University of Zurich[UZH]

Digital Society Initiative

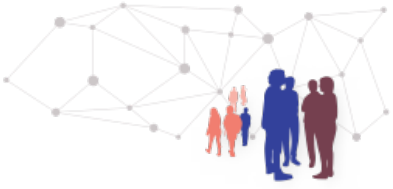

Artificial intelligence in medicine - Case Study «Therapy»

DSI Strategy Lab 2022

Editorial team: Michael Krauthammer, Nikola Biller-Andorno, Markus Christen, Jeffrey David Iqbal, Christian Kauth, Viktor Kölzer, Tania Krones and Claudia Witt.

This case study was created in a systematic and participatory process as part of the «Strategy Lab Artificial Intelligence of Medicine» of the Digital Society Initiative (DSI) at the University of Zurich. The process, which involved various stakeholders, extended from the beginning of 2022 to mid-2023. In particular, a workshop with experts in June 2022, a workshop with doctoral students of the DSI Excellence Program in August 2022, and the testing of excerpts of the case study in various events with stakeholders were decisive for the creation of the case study.

Participants in the expert workshop, in addition to the editorial team, were: Abraham Bernstein, Daniel Eberli, Philipp Fürnstahl, Sven Hirsch, Christian Kauth, Emanuela Keller, Birgit Kleim, Tanja Krones, Titus Neupert, Cristina Rossi, Bernd Stadlinger, Florent Thouvenin and Andreas Wicki. Participants of the PhD workshop were Anais Aeschlimann, Ibrahim Al Hazwani, Joe Baumann, Giulia Frascaria, Marius Furter, Alexandra Ioana Georgescu, Maël Kubli, Alexander Lill, Eanuele Martinelli, Judit Martínez Moreno, Matteo, Micol, Markos Mpadanes, Kimon Papadopoulos, Amina Saleh, Jana Sedlakova, Kateryna Shapovalova, Lukas Tribelhorn, Morley James Weston, Basak Yalman, Federica Zavattaro and Donatella Zingaro. We thank all of these individuals is for contributing to this process; in particular, Jeffrey David Iqbal for providing operational guidance throughout the Strategy Lab process.

Further information: dsi.uzh.ch/strategy-lab.

Structure and use of the case studies

The case studies evolve along a temporal progression with a parallel increase in the level of autonomy of the AI application, as follows: the case study starts in the (fictional) «now» using AI with a very limited level of autonomy. It progresses to the «near future» with greater involvement of AI in the decision-making process. Finally, in a «far future scenario», it will be played out that an AI makes relevant decisions largely autonomously (as an «avatar» or «digital twin»). A precise temporal allocation of the scenarios is not attempted. The scenarios are fictitious in nature and were framed rather as «positive utopias» than «negative dystopias», although several aspects of the case studies are intended to be provocative and they should trigger ethical debates regarding the desirability of the scenarios as outlined here.

Comparatively simple language has been chosen; technical terminology has been avoided as far as possible so that lay people can understand the case studies. Factual statements are referenced only minimally (the scenarios are, after all, fictional in nature), supported by few references that are as

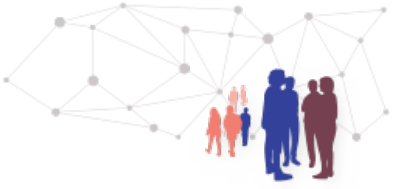

accessible to a general audience as possible. When creating the scenarios, care was taken to ensure that there were connections in terms of content between the fictional case stories.

The case studies will be available «open access» with the intention that they can be used in educational settings and/or workshops discussing AI applications in medicine. The keywords indicate topics that characterize the case studies and should help to decide upon their use in educational settings. Suggested questions after each «step» further guide their use in workshops and discussions.

Keywords: Cancer, Digital Twin, Health Data, Social Media Data, Therapy, Wearables

1 Introductory remarks

Cancer is the leading cause of death in Switzerland, with a sharp increase in incidence among people over 50 years of age [1]. While cancer has been known to be a disease with high morbidity and mortality, recent decades have led to the development of personalized treatment approaches that improve the outlook for many patients. These developments are related to advances in surgery, radiotherapy and targeted treatments tailored to each person suffering from cancer [2] including immunotherapy that uses the patient's own immune system to fight the disease [3]. Cure prevalence across all patients diagnosed with cancer is currently over 65% for men and over 75% for women.

