practice that regulates the secure storage and exchange of health-related information as well as data ownership. The crucial insight here is that our position in time marks only the beginning of the medical 'digital data age', as much of the data is not yet acquired for a prolonged time (such as. all X-ray scans and ultrasound images of a patient's lifetime, or the food intake and physical activity data for the lifetime of a teenager). While, in many areas of medicine, the accumulation of

https://doi.org/10.1017/S106279871600048X Published online by Cambridge University Press

data over time will allow interesting new directions to improve personal health and well-being, there is one area in medicine where digital data have existed for many years, and where some of the problems and new approaches can be studied and discussed. It is the field of radiology. It includes image acquisition for the purpose of screening (e.g. active screening programmes for breast cancer, screening recommendations for lung cancer[1] and skin cancer), diagnostic imaging for cancer and other indications, and patient monitoring after treatment. Arguably, cancer is a major cause for acquiring medical images, and imaging technology evolves quickly to ever better depict cancers as early as possible. The oncologist is probably the most frequent customer of radiological services.

Oncology is a medical discipline where the understanding of causes and effects and the choice of treatment options has recently grown into a highly complex picture, where in some parts, research has only begun to scratch the surface of today's mostly evidence- and guideline-based clinical practice of diagnosing and decision making. Reflecting the increase in knowledge and data against the technological advances and societal changes that concurrently take place, patterns and hypotheses emerge, out of which three will be the focus of attention.

The following will take as a practical example the detection, diagnosis and treatment decision-making in breast cancer, this being a disease with a high societal impact and posing challenges to healthcare economics, as well as being a field of very active global scientific research.

Three Theses on the Future of Medicine

Conquering Cancer? Drowning in Complexity

In the past decades, the study of complex systems has become a scientific discipline. Besides a considerable body of literature covering complex systems dynamics from mathematical points of view, interdisciplinary and cross-disciplinary research institutes and journals have emerged. This goes along with a focus on applications rather than single-discipline fundamental science, and seems to signify two observations: first, the importance of a holistic perspective gets more attention. Second, research with a too narrow a focus can no longer solve some of the most interesting scientific challenges.

A system or entity that cannot be described in an algorithm or recipe shall be called complex. Being complex can be opposed to the notions of being simple, complicated, or chaotic;[2] see Figure 1.

Humans can act and react in different patterns when confronted with complex environments. While, for illustration, applying the patterns to the medical field of oncology, we suggest three basic principles that generally govern human action-taking in the presence of complexity.

Ignore. Act as you are used to, following intuition or superstition. In medicine, this might be the case for some areas of alternative medicine where there exists no theoretical or logical reason for their benefit, nor any empirical proof of efficacy. Even the continued employment of the prostate-specific antigen (PSA) test to determine the growth of prostate cancer detection may be due to a simplification of the mechanisms of

https://doi.org/10.1017/S106279871600048X Published online by Cambridge University Press

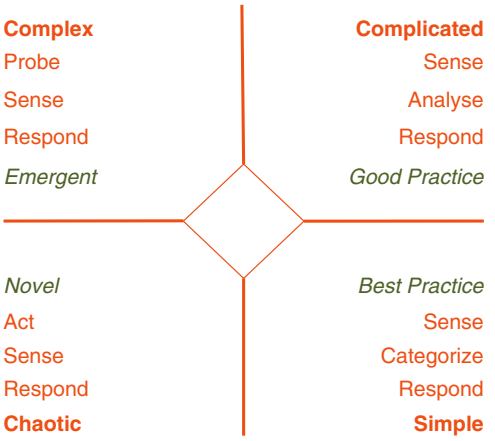


Figure 1. Responding in systems from simple to chaotic. After Ref. 2.

PSA exprimation in the presence of prostate cancer. Women, notably, also produce PSA in the urethral gland.[3]

Simplify. Assess only comprehensible parts of the situation that can be described, logically connected, and be made the basis for a rule set. Typically, in medical practice, guidelines for diagnosis and therapy summarize the consent of a medical discipline. Often enough, controversies arising from scientific differences are resolved by finding the least common denominator. Reimbursement follows the guidelines, and cost considerations limit a doctor's impetus to prescribe imaging or treatment beyond the insurance-covered substances and procedures.

