

Getting a Diagnosis from Pixels

Computer scientist Bjoern Menze is training computer programs to “learn” how to detect brain tumors and propose individualized treatments. His work shows how medical image analysis is opening up completely new perspectives.

Klaus Manhart


Wenn aus Pixeln Diagnosen werden

Medizinische Bilder von Krebstumoren liefern Ärztinnen und Ärzten wichtige Informationen. Allerdings liegen auf den Bildern die Tumormerkmale nur als unstrukturierte Pixelansammlungen vor. Die qualitativen Eigenschaften von Tumoren – ihre Größe, Lage, Form, Oberflächenbeschaffenheit und Substrukturen – sind schwer in objektive und quantifizierbare Daten zu überführen. Bjoern Menze, Professor für Bild-basierte biomedizinische Modellierung an der TUM, versucht diese Lücke zu schließen.


Der Informatiker entwickelt Verfahren des maschinellen Lernens, mit deren Hilfe in den biomedizinischen Bildern automatisch anatomische Strukturen, Organe und Tumoren erkannt werden. Diese werden dabei nicht nur identifiziert, sondern auch so bewertet und analysiert, dass die extrahierten Tumormerkmale die Diagnostik und Therapieentscheidungen verbessern können. Aktuell konzentriert sich die Forschung des 40-jährigen Wissenschaftlers auf die Erkennung und Analyse von Gehirntumoren.

Mit seiner Forschungsgruppe konnte er zeigen, dass maschinelles Lernen viel mehr Informationen zu Tage fördert, als das medizinisch geschulte Auge erfassen kann. Der Computer wertet Bilder dreidimensional aus. Auf diese Weise werden auch komplexe 3D-Strukturen und Texturen erfasst, die für den Arzt nicht unmittelbar einsehbar sind. Gleichzeitig verbessert die Computeranalyse die Qualität der Information: Die Algorithmen können subtile Unterschiede, wie verschiedene Grauwerte, berücksichtigen und sehr genau bestimmte Muster erkennen, die Menschen entgehen.

Damit verbessern die Lernalgorithmen mit der quantifizierten Auswertung der Bilder die ärztliche Diagnose. Sie geben Medizinern objektive und fundierte Informationen an die Hand und sorgen für eine Qualitätssicherung. Letztendlich wird damit auch die personalisierte Medizin gestärkt und es werden Therapieempfehlungen und automatisierte Entscheidungen möglich. Die letzte Entscheidung wird jedoch immer beim Arzt oder der Ärztin bleiben. □



Tumor areas in magnetic resonance images of the brain are automatically measured using machine learning techniques. The picture shows two different types of tumor segmentations. Top images: Individual algorithms (blue) are consolidated to produce an optimized algorithm (magenta), which matches or exceeds the measuring accuracy of an individual medical expert (yellow). Bottom images show findings established by individual medical professionals (blue) and findings through group consensus (magenta). Computer algorithms are trained to reproduce the experts' consensus segmentations, which are taken as a benchmark.



Picture credit: Björn Menze (source: TUM)

“The algorithm can detect specific patterns with extremely high precision – more accurately than this would be possible for a human.”

Bjoern Menze



Medical images are a mine of information. Radiologists and medical specialists scrutinize computer tomography and magnetic resonance images, looking for anomalies and signs of disease. However, hidden in this vast sea of pixels is another layer of information that is not visible to the naked eye. And transforming this unstructured image information into objective, quantifiable data is no easy task.

Bjoern Menze, Professor of Image-Based Biomedical Modeling at TUM, wants to turn the information hidden in medical images into useful insights by converting the pixels into meaningful metrics. His aim is to support and validate today's qualitative, visual diagnostics with quantifiable values extracted from the imaging data.

This entails using computer algorithms from the fields of machine learning and computer vision to detect, analyze and evaluate anatomical structures in images from radiology, neuroradiology and nuclear medicine. And Prof. Menze's research efforts are already beginning to pay off, yielding better markers for tumor assessment and improvements to radiotherapy for brain tumor patients.

Computer-based tumor analysis

Menze's research team started out by programming a computer to recognize specific organs of the body. Here, the computer examines the pixels in medical images and classifies particular pixel structures as heart, liver or kidney. To train the computer, researchers first show it relevant images with the anatomical structures identified by an expert. Using artificial neural networks, the algorithm then learns how to categorize these organs and is ultimately able to interpret new images by itself (see FAQ on Machine Learning p. 27).

