Professor Jürgen Hennig is Co-Chairman and Scientific Director of the Department of Diagnostic and Therapeutic Radiology and Medical Physics at the University Hospital Freiburg, Germany. He is known as the inventor of the RARE sequence, aka Turbo- or Fast Spin-Echo. His research activities are, however, wide-ranging and include methodological developments in MRI and spectroscopy, functional neuroimaging, molecular imaging, MR in oncology, cardiac and cardiovascular MR, and MR in metabolic disease. Professor Hennig's work has been recognized with numerous awards: Among others, the Max Planck Research Award, the Einstein Professorship of the Chinese Academy of Science, the Gold Medal of the Society of Magnetic Resonance, and the Hounsfield Medal of the Imperial College London. Over 480 publications are evidence of his focus on education.

Dear readers and colleagues,

It's that time of the year again and ISMRM is drawing near. A brief look at the program shows, unsurprisingly, that artificial intelligence and deep learning will again dominate the meeting. There is no doubt that AI will, and has already started to, permeate and transform all aspects of our field, from image reconstruction to diagnostic reading and beyond. It is less clear which of the multiple approaches and applications will ultimately take MR to the next level. The development of MR over the roughly 40 years since the first human scanner has never been a straight line. Whenever people thought, now we know everything and the rest will just be improved engineering, a new idea came along and transformed the field. This was the case with the advent of parallel imaging in the late 1990s and compressed sensing about 10 years later, to name just two examples.

Early days

I entered the field in 1984. I was already 33 years old, an age when others were already fairly far along their career path. I was still searching and looking for opportunities. I hadn't exactly been an end-focused student and took my time to complete my studies in Stuttgart (Germany), Imperial College London (UK), Munich, and finally Freiburg (both in Germany). To me, being a student first meant finding out who I am and only second acquiring knowledge and expertise about my chosen subject. At suitable junctures before and after my PhD project as well as before and after my postdoc in Zurich, Switzerland, I took off for typically three months at a time to go backpacking in Nepal, Peru, Indonesia, and Central Africa, the latter with my then fiancée and now wife Annemarie, the first three on my own. Getting into MR was, like many things in my life, rather serendipitous. I really wanted to go into industry and applied for a job with Bruker NMR. They didn't have an opening but Bernhard Knüttel, CEO at that time, told me about a project to install what was to be the 3rd MR scanner in Germany, and Professor Werner Wenz, head of radiology in Freiburg, was looking for a specialist with a PhD to help set up the system. When, after returning from Africa, I started work in Freiburg, I found that preparation of the scanner room had been seriously delayed (some things never change) and that the scanner was still sitting in the factory. So, for the next six months I commuted to Karlsruhe, Germany, to familiarize myself with the equipment and sequence programming. My first project under head of application development Bernhard Ströbel and chief engineer Dieter Ratzel was to implement a radial multiecho sequence for relaxation spectrometry.

Increasing speed

The overriding topic of conversation during coffee breaks and lunches were the long acquisition times of 10 to 15 minutes per scan and possible ways of overcoming them. Sir Peter Mansfield had already shown that this can be done with echo-planar imaging (EPI) and in one of his early publications even mentioned that multi-echo imaging should also be possible with multiple spin echoes. Larry Hall had just published a paper on using a long echo train to acquire all projections for radial reconstruction in a single echo train, but it was clear that this would not work on a clinical scanner since it required exact 180° pulses

We should by all means take a fresh look at 'old' topics in view of everything that we have learned in the meantime, but of course it would be prudent to consider insights already gained rather than wanting to start from scratch.

and only works for very long T2s. At that time, Bruker had decided to switch from radial encoding to Fourier encoding but still use CPMG multi-echo readout to perform relaxometry. Bernhard Stroebel was in charge of developing the 'bread-and-butter' sequence with identical phase encoding in all echoes of the echo train and I thought I'd try using this for echo imaging. We already had experience in adapting the CPMG conditions to the use of switched gradients through previous work on radial imaging, so rewinding the phase-encoding gradient before each refocusing pulse seemed like the natural thing to do and didn't feel very 'inventive'. The first human image we took was of the eye of a volunteer who happened to be scheduled on the day that we completed the first successful phantom experiment. Pulse programming was pretty straightforward in these days, the infamous Z17c process controller allowed only 16 lines of code, so everything had to be compressed into 16 instructions. The array processor used for reconstruction had to be programmed in assembly code, so without the enthusiastic support of chief programmer Arno Nauerth I would not have got very far.