These personalized therapies are part of new developments in cancer treatment in which the molecular nature of a patient's cancer (i.e., the changes in the genetic code that cause a cell to behave incorrectly) can be decoded and used to make treatment decisions. These decisions are usually made in what is known as a «tumor board», a panel that brings together experts from oncology, pathology, radiology, etc., to discuss a cancer case and make an optimal decision. As deciphering a patient's cancer generates a large amount of data to be processed by the experts on the tumor board, computers and artificial intelligence (AI) can play a crucial role in supporting the decision-making process.

This hypothetical case study will focus on how AI systems can support and may ultimately replace such tumor boards in the search for an optimal therapy. The case study does not exist in this form, all names of persons and companies are fictitious.

2 Case study level «Now»

2.1 Situating the scenario

The understanding of what constitutes «cancer» is changing rapidly. It is becoming more and more apparent that the «molecular profile» of a particular cancer, i.e., the DNA changes in the genetic material of the cancer cells, as well as changes at the protein level can be shared across different cancer subtypes and can thus be targeted by a disease agnostic treatment. Computers are already indispensable for processing such information. Moreover, the digitized data from both molecular analysis and the electronic health record of patients can be exchanged and compiled into databases of cancer signatures.

The cancer consortium called CanCon is an association of the tumor boards of the largest university hospitals in a fictitious country. The goal of CanCon is to build a «learning healthcare system» in which data is continuously exchanged, processed and analyzed, with constant feedback: Whenever a cancer patient is treated, the means and outcome of the treatment should be digitally recorded and linked to the cancer signatures database. CanCon has succeeded in getting some university hospitals to agree

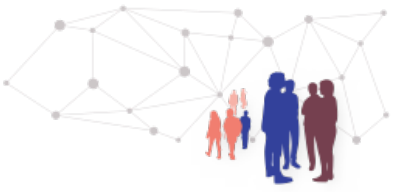

on common standards for sharing and handling data, and the General Consent system ensures that patients' consent to treatment includes consent to anonymized use of all data obtained through diagnosis and treatment. In collaboration with software engineers and programmers, AI algorithms are being developed to analyze these large and ever-growing amounts of data with the goal of finding the best possible treatment for each person with their specific cancer signature. Currently, CanCon is trying to involve more hospitals into the consortium to increase the size of the database available for AI.

2.2 How it could work

A key figure in our story is Carmen, a young oncologist who has only recently taken up her post at a renowned university hospital and is now participating in the tumor boards there. It became clear to her during her training that much is changing in medicine – and especially in oncology. Carmen is part of the first generation of physicians where big data analysis is a natural part of training and she knows how to consult cancer databases to match a cancer signature with the optimal treatment guideline. She firmly believes that novel molecular biology methods will make cancer more and more a curable disease, if only enough investment is made in technological advances. Therefore, she immediately volunteered to be a CanCon «ambassador» at the clinic and regularly reports on the progress of CanCon development at the tumor board.

The pathology department of the University Hospital has driven this medical progress and has been routinely recording the molecular biological signatures of cancer cases. The hospital's own database already contains many ten thousand cases of such signatures, including data on the prognosis of the patients concerned. Carmen has been advocating for some time that this database be merged with CanCon's much larger resources – but privacy concerns as well as technical standardization problems have so far prevented that. Now there is finally a ray of hope: legal clarifications have shown that data integration is possible, subject to defined safeguards, and the technical obstacles can also be overcome; the hospital management has just approved the necessary resources. Nevertheless, experience shows that it will certainly take another year until the integration of the clinic data in CanCon is completed and the AI algorithms of the consortium can use this data to optimize the search for therapy options. Until then, the tumor board will continue to rely on the clinic's own data when decisions become necessary.

Such a decision is now necessary in the case of Sebastian, who suffers from a serious type of skin cancer (melanoma). The molecular biological signature points to a rare variant for which there are only a few analogous examples in the clinic's own database so far – and in all those cases the therapeutic measures taken at the time have had little effect. A look at the literature provides clues to new alternatives – but here, too, the evidence base is thin. Carmen – as a vigorous advocate of modern AI methods in oncology – now advocates bringing forward the integration of the clinic's melanoma dataset into CanCon with the hope that the integral dataset will allow better individualization of therapy through CanCon algorithms. However, head physician Kaiser cautions that the processes for data customization have already been defined and that planning cannot now simply be redone for an individual patient. However, Carmen insists and wants to be present at the next conversation with Sebastian Krug.