The team then took these algorithms developed to recognize organs and applied them to cancerous tumors in the human brain, so-called gliomas. "In terms of methodology, the task of locating an organ is very similar to that of identifying a tumor or tumor substructure," explains the 40-year-old researcher. "Rather than taking an array of pixels and asking whether they represent the left kidney or the liver, we now ask: Is this a tumor? And perhaps a particular tumor substructure?" >

Prof. Bjoern Menze

Medical computer scientist with an interest in archeology

Bjoern Menze is Professor of Computer Science and heads the Image-Based Biomedical Modeling at the Munich School of BioEngineering. The 40-year-old computer scientist is a fellow of the TUM Institute for Advanced Study (TUM-IAS) and conducts his research at the Center for Translational Cancer Research (TranslaTUM).

He develops algorithms that analyze biomedical images using models from computational physiology and biophysics. The emphasis of this work is on applications in clinical neuroimaging and the personalized modeling of tumor growth.

Bjoern Menze originally studied physics in Uppsala (Sweden) and Heidelberg (Germany), going on to complete his doctorate in computer science. Following postdoc periods at Harvard University in Cambridge and Harvard Medical School in Boston (USA), he continued his research as a member of the CSAIL Medical Computer Vision group at the Massachusetts Institute of Technology in Cambridge (USA), the Asclepius team at the Inria research center in Sophia-Antipolis (France), and the Computer Vision Lab at ETH Zurich (Switzerland). His work to transfer medical image analytics to Near Eastern archeology applications has been featured in Nature, Geo Magazin and Spiegel.

Image-Based Biomedical Modeling Group

www.translatum.tum.de/research-groups/image-computing/



Imaging techniques such as magnetic resonance imaging are used to generate numerous images of a brain tumor before treatment decisions are made. Bjoern Menze (left) and his colleague Dr. Jan Kirschke, senior physician at the Department of Diagnostic and Interventional Neuroradiology at TUM's university hospital Klinikum rechts der Isar, discuss the assessment of magnetic resonance images of a glioma patient.

The really crucial step, however, involves measuring the tumor once it has been identified. Here, the algorithm records its size, shape and relative location in the brain. Key analytical indicators include the image intensity and surface properties of the tumor in question – that is to say, its texture. The computer can differentiate between light and dark areas, detect whether the tumor is homogeneous or inhomogeneous and identify subareas worthy of attention. Training examples are initially used to teach the computer these core indicators, which it is then able to apply to new data sets.

Improving diagnostic accuracy

This type of automation benefits medicine in a number of ways. Machine analysis brings significantly more information to light than doctors are able to extract by themselves. While medical professionals generally focus on particular image sections and surface characteristics, the computer conducts a three-dimensional image analysis. This also identifies complex 3D structures and textures not immediately visible to the

physician. At the same time, computer analysis improves the quality of the information: “The algorithm can factor in subtle variations, such as different gray tones, and detect specific patterns with extremely high precision – more accurately than would be possible for a human,” underscores Menze.

By delivering detailed and quantified image analysis, the learning algorithms can enhance medical diagnostics. They provide doctors with additional information and provide a quality assurance net. They also facilitate well-founded treatment recommendations and can automate decision-making. Based on the pooled analyses, for instance, the computer can calculate the probability that a tumor belongs to a particular subgroup – perhaps determining that subgroup 6 is more likely than subgroup 4. The learning algorithm can also integrate additional information, such as genetic or histological data. Ultimately, however, the doctor will continue to reach a diagnosis and make clinical decisions. “We are supporting professionals in their assessment, not replacing them,” emphasizes Menze. ▶