Get help. To grasp the full complexity of medical cause-effect interrelations and to turn this knowledge into actionable advice might be beyond the capacity of humans. Instead of ignoring or simplifying the situation, a third option is to divide the complex system up into building blocks until the granularity is sufficiently fine to model each part independently. From these modules, inferences may be combined into higher-level models, finally rebuilding the complex system from scratch. This approach has been followed in the Human Brain Project, attempting to recreate the brain in silicio to study its complex behaviour. We have reached to date 'a detailed computer representation of about a third of a cubic millimetre of brain tissue containing about 30,000 neurons connected by nearly 40 million synapses'.[4] This has been enabled through computer-based analyses of brain connectivity. Computers are very versatile when it comes to crawling through large amounts of data. Modern methods will briefly be reviewed below (in the Section 'Deep Learning and Cognitive Computing in Medicine'). Computers can, without an expert guiding the way, find meaningful combinations of observations that predict outcomes. This strength became apparent only with the emergence of high-performance computing power, such as IBM's Watson technology. Computers will be ever more helpful as they are provided with more and more data, which helps patterns to emerge. This is one reason that more attention is being paid to patient-owned and patient-provided data.

https://doi.org/10.1017/S106279871600048X Published online by Cambridge University Press

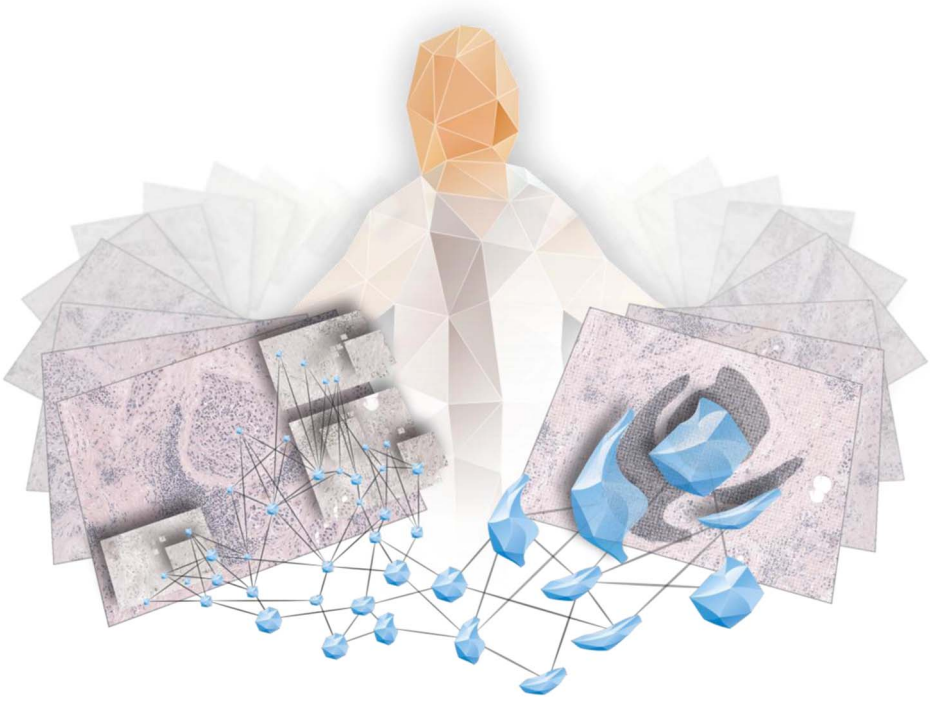


Figure 2. Doctors supervise self-learning algorithms. Deep learning uses many-layered neural network architectures, in which the low layer (here: left in the image) finds salient low-level image features, while higher layers correspond to concepts.

The complexity of cancer, the hallmarks leading to its genesis,[5] its diversity of genotypes and associated phenotypes, increasingly reveals itself through genetic profiling and various clinical and non-clinical imaging methods. In this light, no two cancers are the same. Consequently, they shouldn't be treated the same. A wide variety of treatment options has emerged, evolved and diversified. This furthers a clinical decision-making problem: doctors obtain extremely comprehensive information about patient demographics, the disease and its co-morbidities, therapeutic strategies and their impact and likelihood of outcomes. Summarizing, there is too much information to consider, and too many options to choose from.