Therefore, she sits between the two during the next consultation, which involves discussing the tumor board's recommendations Sebastian. Head physician Kaiser explains Sebastian that he suffers from a rare variant of melanoma with which there is hardly any experience. The treatment outcomes at the clinic are discouraging – and in the literature, there were only few other cases, everywhere with only moderate success. Sebastian asks only a few questions about Kaiser's remarks. Then Carmen speaks up. She describes to him the great advances in the field of oncology and the increasing use of artificial intelligence in determining possible individualized therapies. Full of enthusiasm, she even explains to

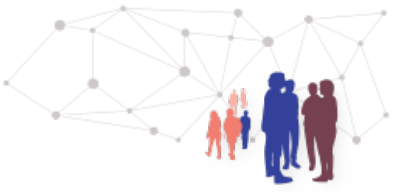

him the upcoming integration of clinical data into the CanCon network, which might offer even better opportunities. Kaiser did not like hearing that – after all, they had clearly told her in the tumor board that accelerated integration of the melanoma data was out of the question.

Sebastian interrupts Carmen's flow of words: «Are you really sure that something like 'artificial intelligence' can find a better therapy than all these doctors? I read about 'artificial intelligence' all the time in the newspapers – but does it really work? There was that autonomous car that ran someone over, right...?» sais Sebastian. «And anyway, how does this artificial intelligence even know what I want, what's important to me? That is all too technical for me. It is more important to me that I can trust the doctor who treats me; that I can talk to him. The fact that, statistically speaking, one therapy is a few percent better than another, is to be honest, all the same to me...». Carmen looks somewhat sobered in the face of these words. «And how do we want to deal with this situation now, what would be your preferences regarding a possible therapy?» asks Kaiser. «Well», Sebastian replies, «if I understood both of you correctly, all existing therapeutic approaches give little cause for hope. In that case, I would like to choose the option in which I can communicate with you and the nursing staff as human beings. I do not know how I will feel in a month. But I want to be sure that someone is always there to whom I can communicate my fears and concerns – and who will then also adjust the therapy accordingly.»

2.3 Possible questions for discussions

Do you agree with the use of data and AI in the context of the tumor board? Are we technically able to do that today? If not, what needs to be changed to make it possible? Is this something we would already be allowed to do under the current legal framework? If not, what should be changed? Would you also speed up the integration of clinic data into CanCon? What do you think of the behavior of Carmen, Kaiser, and Sebastian?

3 Case Study Level «Near Future»

3.1 Situating the scenario

It has been many years since the CanCon project was completed. Lauded as a precursor to comprehensive digitization of oncology, the project has resulted in the former tumor boards becoming digital centers where patient data is continuously shared and analyzed for optimal treatment decisions. In the meantime, the problem of data sharing has been largely solved at the national and international level, especially since harmonization of cancer data allows seamless comparison of data, treatments and patient outcomes in cancer. A wide range of data is collected: in addition to genetic signatures, information on the survival of those affected, their quality of life after treatment including psychological factors, the exact circumstances of treatment, but also socio-economic factors that affect cancer risks as well are recorded.

The AI systems themselves have also undergone a rapid technological development. Techniques such as federated learning have become widespread. They allow AI models to be trained on multiple devices, each with its own local dataset that is not shared with other participants. This technique has made many privacy fears obsolete. Intriguing image generation techniques have evolved so that AI systems can generate meaningful illustrations and movies for all forms of medical interventions, which patients can generate themselves to better understand what to expect.

The committees in the clinics that used to be called «tumor boards» have also changed. It is a network of healthcare professionals from the hospital and external parties who subject the computer-generated

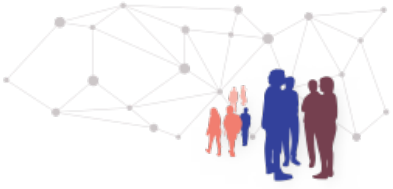

therapy recommendations to a second assessment in an online session – the current standard procedure based on guidelines issued by the professional medical societies that require a human expert committee to always give an assessment of the AI therapy recommendations. These network discussions, abbreviated as «CanThS» (Cancer Therapy Sessions), are therefore also called «Kant» in hospital slang, as a reminder of the philosopher for whom human autonomy was so important. After all, the original fears of loss of autonomy of patients and the medical profession have not come true. Instead of using guidelines based solely on clinical trials, the current data integration and AI curation allows for individualized guidelines. They are still discussed by the CanThS very much in a way similar to the discussion of old in the tumor board.