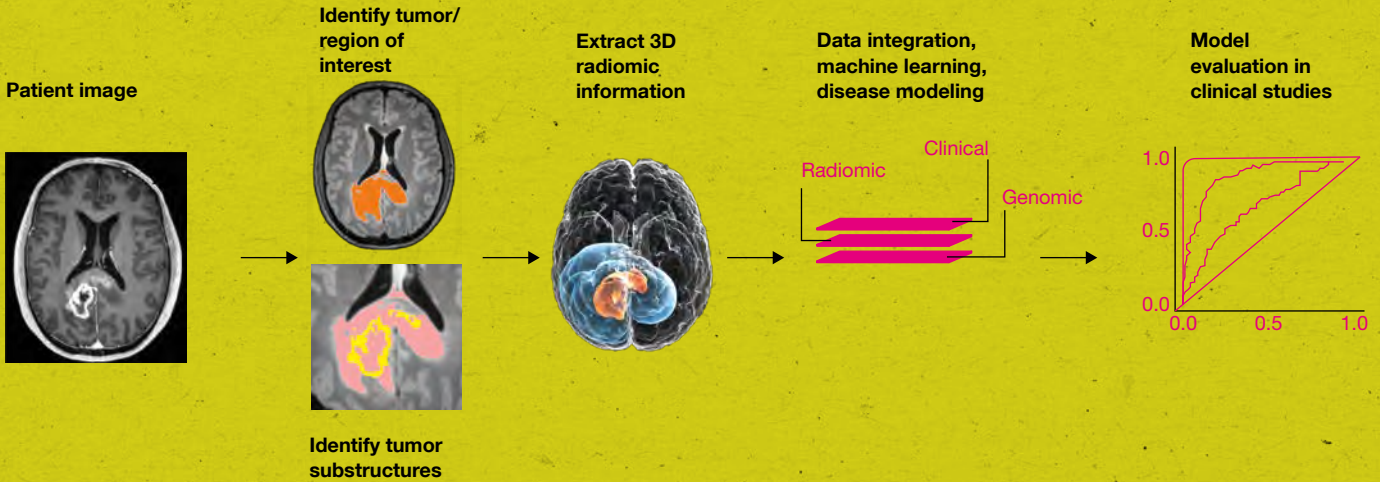


Computer vision and machine learning methods enable automated recognition and measurement of anatomical structures in computer tomography images. Data such as lung size, shape and texture are converted into a series of measurements, known as vector, that can subsequently be analyzed with statistical methods.


Computer Vision

Computer vision is a discipline that focuses on teaching computers specific capabilities that approximate the human visual system. At present, the primary focus of this field is the ability to identify specific objects within images, describe them, measure their properties and classify them. The outcomes can then serve as a basis for decision-making or determining subsequent processes. Many of the methods used by Bjoern Menze for medical image recognition and analysis stem from this area. Computer vision now also plays an important role in the development of self-driving cars. Here, the computer not only needs to “see” the street and everything happening on it, but must also be able to draw conclusions and learn from what it sees.

Extracting diagnostic image information and predicting disease progression



The flowchart shows the different steps in the process of image-based biomedical modeling that lead to the automated analysis of medical images. The information this yields could improve radiotherapy plans and reduce the side effects of radiation (lower image).



Improving radiotherapy plans for glioma patients – the target area and local radiation dose is adjusted to the tumor site and individual anatomy. In standard cases, the tumor is uniformly irradiated with a surrounding safety margin that is the same for all patients. Left: Conventional radiation chart. Right: Refined radiation chart derived from patient-specific tumor growth model.

Following through on major advances in terms of technology and methods, Menze's next big challenge lies in clinical applications. The computer scientist also views his work as contributing to progress in personalized medicine. Here, patient populations are divided into small subgroups to optimize treatment delivery. With the availability of imaging data, patients can be assigned with greater precision to subgroups that respond well to a specific type of treatment.

Adapting learning algorithms

Menze's research group often adopts mature algorithms – including a deep learning algorithm from Google – as well as ones devised by other researchers, such as that developed by computer vision expert Prof. Daniel Cremers, also based at TUM. “The fact that big data and techniques like deep learning are also hot topics in business right now is certainly to our advantage,” says Menze.

These algorithms cannot simply be adopted as is, however. They have to be adapted and further developed for biomedical applications. Menze and his team thus work in close collaboration with neuroradiologists, radiation therapists, neuropathologists and neurosurgeons to incorporate clinically relevant factors such as findings from tissue analysis or genetic markers into the learning algorithms. As well as healthcare professionals, the research team also liaises with physicists, biologists and mathematicians – who are working on tumor growth models, for instance.

Menze sets great store by this strong culture of interdisciplinary collaboration: “It's tremendously exciting to interact with so many different experts,” he says. “All of these stakeholders are working with these images because they hold the information they need. But we are the ones capable of extracting quantitative data from them.”

Even seemingly unrelated disciplines are benefiting from Menze's research efforts. As a PhD student at an archeology lecture in Boston (USA), Menze realized that the methods he was devising for medicine could also be applied to archeology. Together with archeologists from Harvard University, he used satellite images and his image recognition algorithms to locate ancient settlements – just like tumors in the human body. The computer scientist continues to pursue this “hobby” from his student days – whenever he can find the time.

Klaus Manhart

“The fact that big data and techniques like deep learning are also hot topics in business right now is certainly to our advantage.”

Bjoern Menze

Brain tumors

Brain tumors occur relatively rarely compared to other types of cancer such as lung, breast and colorectal cancer. Worldwide, more than 256,000 brain and other Central Nervous System tumors were estimated to have been diagnosed in 2012, with incidence rates varying across the world (Cancer Research UK).

The groups most frequently affected are people between the ages of fifty and seventy and children (to a lesser extent). Benign gliomas and malignant medulloblastomas are the most prevalent forms in children. Conversely, malignant gliomas and benign meningiomas predominate in older people.

The outcome of a brain tumor cannot generally be predicted. Some forms of tumor have favorable treatment prospects, whereas others do not. The chances of successful treatment depend on the location of the tumor, the type of tumor cells, and sensitivity to radiation and chemotherapy.