Our data acquisition tools already allow us to measure more than can be clinically assessed, with efforts to automate the acquisition processes further. Researchers are beginning to face a 'big data' problem, as many of the tasks needed for information extraction and sense-making can't be automated to a matching degree, leaving the tasks of information extraction and sense-making to the human observers. We create a historically unparalleled wealth of multi-scale, multi-level, and multi-disciplinary information that cannot realistically be described by conventional statistical analyses. Take as one example among several the existing collection of cancer genomic data that is being collected by the NCIs Cancer Genome Atlas project: 25,000 TB of data have been collected so far, and little of it has been tied to clinical information.[6]

https://doi.org/10.1017/S106279871600048X Published online by Cambridge University Press

One possible human reaction facing this complexity is to resort to simplifications and categorizations based on reduced information, or, in the worst case, on subjective opinions or best practices. In the case of breast imaging, for example, reporting of imaging findings follows a catalogue of predefined terms. This not only reduces the continuum of potential findings to nominal categories, but also isn't open for non-standardized findings. Another example is decision-making for chemotherapy. There is a limited number of potential drugs that can be combined into a limited number of therapy regimen. The choice of regimen is taken based on about four parameters derived from the huge and extremely feature-rich histopathology images, and these parameters are given in categories ('low – medium – high') rather than as a number. The choice of drugs is due to national and international guidelines, which represent the minimal consensus that could be reached between competing scientists worldwide.[7]

Empowering the Patients

In parallel, societal and legislative changes influence (and are influenced by) the 'big data' situation. Health data are personal, and hence it is a natural demand that such data should be owned by the patient. This is not the practice today, however. Instead, the value of data is increasingly acknowledged by solution providers in the medical area. This applies to well-known players in the field, but in particular small start-up companies have started offering services that enrich data with semantic information or provide data in return for a fee. This addresses the needs of solution providers who want to aid clinical decision-makers. A new branch of industry has developed, focusing on data-related services. It is remarkable that, while in the past, a few very thoroughly understood and curated data items bore a high value, today the greatest hopes are instead put into the largest possible amounts of raw data, assuming that, in total, such data exhibit patterns that can be detected automatically. This has multiple consequences for the patient. First, curating data is no longer the prime concern; raw data are sufficient for many purposes. Second, one patient case alone loses value. Consider this in the context of large-scale clinical trials for drug or medical device approval. Typically, companies conducting such studies struggle with the inclusion of sufficiently high numbers of patients, and have to set aside large amounts of money to finish the trials. In this scenario, each patient case counts, and each case is only valuable when it is properly documented and followed up on. In opposition, in the era of 'big data', numbers count. Since patients don't necessarily (or do not willingly) reveal much personal information, awareness of what personal data go where is low. But seeing a physician or visiting the hospital creates data, sometimes images along with reports, sometimes clinical information. These data are digitally available, and patients are at any time free to request them.

This is exploited by companies offering web-based or personal data collecting services, including the health infrastructures set up by companies such as Google, Microsoft and Apple. The promise is the easy and persistent storage of health-related information, and the monitoring of personal health indicators. The platforms have one thing in common: in return for the services they offer, the terms of service demand a

https://doi.org/10.1017/S106279871600048X Published online by Cambridge University Press

certain usage right on the data. The same is true for patient-oriented web communities, such as the very popular patientslikeme.com. For the patient, this amounts to giving up a level of control of the data, and perhaps participating in the profits derived from the data too little.[8] The underlying reason is that the permission to use the data is not given on an on-demand basis, but wavered unconditionally under terms of use. Providers of hospital information systems have hopped on the same train, offering patient-oriented web-based data access and services on the data the hospital stores of them. While this is certainly in the interest of both the patient and the health-care providers and simplifies the exchange of patient data, for example between the hospital and the general physician, again the service comes with strings attached.