I can't say that RARE (aka TSE or FSE) was an immediate success. On the contrary, for many years I presented the latest developments at the obscure 'Rapid Imaging' sessions at SMRM meetings without much interest shown, alongside Sir Peter Mansfield's equally obscure EPI images. The first attempt at publication was rejected with the remark 'This has been tried before, it doesn't work'. And indeed, there had been previous attempts at spin-echo-based echo imaging, which used different types of phase-cycling schemes to deal with the multiple spurious signals that suddenly appear after multiple refocusing pulses. Luckily, I was unaware of these at the time, otherwise I may well have tried to work along those lines. So sometimes it does pay to not let too much information stand in the way of freethinking ...

Hot topics

My first Society of Magnetic Resonance in Medicine (SMRM) meeting was 1985 in London. It was my first major scientific conference, so I had no idea what to expect. One of my lasting impressions was attending the notorious session when Raymond Damadian used his role as moderator for a lengthy tirade against Paul Lauterbur, which in the end got him expelled from the society.

1 Extract from the table of contents for the 1985 SMRM.

Scientifically, THE hot topic at the time was quantitative MRI (sounds familiar?). A look at the table of contents for one session (Fig. 1) reveals that all the hot topics discussed today were already there, albeit under somewhat different names.

In spite of these similarities, it is quite wrong to say that MR has gone around in a circle. I like to think the path is more like a helix: We have returned to the same point with respect to some coordinates, but have considerably advanced in another dimension. So we should by all means take a fresh look at these 'old' topics in view of everything that we have learned in the meantime, but of course it would be prudent to consider insights already gained rather than wanting to start from scratch.

Relaxometry

Relaxation in tissue is somewhat more complex than the values we get from fitting our data to a basic version of the Bloch equations, irrespective of whether we do that based on 'traditional' exponential fit or more fancy techniques like MR Fingerprinting[1]. Due to tissue heterogeneity, relaxation times are inherently non-monoexponential. As an illustration, Figure 2 shows the T2-spectrum of the human brain acquired from multiexponential analysis of a long echo train with 6 ms echo spacing.

A practical consequence of this is the fact that measured T2s will depend on the actual measurement parameters: Measuring at different TEs will influence the result of a monoexponential fit, a change in TR may lead to different degrees of saturation of different compartments, which will then also show up as a change in the measured T2.

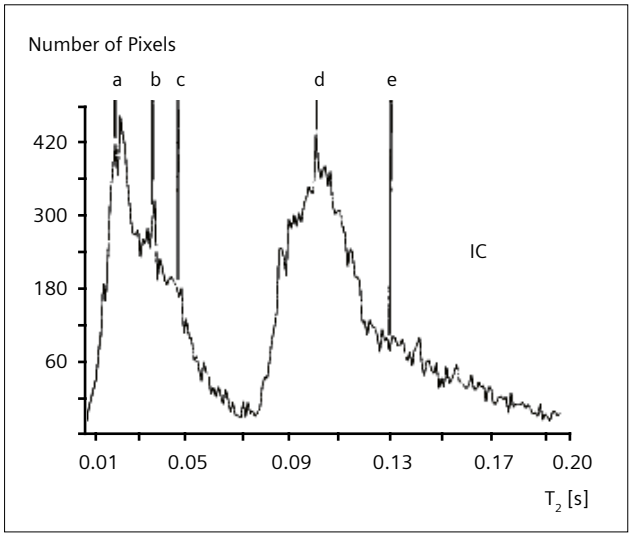

2 T2-spectrum of the brain. *(From K. Gersonde, F Elsberg, Tolxdorff, D. Ratzel, DB. Stroebel. Analysis of Multiple T2 Proton Relaxation Processes in Human Head and Imaging on the Basis of Selective and Assigned T2 Values. Magn Reson Med 1,463–477 (1984)).*

T1 is more innocuous for the simple reason that T1s of most tissues are rather long, which means that spins have had a chance to visit multiple compartments during the T1 decay. Given a typical diffusion constant, the 'averaging radius' over 1 s is about 50 µm, enough to average out most tissue heterogeneity.