An important input for these bodies comes from a new profession called «H Coaches»: to each patient now such a «health coach» is assigned. This new profession is located between information technology, psychology and medicine. Their task is to work with patients to ensure that the diverse data collected by patients themselves (for example, through their smartphones) is of sufficient quality and standardized. This can be very simple things – such as making sure that people also operate the wearables correctly – but also includes structured conversations, for example, to record soft factors such as individual preferences in as standardized a way as possible. In a sense, H Coaches ensure that the interfaces between patients, doctors and AI systems (including the software engineers who develop the systems) function as smoothly as possible. Many people appreciate dealing with H Coaches and, overall, the introduction of this professional group was important for the healthcare labor market, as a number of medical tasks – such as differential diagnosis or even psychiatric therapy discussions (think, for example, of the YBDT app, case study «Allocation») – were largely taken over by AI systems in the wake of digitization.

In view of the success of digitized diagnostics in other areas of medicine – think, for example, of the success of the interactive cancer prevention file ICPR (case study «Prevention») – a even more radical change is emerging here: more and more research is being conducted on how to increase the degree of autonomy of wearable devices, for example, so that instructions for their correct use are no longer required. Alternatively, instead of structured conversations with patients, some researchers tries to determine human preferences directly from behavioral data. All this could mean that the H Coaches would no longer be needed as «translators»; after all, they are also a relevant cost factor.

This points to an emerging fault line in society: One part of the society has no problem using wearables or recording their daily activities at home using environmental sensors. This gives patients more opportunities to access relevant data themselves – and technology companies support them by offering them tailored health information. However, other people are skeptical of this approach, being afraid to monitor their individual parameters on a daily basis, which might cause health worries. They swear by H Coaches and the associated human contact, while others even completely refuse AI-assisted medicine. It thus remains controversial to what extent medical care must adapt to data-driven medicine.

3.2 How it could work

Carmen leans back, relaxed. She has just finished her last «Kant»; a check of a therapy recommendation for Pia, who has a hereditary breast cancer – one of those cases that can be easily recognized but where preventive measures such as a prophylactic mastectomy represent a very difficult and life-changing personal decision. Carmen particularly likes the conversations with Florian, Pia's H Coach with a lot of life experience. Carmen once learned that he had once been severely depressed himself (see case study «Allocation»). After a solid computer science education, he had undergone further training to become

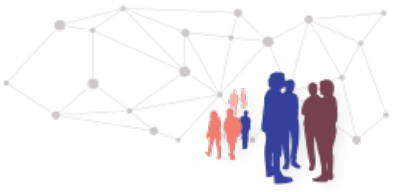

a coach who sometimes knows the patients' preferences better than they do themselves, she smiles to herself. «Florian is simply a natural at capturing a person's personality», she thinks.

In moments like these, she recalls her then groundbreaking study. Inspired by her experience with one of her first patients in the clinic – Sebastian was his name, she recalled – she had been working on how to put people's preferences into a form that could be understood by AI, so that it could make therapy recommendations not only on the basis of (individual) clinical data, but also taking these preferences into account. At that time, she had recorded the clinical dataset of Sebastian and all other patients herself and integrated it into the CanCon dataset for melanoma in many night shifts. The AI she trained with it was then able to shape the patient's therapy process, along with his changing preferences, in an optimal way. Her study revolutionized the use of AI in oncology, as it was the first to show how to incorporate such personality aspects and the prediction of their interaction with the therapeutic process into the individualization of a cancer therapy.