Electronic Health Records (EHR) stored or controlled by the patients themselves offer an opportunity to prevent single-sided bargains. In this scenario, each use of his or her data has to be acknowledged by the patient. The health data exchange infrastructure of course becomes more vulnerable, since with freedom of choice (of system and service providers) there comes the need of interoperability, interfaces, openness and transparency. However, with proper provisioning of secure access and exchange systems, safety concerns can be overcome,[9,10] but questions concerning medical liability and responsibility, and psychological effects on the personal level remain unanswered. The concern is, how will services offered for patient-centric data analysis be certified, accredited, approved and controlled? The services may cause medical errors, such as misinterpretation of the data, wrong predictions based on the data, or wrong diagnoses leading to harmful therapy choices. The reasons can be manifold: misleading, incomplete, or insufficient data; services that alone or in combination misinterpret the data, or combinations of data and service limitations. Liability consequently becomes a multi-faceted issue.

Further ethical questions arise when envisioning the medical histories of millions of humans to be jointly analysed for patterns, striving for precision medicine.

Conversely, in some clinical disciplines, the integrity of the doctor is in question when patients are granted by law the right to see all notes and documents created about them. Perhaps the most striking example is that of clinical psychology and psychotherapy. As in some forms of psychotherapy, the focus is on the relationship between patient and therapist, so the notes not only entail the assessment of the patient's psychological state, but also the therapist's, revealed through self-inspection. Therefore, if revealed to the patient, the personal rights of the therapist are potentially violated. A delicate balance between the rights of health care professionals and patients should protect both sides.

Further, with the possibility to relate an individual's aggregated health data with the whole population, health insurance companies might start to offer individual benefits if in return their data prove a healthy lifestyle. In a pessimistic scenario, self-monitoring might even be made a prerequisite to obtain coverage by a health insurance altogether, and be made the criterion upon which insurance companies calculate their dues.

So far, we have described the effects of large amounts of digital health-related data, and projected the future changes this might cause. The patient appeared as a data

https://doi.org/10.1017/S106279871600048X Published online by Cambridge University Press

provider, augmenting a global, population-wide dataset from which a plenitude of parameters, information, and knowledge can be derived. In yet a different light, the individual patient will potentially profit directly from such a data pool. This is because, when we acknowledge the deeply individual nature of disease, the future of evidence-based medicine might be at stake. This paradigm assumes that all influential covariates of a disease outcome can be controlled, so that the influence of an intervention (for example, a drug) can be singled out to prove its effectiveness. Patients are thus categorized. Diseases are likewise stratified according to measurable criteria, and usually into only a few categories. For breast cancer, only a handful of distinct groups are counted. Still it is obvious that cancer is much more diverse, and the environment – the local stroma, and the life circumstances of the patient – has a notable influence on the outcome. Rightfully, patients ask for personalized treatment plans and drugs, and may become the driving force advancing population-wide data analysis. However, the purely technical barriers to overcome to share the data across centres, cities, countries and, ultimately, the world, are today unsolved or come at excessive integration costs – let alone the ethical issues mentioned before. Sharing, however, is already a necessity in today's setup of phased clinical studies; for example, to collect sufficient numbers of cases showing the genetic mutation of interest[6] in the assessment of a new drug. Sharing will be even more important to support personalized drug design by identifying matching individuals, and it has, by the same mechanism, an application in prevention, since bearers of similar traits can be identified before the outbreak of disease, once the fingerprint of the disease is known.

Precision medicine for the patient's benefit is promising[11] but costly, so who is going to cover these costs? Interestingly, simple health-care cost estimates with only few assumptions already point out a potential for economic benefits with simultaneously improved outcomes. For example, in breast cancer, the potential saving in chemotherapy costs would by far outweigh the development and integration of an imaging biomarker-based test that determines the effectiveness of the chemotherapy treatment after the second cycle. For example, the change of diffusivity of the tissue in and around the tumour seen in MRI before and into the treatment is known to be a reliable predictor of treatment success. As, in some subgroups, there are no or only a few predictors ahead of the start of treatment, the search continues. This is why it is possible that pharmaceutical companies become the primary stakeholders in such developments. They benefit from large-scale data pools to assess their drugs' efficacy. Using large-scale data aggregation, striving for the 'right', the optimal and personalized drug can be pursued. Areas of salient future research can be identified.