On the one hand, the richness of information contained in a full analysis of multiexponential decay offers the possibility of more specific tissue characterization, on the other hand, such relaxometry measurements take time and therefore are not (yet) suitable for practical clinical application. It is, of course, perfectly legitimate to explore whether simplified monoexponential T1s and T2s could be useful as 'biomarkers' for disease. However we should be aware that, especially for T2, the measured values are only approximations, which also depend on the actual measurement sequence used, so the universality implied by the term 'quantitative MRI' is just a fiction. Calibration measurements typically presented to verify the validity of some new method are typically performed on homogeneous solutions and are not really pertinent since they do not reflect tissue heterogeneity.

FLASH

Relaxometry quickly lost general appeal in the late 1980s. This was not so much due to its inherent value but more to the advance in fast imaging, most notably the arrival of fast gradient echo imaging (FLASH). Jens Frahm and Axel Haase had already presented a poster at ISMRM 1985 (p. 980) on fast imaging techniques showing FLASH together with single-shot STEAM imaging. This really took off like a storm the following year. The possibility of acquiring rather interesting looking images in seconds rather than minutes caught the imagination of the community, so already at SMRM 1986 there were multiple abstracts inventing and reinventing FLASH under more or less fanciful new names: I particularly liked GREED for gradient reversal echo, equilibrium driven. Then and thereafter, contrast behavior was investigated and understood, leading to spoiled FLASH, refocused FLASH (FISP), PSIF and ultimately TrueFISP for the fully balanced version – of course, with a full set of unique acronyms for each vendor.

The mixed contrast of gradient echo techniques with variable contributions of T2*, T2, T1, depending on type and parameter selection, limited the use of FLASH to standard imaging, but it has spawned a tremendous growth in new techniques where FLASH is used as a vehicle for new areas of application – MR angiography, Cine Cardiac MR, flow imaging, susceptibility-weighted imaging, MR fluoroscopy – to name but a few. In the following years,

[1] *Work in progress: the application is currently under development and is not for sale in the U.S. and in other countries. Its future availability cannot be ensured.*

SMRM programs were dominated by the development of new sequences and their clinical application. Indeed, nearly all sequences in MR scanners today originate from researchers at academic institutions. There are a few notable exceptions – Gerhard Laub and Chuck Dumoulin come to mind who shared the Gold Medal in 1993 for their work on MR angiography.

Hardware

Hardware development was different in the early years. There always was some RF coil development in academia, but the rest (magnets, gradients) was driven by industry. The advent of shielded gradients in the early 1990s signaled the next transformative step in MR. Liberating MR from the obnoxious eddy currents inherent to the use of unshielded gradients in a superconductive magnet with its many layers of conductive copper and steel led to new degrees of freedom for MR sequences. I was one of the first beneficiaries of that. RARE went for many years largely unnoticed – the challenges for implementation on supercons were the main but not the only reason for this.

Bob Mulkern at Brigham and Women's Hospital, Harvard Medical School, was the first to implement RARE on a 1.5T GE scanner. The first paper on what he originally called FAISE led to an explosive growth, similar to FLASH five years before. Bob could have easily taken the praise and reduced my original work to some insignificant prior art (this has happened in other cases), but he did not. For this, I am forever indebted to him and, in particular, to Ference Jolesz.

Shielded gradients finally brought EPI into clinical application. As early as 1987, Richard Rzedzian and Ian Pykett, both from Peter Mansfield's lab in Nottingham, had introduced the Advanced NMR system, an EPI-only MR system with a resonant gradient system, to the market. Alerted by this unwelcome new competitor, all major vendors intensified their effort to make EPI happen. GE introduced Advanced MR, Siemens followed with the MAGNETOM Vision (regarded by many users as the best MR system ever), which was EPI-ready in 1994.

EPI not only finally facilitated the practical use of diffusion-weighted imaging, its advent also coincided with the introduction of BOLD-based fMRI. Ken Kwong, who was the lucky user of one of the Advanced MR EPI gradients, was the first to perform BOLD-based EPI for fMRI studies.

The introduction of fMRI was not so much a transformational step, it was more like opening the door to a vast new field of application for MR. It also opened up a new and vast area that was not primarily driven by clinical application but by neuroscientific research.