A few weeks later, Pia checks in with Carmen again. Pia is currently following the AI therapy plan confirmed at «Kant» a few weeks ago, but she wants to discuss something with Carmen: A few days ago, she received a request from Pineapple Health, one of the world's largest health technology companies: she was invited to participate in a study for a new technology to dynamically capture personality traits, which would also allow her personality to interact with therapeutic progress. Pia is still uncomfortable with the idea that one component of the therapeutic approach involves «egg freezing» combined with prenatal diagnostics to ensure that only eggs not containing the breast cancer gene – so her therapy will not only affect her, but also her offspring. This constellation had unintended consequences for her marriage, as she is discussing this issue also with her H coach – but her husband did not want to be present at the discussions with her H Coach; she suspects a latent jealousy here. However, using the new technology, she could not only better understand her preferences, but also discuss them directly with her husband without involving the H coach, which might solve this unintended consequence. Nevertheless, she is not sure whether such an adaptation of the therapeutic procedure – especially a longer-term replacement of the conversations with her H Coach by the new Pineapple technology – really suits her. Also Carmen is unsure how to advise Pia with respect to this constellation.

3.3 Possible questions for discussions

What do you like about the combination between AI and H-Coaches, what do you dislike? Would you see an issue if the current gatekeeping function of the doctor is reduced in order to strengthen the coaching function? Would you feel comfortable leading in such a therapeutic setting? Would you prefer to move to Pineapple Health's new technology for direct insight into what is happening therapeutically without involving the H coach?

4 Case Study Level «Far Future»

4.1 Situating the scenario

Once again, many years have passed. Medicine has become an information science, and large technology companies offer infrastructure, knowledge (data), and processing power that no single hospital or physician can provide alone. New and powerful tools have added to the arsenal of medicine. One such tool are digital twins: comprehensive simulations of a person's body from the cellular to the organ level. The digital twin is continuously fed data from the «real» twin via sensors, some of which are implanted when a person is born. People have direct access to their digital twin as well as to a

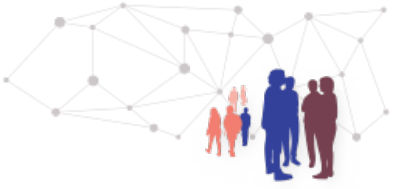

personalized AI coach that has replaced the previous H Coaches. The AI coach is in constant interaction with the sensor output and the patient's digital twin, allowing it to recognize when a medical problem is occurring. There is even a biological extension of the digital twin: cloned tissue of the real twin grows in nutrient fluid, in the form of organoids (simplified organ systems), for example. The AI coach can thus carry out preclinical drug tests in the tissue of the person concerned without endangering him or her – which, in the case of cancer, for example, provides information about the optimal cancer therapy.

This technological leap has fundamentally changed the relationship between patient and healthcare professional within a few decades. Gradually, AI systems have gained autonomy in making treatment decisions. Initially, this was only done in simple cases – for example, when it was a question of whether a patient should take an anti-nausea medication if a cancer therapy had side effects. This was followed by a period in which independently trained AI systems made a decision on more complex medical problems and autonomously reached a consensus, which human experts only had to review. It gradually became normal for patients to seek AI second opinions even when interacting with a healthcare professional. This also transformed the financing of the healthcare system: fewer and fewer «traditional» physicians were trained and corresponding personnel costs dropped significantly. Money increasingly flowed to the technology companies that provided the AI coaches. In the process, legislation ensured that no monopoly could occur – for every medical problem group, there had to be systems from at least three different providers. Excessive corporate profits were consistently skimmed off to control healthcare costs.

In the fully digitized healthcare system, the digital twin has become a basic right: every person is entitled to such a twin with certain basic functions – analogous, for example, to the earlier idea of basic insurance. The digital twin is part of the free «basic package» to which every citizen is entitled and which is financed by the state. In this setting, usually the patients themselves make therapeutic decisions in conjunction with their AI coach. Care has increasingly shifted to the home; patients only come to a health center for invasive procedures.

The role of physicians has also changed largely. There are only private training centers that continue to offer classical training for physicians who work primarily for the few remaining digital skeptics. Physicians are still involved in a narrow but diminishing advanced care setting – and are increasingly taking on the role of the former H Coaches as translators between digital twin predictions and the patients. Accordingly, doctors exchange information primarily with digital twins and contact real people only if there are clear indications that a person cannot cope with his or her digital twin. The state universities train specialist physicians as a form of information scientist whose task is to make sure that digital twins represent real people as good as possible. Particularly, some digital twin prediction may not match up with a patient's true health trajectories. In such cases, specialist try to train a twin's model with more adequate data or attempt to add new biomedical sensors.