Then, the question remains, how will the empowered, participating patient affect medicine's quality through his/her informational self-management. Will a popular medicine trend emerge, and how will that impact the profession of the doctor?[12]

Deep Learning and Cognitive Computing in Medicine

There are many clinical disciplines potentially involved in the treatment of a breast cancer. There is the radiologist, who creates the images for screening, for diagnosis, and

https://doi.org/10.1017/S106279871600048X Published online by Cambridge University Press

for treatment planning. Radiologists usually don't start to work on their own behalf (except for screening), but are ordered by another discipline, such as the oncologist, or the surgeon. In addition, although they suggest a diagnosis conforming with the images, it will be the pathologist who determines the final diagnosis. In the scenario of empowered patients and shared data, this picture might change. A prerequisite could turn out to be the development of machine learning of a new kind.

The confluence of heavily parallelized computer architectures and ingenious algorithmic developments in the field of machine learning has recently enabled a key technology advancement closely linked to 'big data'. The community of machine learning algorithm developers are creating algorithms that perform at the same level as humans, and in several scenarios already even better than human. A very prominent competition has been held over several years, where computers are to say what is seen in an image, and to segment it. More than 14 million images have been collected; and in a huge effort, the depicted object of interest has been annotated in each single one of them.[13] After some stagnation, a recent development[14] boosted the algorithm performance to superhuman accuracy in telling apart even the most subtle differences between objects.

Meanwhile, huge international projects strive to rebuild the architecture of the human brain in silicio. While some pursue the bottom-up approach, depicting the brain in the finest possible resolution, others emulate some of the brain cell functions on a more abstract level. Of those, some are algorithmic models of neurons, others are hardware designs inspired by a similarly abstracted idea of brain function. IBM, for example, markets a chip architecture modelled on neurons connected by synapses. This design is much more energy efficient than software-based emulation. On the other hand, the design is tightly intertwined with the computations intended to run on it. When the paradigm changes (and this happens frequently in this field), the architecture is inflexible. This approach is consequently critiqued on the technical level. It is largely undoubted, however, that computational or hardware emulations of neural networks are very promising approaches to solve all kinds of cognitive tasks.

The deepest and largest of such artificial neural networks still only achieves the cognitive performance level of a small rodent brain (but at the price of orders of magnitude more energy consumption), but given the international efforts and the fast development of high-performance computing the technical basis for ever deeper and wider networks has been laid, and novel developments are announced almost daily. An accompanying growth in data volume is a necessity to train such networks, as their degrees of freedom (the weights to learn between the layers) grows with each neuron that is introduced.

Being one paradigm within the field of hypothesis-free statistical analysis, 'deep learning' uses a method shaped on learning patterns observed in the human brain.[15,16] 'Deep' here refers to the number of network layers stacked 'above' each other. Each layer has a number of neurons, which have an activation function that is usually nonlinear. Between the layers of the networks, weights determine how the activation transports from one neuron to another. The weights hence resemble the synapses. Learning the weights is the task accomplished during the training of the network.

https://doi.org/10.1017/S106279871600048X Published online by Cambridge University Press

In the past few years, the complexity (the number of the connections with weights) of the networks rose at largely constant rate.

On the other hand, several authors have shown how deep architectures are prone to overfitting even in data-rich scenarios. Hence, it has been proposed to reduce the complexity of the deep architecture after training. This is possible in several ways, of which the most intuitive one probably is to reduce the number of connections (and consequently also neurons) based on their contributions to the decision. This is known under the term 'network pruning', and has been proven to enhance the generalization ability of the network.

Another approach to drastically reduce the network complexity is by creating an adjunct shallow network architecture (or any other kind of learning method) that observes the input, weights, output and desired output of the deep network when it is presented with testing examples. The shallow net will, in such a setting, not be able recreate the empiric model. Instead, by its much smaller number of degrees of freedom, it will be forced to learn the underlying function that is modelled by the deep architecture's weights.

Only very recently it turned out that, by reducing the number of weights between layers but increasing the number of layers drastically, a huge boost in performance can be achieved. Interestingly, significantly fewer parameters are now required, and the network architecture is also only composed of very simple building blocks stacked over and over again until the desired depth and performance is obtained.