Now, finally, MR seemed to have reached its destination. All the main acquisition modes, FLASH, EPI, and RARE, were working and gave great results, scanners had reached a state of stability and maturity, applications were flourishing. What else was there to do?

The next wave in speed

Then came ISMRM 1997 when Dan Sodickson presented SMASH, a new technique for parallel imaging, which was the start of Parallel MRI, a whole new dimension in accelerated MRI. A few months later, I organized a fast imaging workshop in Asilomar together with David Feinberg. One of the posters was submitted by a young PhD student from Zurich named Klaas Pruessmann on a technique he called 'SENSE'; and although I admit that I didn't get through all the math, I am happy to say that our two-man poster award team recognized this achievement with the poster award (an Asilomar T-shirt, if I remember correctly).

What has happened since then? Not much or a lot depending on how you look at it. It is a sobering thought, but the extremely dynamic and rapid development of all basic technologies with respect to sequences, RF coils, and gradients, which took place between the early 1980s and the mid 1990s, was more or less concluded about 20 years ago. I still operate a 1.5T MAGNETOM Symphony upgraded from a MAGNETOM Sonata bought nearly 20 years ago, and it basically knows most of the tricks of our latest scanners. It is a bit slower, the gradients are not quite as fast and strong, it has a few receive channels less, but basically it is the same machine and perfectly adequate for clinical routine use. Since then, we have, of course, seen the introduction of 3T and more, culminating so far in the 7T MAGNETOM Terra. Translation and adaptation of methods and technologies to higher fields has been an astonishing effort requiring tremendous expertise and skills, but it has been a translation – nothing more. The one fundamental challenge is the fact that the RF wavelength at higher fields approaches the dimension of organs within the body, leading to inhomogeneities of the RF field. As a consequence, image intensity and image contrast may vary across the field. What is worse, the non-negligible E-field contribution in this not-quite-near-field-anymore regime has led to tremendous and still largely unsolved SAR problems. Much creativity and effort has been applied by many groups to these problems, but there is still much more to do.

Free thinking

For me personally, ISMRM 2005 was a real turning point in my scientific work. I had the honor and pleasure of giving the Mansfield lecture in the famous Jackie Gleason theater, home of the famous Jackie Gleason show in the 1950s and 60s. For whatever reason and the first time ever in my career, I went to the theater to test my presentation on stage

rather than just checking it in the preview room (where it worked perfectly). I had a few videos and animations and, what shall I say, not one of them worked. Together with the technician, I spent the whole afternoon reformatting my files and downloading new codecs until everything went smoothly. Other speakers of plenary sessions in the same location were not so lucky and had to spend a lot of their precious presentation time explaining to the audience what they should be seeing, but weren't.

The topic given to me for my lecture was 'Fast Imaging Horizons in Rapid MR Imaging'. In the preparation for the talk, I wanted to end on a cliffhanger to indicate that there is still much to come. So I thought, what's the limiting factor for imaging speed? Gradients, of course. So what about omitting gradients, how can we still get spatial information? By making use of the spatial sensitivity of multicoil arrays. Put like that, imaging speed is only determined by the speed at which one can sample data, which is extremely fast, at megahertz sampling rates, if necessary. In extreme cases, one can allocate each image pixel to a different coil. This led to the term OVOC for 'one voxel one coil imaging' (I later found out that this idea had been published before by Mike Hutchinson in 1988 as a pure thought experiment and I should also like to mention that Fa-Hsuan Lin independently developed a very similar concept with his inverse imaging technique).

Rather than just presenting signal traces corresponding to the time-dependent signal under each coil, one can first acquire a high-resolution image and then use the coil images as a basis for dynamic images generated by the sum-of-squares combination of the reference images with the instantaneous signal. This way one can, in principle, produce images of arbitrary resolution (which is defined by the reference data) at arbitrary speed – spatial resolution appears to be decoupled from temporal resolution. I didn't really expect much when I performed the test experiment, so I was quite pleased and surprised to see that even with the crude 8-channel coil setup that I had, I could clearly see spatially variable ECG- and breathing variations. So this was really the starting point of my still ongoing work on ultrafast fMRI.

Chuck Mistretta generously claims that my talk inspired him to develop what he called the HYPR-technique. He also coined the term 'Hennig limit' for making an image from a single measurement point (or, ultimately, from no data at all, which brings us dangerously close to deep learning and AI).