4.2 How it could work

Many years have passed. Technology has advanced considerably, with new sensor technologies that can continuously monitor biological processes on the cellular, organ and system or body level, way beyond the single-sensor technology (such as insulin sensors) that have been common-place few decades ago. These sensors are mostly implanted but may also take the form of genome-edited cells. Likewise, compute resources have steadily increased and allow the modeling of complex biomedical processes. The challenge is no longer to gather data across as many patients as possible to train robust AI algorithms, as was the case few decades ago, but the real-time integration of personal data streams within a so-called digital twin system. Digital twin systems consist of two main components: the first is

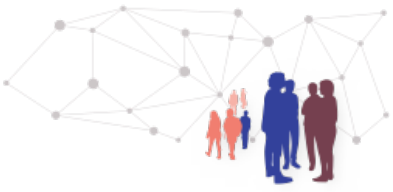

a full-scale and digital replica or model of an individual patient that is able to predict future health states and behaviors. The model is initialized using a pre-trained base model (an average model from existing digital twins) which is subsequently fine-tuned using continuous collected individual data sets. The digital twin is constantly optimizing its own prediction ability by comparing its predictions with actual (future) biological states or medical events.

Predictions are also refined by aligning results of organoid drug experiments with results from a patient's drug treatments. To that end, the digital twin has access and request organoid drug screenings when necessary. The digital twin's second component are interfaces for answering health questions and communicating health predictions, tailored for different stakeholders (patients, physicians, researchers). Physicians may interact with twin by studying diverging future twin versions based on different treatment choices. Patients interact via the AI coaches, virtual machines themselves that have access to the digital twin's forecasting capabilities and that translate health questions into meaningful treatment advice for the patient. Finally, researchers may interact with a twin, or set of twins, by testing new drugs in virtual clinical trials.

Carmen has retired at 75, but she continues to follow with interest the transformation of medicine into an information science. Statistically, she has another 25 years of life ahead of her and enough joy and energy to attend her son Kilian's virtual continuing education events as a guest. Kilian has become a respected member of DigMed, the worldwide association of the best digital physicians. His work on the differential diagnosis of digital twins is groundbreaking and it has significantly advanced fast-forward simulation methods. It has become possible to rapidly model subtly different variants of digital twins and offer predictive therapy based on this: the biological twins receive behavioral suggestions that extend several months into the future and can thus prevent the emergence of diseases. But it doesn't quite work out that way in practice, because people's stubbornness sometimes counteracts the well-intentioned advice. «It's easier to keep the digital twin healthy than the biological twin,» says Kilian.

Sometimes Kilian talks to his (real) mother about the medicine of the past, which is still kept alive in parallel societies. He is proud that Carmen was a pioneer, but sometimes also smiles at her nostalgic memories of the H Coaches, this «human translator» that is no longer needed in this form today. People have simply evolved, he says: just as people used to regard clothes as part of themselves, it is now the composite between biological original, digital twin and AI coach that form a human being, so to speak. As physician, you need to navigate in those different spheres, real and virtual, to be an effective health care provider.

In these discussions, however, Carmen then always counters what should be done if the biological and the digital twin develop in different directions. After all, everyone goes through autonomy training in their schooling (see the «Allocation» case) in order to be able to emancipate themselves from their digital twin if necessary. However, Kilian then says that such a contradiction cannot actually be possible, after all, the digital twin lives from the data input of its biological counterpart. Of course, he recognizes the problem that his research on preventive therapy sometimes collides with people's stubbornness. «Probably,» he thinks, «you have to adapt the curriculum for autonomy training...»

4.3 Possible questions for discussions

What do you like about such a fully digitalized healthcare system, what do you dislike? What role do you see for medical professionals in such a system? Would you want to live in such a world? If no, why not?

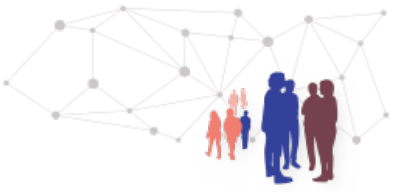

5 References / Links

[1] https://www.bfs.admin.ch/bfs/de/home/aktuell/neue-veroeffentlichungen.assetdetail.19305696.html

[2] https://www.usz.ch/individualisierte-krebstherapien-nach-langem-kampf-vom-krebs-befreit/

[3] https://www.krebsliga.ch/ueber-krebs/therapien/immuntherapien-und-zielgerichtete-therapien