The ultimate promise of knowledge-based medicine is precision medicine: the optimization of diagnostic tools and therapeutic choices for individuals.[17–19] This is based on an objective and holistic data analysis, resulting in personalized, intuitive visualizations and comprehensive information use – something only automated computer analyses might be able to provide.

With a powerful cognitive computer at his or her side, the clinical practitioner achieves new prospects. It is possible to substitute parts of the expertise of many clinical disciplines involved in cancer care with cognitive computing. Most prominently, image interpretation can be accomplished by a computer. This puts the existence of radiology as a diagnostic discipline into question. Cognitive computing can in principle do the same for digital pathology image interpretation, and generate the pathology report for a biopsy specimen automatically. In addition, selecting the appropriate drug or designing the fitting treatment plan from the clinical information of the patient can strongly be supported and supervised by a decision-making computer program. Two direct consequences spring to mind. First, for patients, a service might become available that gives them computer-generated second opinions on their data. Second, the separation of clinical disciplines (radiology, pathology, oncology) may diminish, and the discipline of general breast diagnostics may emerge: a clinician knowing a lot about all the disciplines, strongly supported by a computer to assure no details are missed.

Together, data complexity, patient empowerment, and computing power close a loop: the hope to disentangle the patterns hidden in the data lies in collecting heaps of it, and analysing it with state-of-the-art technology. Collecting enough data will be possible by empowering the patient, who may in return benefit from the analyses of

https://doi.org/10.1017/S106279871600048X Published online by Cambridge University Press

his/her personal data – automated, reliable, on demand for the patient, and cutting health care costs for society.

Conclusions

We make the case for 'P4 medicine': predictive, preventive, personalized and participatory. As some put it: 'biology drives technology drives computation'.[20] As the biology of cancer begins to reveal itself through technologies invented over the past decades, data-driven computational tools have to be devised and understood to cope with the breadth and wealth of knowledge and information. On the legislative, ethical, economic, and psychological levels, we have pointed out pressure points that need to be addressed by the involved interest groups together with the respective legislative bodies and a cross-disciplinary board of scientists. Shaping the future of medicine is, seen in this light, a true societal task. The synergy of humans and technology manifests itself when we accept that our minds are not limitless, and that we need the power of technology to help us grasp and understand the wealth of information present in our data. This might enable us to use our human creativity to the better of patients.

References and Notes

1. V.A. Moyer (2014) U.S. preventive services task force: screening for lung cancer: U.S. Preventive Services Task Force Recommendation Statement. *Annals of Internal Medicine*, **160**, pp. 330–338.
2. D. Snowden (2000) Cynefin: a sense of time and space, the social ecology of knowledge management. In: C. Despres and D. Chauvel (Eds.), *Knowledge Horizons: The Present and the Promise of Knowledge Management* (UK: Butterworth Heinemann).
3. M. Zaviačič and R.J. Ablin (1998) The female prostate. *Journal of the National Cancer Institute*, **90**(9), p. 713.
4. H. Markram, E. Muller, S. Ramaswamy, M.W. Reimann, M. Abdellah, C.A. Sanchez, A. Ailamaki, L. Alonso-Nanclares, N. Antille, S. Arsever, G.A.A. Kahou, T.K. Berger, A. Bilgili, N. Buncic, A. Chalimourda, G. Chindemi, J.D. Courcol, F. Delalondre, V. Delattre, S. Druckmann, R. Dumusc, J. Dynes, S. Eilemann, E. Gal, M.E. Gevaert, J.P. Ghobril, A. Gidon, J.W. Graham, A. Gupta, V. Haenel, E. Hay, T. Heinis, J.B. Hernando, M. Hines, L. Kanari, D. Keller, J. Kenyon, G. Khazen, Y. Kim, J.G. King, Z. Kisvarday, P. Kumbhar, S. Lasserre, J.V. Le Bé, B. R.C. Magalhães, A. Merchán-Pérez, J. Meystre, B.R. Morrice, J. Muller, A. Muñoz-Céspedes, S. Muralidhar, K. Muthurasa, D. Nachbaur, T.H. Newton, M. Nolte, A. Ovcharenko, J. Palacios, L. Pastor, R. Perin, R. Ran-jan, I. Riachi, J.R. Rodríguez, J.L. Riquelme, C. Rössert, K. Sfyrakis, Y. Shi, J.C. Shillcock, G. Silberberg, R. Silva, F. Tauheed, M. Telefont, M. Toledo-Rodriguez, T. Tränkler, W. Van Geit, J.V. Díaz, R. Walker, Y. Wang, S.M. Zaninetta, J. DeFelipe, S.L. Hill, I. Segev and F. Schürmann (2015) Reconstruction and simulation of neocortical microcircuitry. *Cell*, **163**(2), pp. 456–492.
5. D. Hanahan and R.A. Weinberg (2011) Hallmarks of cancer: the next generation. *Cell*, **144**(5), pp. 29–29.