As another follow-up, I thought that it would actually be neat to sacrifice some of the unnecessary high speed in order to introduce some 'homeopathic' dose of gradient encoding (still not really thinking about producing an image but rather just putting some more spatial discrimination into the multi-sensor readout). For this purpose, it seemed like a smart idea to have each RF coil associated with an individual gradient coil oriented toward the coil coordinates. For multiple coils, these individual gradient fields combine to form a highly non-linear spatial encoding field, which is how PatLoc was born. This led to a further highly fruitful line of research on non-linear spatial encoding fields, which has meanwhile been taken over by Maxim Zaitsev in my group, and several others elsewhere, culminating in his 84-channel gradient matrix coil.

This free thinking to come up with a (however unrealistic) attention-grabber for a lecture has given me the basis for much of my scientific work for more than a decade. This is not the first time that thinking about the fundamentals has led to new ideas – that's why I really enjoy giving introductory lectures or presentations to a more general audience, both of which force one to really think and maybe rethink the basics. (I am still waiting to see what inspiration will come from writing this editorial!)

Looking ahead

The elephant in the room in this account of the technological and methodological development of MR is, of course, the magnet. Current magnets all look the same, big beasts, more or less short, 60–70 cm-wide supercons with 1.5 or 3T field strength. This is a result of the fact that building a superconducting magnet has always been and still is an extremely demanding task defined by very narrow constraints set by materials, currents, and forces. Except for making magnets somewhat shorter and bore sizes somewhat wider, not much seems to have changed since the first supercons were introduced in the mid 1980s. But this is about to change.

Currently, a first step is being taken by the introduction of low-helium magnets, which not only saves on helium costs but reduces the installation effort by avoiding the need for a quench pipe. This is a first step toward making MR more accessible to a wider community. The next step is already on the horizon. Stefan Röll, well-known to many Siemens users, has created his own company Neoscan with the aim of building an MR system based on high-temperature superconductors. The HTS wire is able to sustain unprecedented high current density, which in principle allows the construction of extremely compact magnets. Once this technology has matured, it will be possible to build very slim shielded magnets, where the magnet plus shielding will be only a few tens of centimeters thick (much like current gradient coils.) This will also allow the construction of tailor-made, extremely compact magnets, e.g., for use in areas without much supporting infrastructure. This will take a while to become routinely available. The current wire can be manufactured only with a limited length, too

short to wind a magnet from one strand of wire. Also, the mechanical stability of the wire in strong magnetic fields is still a challenge, to name but two of the difficulties. The 'ISMRM Workshop on Accessible MRI for the World', which is taking place in India as I write this article, shows that new concepts to bring MRI to a larger community is one of the future growth areas for MR.

The one big transformative step that I have not yet addressed is the introduction of compressed sensing presented by Miki Lustig at ISMRM 2006 and published in December 2007. In 2008, it already had 41 citations and 78 in the following year (and nearly 3000 as of today), which demonstrates the immediate huge interest it generated. The promise of acceleration factors way beyond what is achievable with conventional parallel imaging techniques immediately fascinated the community with the promise of finally breaking through the Nyquist limit of 'conventional' imaging (which we also took as the motto for a series of workshops that we held together with UW Madison – watch out for the next installment).

Artificial intelligence

And now AI has arrived to take MR (and not only MR) to the next level of development. Neural networks have been around since the 1980s, so the basic principles have been known for a long time. But new algorithmic approaches like supervised backpropagation as well as the rapid increase in computing power and data storage have now made deep learning feasible for many applications –self-driving cars, language translation, face and speech recognition, to name but a few. In medical imaging and MR, AI got off to a somewhat rocky start, triggered by an ill-informed statement by Geoffrey Hines, one of the pioneers of backpropagation, that AI will make radiologists obsolete within five years. Well, nearly three years have passed since then and radiologists are still around. AI will certainly transform radiology and it will also transform MRI, but in which way is still unclear. It will make some tasks previously performed by humans obsolete, but this will most likely affect the simpler and thus more boring tasks first. Rather than shying away from this new challenge, we should embrace and explore it, and it is good to see that the ISMRM community is doing exactly that. AI has been the hot topic of the last few annual meetings, on which there have already been three workshops, and I think that's exactly the spirit: There is a new opportunity, we should grasp it and explore it even though we still can't see where it will take us. The explosive growth of work on AI has already produced some do's and don'ts:

Using AI for further acceleration seems like one very promising application. Several groups have already shown that a CNN, even when trained on a modest dataset, can produce near-to-perfect images from highly undersampled data. Image segmentation also works extremely well. Even tough tasks like segmentation of hippocampal subregions, which have been tackled by generations of computer scientists with mixed success, seem to work extremely well. (If Geoff Hines had talked about computer scientists working on segmentation he would have been correct.) It has already been shown that AI can be used to generate realistic-looking CT images from MR data (which can be helpful in, e.g., radiation therapy or in MR-PET) or even nice-looking 7T images from 1.5T data (which seems to be less compelling). The possibilities seem unlimited.

I am not so convinced about the use of AI directly on raw data, effectively omitting Fourier transformation. It doesn't seem to make much sense to replace just this part of the reconstruction process by one that has to be trained first. After all, we know exactly how it should be done. I still think that, in problem solving, it is a good strategy to approach the solution in a way that is as close as possible to using everything we know, deploying AI to take the last missing step. This doesn't mean that exploring how FT can be replaced by AI is useless. As we have a 'basic truth' in this area to assess the result, this may actually help to develop some intuition about this mysterious 'black box'.

There is a fundamental difference between innovation brought about by AI and all previous innovations: Previous progress was made by new insights and better understanding. Scientific problems were solved by analyzing and understanding the subject matter, and this understanding could be used to solve new problems. The motor that has driven science over the last few centuries was the fact that solving some problem invariably not only led to an improvement in the issue addressed, but also to new and previously unexpected areas and directions. AI skips the 'understanding' part of the problem-solving process, it produces solutions (and often very good solutions) directly and without generating any understanding. So, without new insights, where will future innovation come from? This is not meant to be a rhetorical question, but only time will tell.

Back to MR. Given the similarity of MR scanners across all vendors, one may think that the technology is mature and, apart from incremental improvements, not much more can be expected in the future. I think that this is wrong and the current situation is more like the calm before the storm. For many years, development in MR followed a path set by GPS (GE – Philips – Siemens). Now companies like SSI, Tesla, ASG, Neoscan have started to develop new magnet concepts, electronics and devices to build MR, from power supplies to spectrometers to RF coils and gradients, become cheaper and better. Last but not least, new and serious competitors from China and

elsewhere are entering the arena. So I am pretty sure that there is a lot of change to come, which may even include a change to the whole business model of MR. I very much agree with the concept Toby Block promoted in his plenary talk last year: It is very appealing to consider a world in which sequence development takes place in a universal, open space, in which any sequence can be executed on any scanner by any vendor (within the hardware constraints of each system, of course). 'Sequence' means not only the pulse program but the whole measurement workflow, including real-time functionality and feedback loops. Having all the bright minds of the scientific community work toward improving MR will accelerate development considerably and will facilitate cross-platform comparisons and multicenter studies as a big step toward standardization as a basis for personalized medicine. Maxim Zaitsev has taken a first step toward such universal programming with his Pulseq platform. There are others, too, and I am very enthusiastic about it for very selfish reason. It would be futile for me to try and master in-depth pulse programming in IDEA, too little time, too long out of touch. But to write a Pulseq sequence in MatLab is doable, even for me. I recently implemented RARE using Pulseq, it took a few hours and 72 lines of code, more than the 16 I originally had, but much less than the implementation on any existing platform.

Is an open and universal sequence development platform a valid business model? Microsoft has shown that it is: It hasn't done badly by providing an open platform with Windows® with the opportunity for anybody to develop their own software while still making a success of its own, very profitable, software business with the Office® package and others. In medical imaging, this is even more straightforward since users can easily develop new sequences, but very few will have the means to go through FDA approval. In this way, the otherwise costly development part will be outsourced, but the business part will still stay with the vendors.

Will that be the future? I don't know. If I have learned one thing over the years, it is that development very rarely takes the most logical and scientifically rational path. The only thing I am pretty sure of is that we can expect many more (radical) developments in MR in years to come.

Jürgen Hennig

We appreciate your comments.
Please contact us at magnetomworld.team@siemens-healthineers.com