https://doi.org/10.1017/S106279871600048X Published online by Cambridge University Press

6. M.A. Rubin (2015) Make precision medicine work for cancer care. *Nature*, **520**, pp. 290–291.
7. A. Goldhirsch, E.P. Winer, A.S. Coates, R.D. Gelber, M. Piccart-Gebhart, B. Thürlimann and H. J. Senn (2013) Panel members: personalizing the treatment of women with early breast cancer: highlights of the St Gallen International Expert Consensus on the Primary Therapy of Early Breast Cancer 2013. *Annals of Oncology: Official Journal of the European Society for Medical Oncology / ESMO*, **24**(9), pp. 2206–2223.
8. For example, patientslikeme.com very plainly states in their FAQ, that selling the health-related information is the primary business model. Compare https://support.patientslikeme.com/.
9. U. Hübner, G. Schulte, B. Sellemann, M. Quade, T. Rottmann, M. Fenske, N. Egbert, R. Kuhlisch and O. Rienhoff (2015) Evaluating a proof-of-concept approach of the German health telematics infrastructure in the context of discharge management. *Proceedings of the 15th World Congress on Health and Biomedical Informatics* (MEDINFO 2015), Sao Paulo Brazil. In *Studies in Health Technology and Informatics*, 2015, **216**, pp. 492–496.
10. B.K. Atchinson and D.M. Fox (1997) The politics of the health insurance portability and accountability act. *Health Affairs*, **16**(3), pp. 146–150.
11. H. Beltran, K. Eng and J.M. Mosquera *et al*. (2015) Whole-exome sequencing of metastatic cancer and biomarkers of treatment response. *JAMA Oncology*, **1**(4), pp. 466–474.
12. E. Topol (2012) *The Creative Destruction of Medicine* (New York: Basic Books).
13. Compare http://www.image-net.org.
14. K. He, X. Zhang, S. Ren and J. Sun (2015) Deep residual learning for image recognition. arXiv.org (December).
15. G.E. Hinton, S. Osindero and Y.W. The (2006) A fast learning algorithm for deep belief nets. *Neural Computation*, **18**(7), pp. 1527–1554.
16. Y. Bengio, P. Lamblin, D. Popovici and H. Larochelle (2007) Greedy layer-wise training of deep networks. In: B. Schoelkopf, J. Platt and T. Hofmann (Eds.), *Advances in Neural Information Processing Systems 19: Proceedings of the 2006 Conference* (Cambridge, MA: MIT Press) pp. 153–160.
17. C.M. Kelly and K.I. Pritchard (2012) Personalized medicine: What exactly is it and can we truly measure it? *Journal of Clinical Oncology : Official Journal of the American Society of Clinical Oncology*, **30**(18), pp. 2173–2174.
18. National Institute of Health: Precision medicine initiative (2015).
19. Global Siemens Healthcare Headquarters: syngo.via frontier (2015).
20. L. Hood (2013) Systems biology and p4 medicine: past, present, and future. *Rambam Maimonides Medical Journal*, **4**(2).

About the Author

Markus Harz is a computer scientist with ten years of experience in image analysis for medicine. He has managed several international research projects in the field of breast cancer diagnosis and therapy. A stay of six months at a breast cancer centre in Florida, USA, added invaluable practical insights. In his PhD thesis at the University of Bremen, Harz proposed machine learning methods to support clinical decision making in complex tasks. He continues to focus on computer learning methods to automate specific detection, segmentation, and decision-making tasks in medicine.

https://doi.org/10.1017/S106279871600048X Published online by Cambridge University Press