Industry Corner
Debdeep Paul

In this "Industry Corner" column, we interview Dr. Ivan Tashev, a Partner Software Architect at Microsoft Research (MSR), where he currently leads the Audio and Acoustics Group and coordinates their Brain-Computer Interfaces project.

Dr. Ivan Tashev received his Diploma Engineer degree (master's equivalent) in Electronic Engineering and his PhD in Computer Science from the Technical University of Sofia, Bulgaria, in 1984 and 1990 respectively. He was assistant professor at that university, teaching "Data and Signal Processing" and "Real-Time Systems Programming", when he joined Microsoft in 1998. Ivan Tashev is currently a Partner Software Architect and leads the Audio and Acoustics Group at Microsoft Research (MSR) in Redmond, Washington, USA. His research interests include audio signal processing, machine learning and artificial intelligence, multichannel transducers, and bio-signal processing. He also coordinates the Brain-Computer Interfaces project in MSR. Dr. Tashev has published two books, two book chapters, over 100 scientific papers, and is listed as inventor on 50 US patents. He is an Affiliated Professor in the Department of Electrical and Computer Engineering of University of Washington in Seattle, and Honorary Professor at Technical University of Sofia, Bulgaria. Dr. Tashev transferred algorithms to RoundTable device, Windows, Microsoft Auto platform, and served as the audio architect for Kinect for Xbox, and for HoloLens. He is an IEEE Fellow and a member of AES and ASA. More details about him can be found in his web page https://www.microsoft.com/en-us/research/people/ivantash/.

In this interview, Dr. Tashev will give us a view of the emerging trends in human-machine interfaces, provide some open research problems, and how academia-university partnerships can help tackle these challenges. He concludes with advice for young researchers wishing to enter the HMI field and what topics/skills are needed. We hope you enjoy the interview!

---

1. *Please give us a bit of background about yourself and how you ended up at Microsoft Research.*

In 1997, I was an Assistant Professor at the Department of Electronic Engineering at the Technical University of Sofia, Bulgaria, teaching programming and algorithms in this predominantly hardware department. East-European countries were shaken by the changes in both political and economic life that started in 1989. The transition from a command-driven to a market economy was taking very long, and the country was in a pretty much permanent economic crisis. While I had no concerns about my job at the largest technical university in the country, funding for science and research was scarce.

Before 1989, Bulgaria was specialized in computer hardware and software in the Soviet Bloc, which sparked the interest of the large and fast-growing IT companies abroad. I heard that Microsoft was coming to Sofia for the second time to hire IT specialists and sent my resume. Four months later, I had a work visa in my passport and airplane tickets to Redmond for me and my family. I started as a software engineer in a product team, working on a new feature in Windows 2000, practically a V1.0 product. After releasing it, I went to another V1.0 product, and at the end of the second shipping cycle, I found that Microsoft had a research organization. I immediately applied for an opening, was hired, and started in Microsoft Research two months later. These two and a half years in product teams and the two full shipping cycles there gave me valuable experience and knowledge of how Microsoft engineering teams work. This helped me later in transferring new technologies and algorithms from Microsoft Research to Microsoft products – I was speaking the language of engineers and managers in these teams, knew their concerns, and how to address them.

2. Please give us a bit of background on MSR and its role in developing next-generation human-machine interfaces (HMIs), including your latest work in brain-machine interfaces (BMI). How has it evolved over the last few years?

Microsoft Research (MSR) is a state-of-the-art industrial research organization, one of the largest in the world. Since its creation in 1991, it sits on three pillars: pushing forward the state of the art in corresponding research areas (publishing scientific papers); rapidly transferring created new technologies to Microsoft products (going the extra mile to prepare your technology for a product); and making sure that Microsoft products have a future (monitoring the trends in the corresponding research areas, detecting disruptive technologies). Currently, MSR has approximately 1,200 staff and publishes more than 3,500 scientific papers per year in more than 50 research areas. This is the most academic research organization of a commercial company I am aware of. Usually, in commercial companies, the research organization is a sacred place, and external people are not allowed. In MSR, we do have a summer internship program, considered one of the best in the industry. Every year, more than 1,000 PhD students from all over the world come to work with us on research projects that usually end with a talk and a conference paper.

Microsoft is a software company and, as such, pays a lot of attention to human-computer interfaces (HMI). We have several research groups working on new input and output modalities for interaction with computers and their seamless integration into the overall HMI design. The transition from desktop computers with their screen-keyboard-mouse HMI towards portable devices, handhelds, and wearables required new input and output modalities, such as pen, touch, voice, gesture, gaze, haptic, and spatial audio. First, these new modalities must be enabled technologically – all the under-the-hood signal processing that provides high-quality output. Examples here are accurate pen coordinates and good quality clean voice signals suitable for speech recognition. Second, given sufficient quality underlying technologies, the question is how and when to use one or another modality, based on the device formfactor and usage scenarios.

Brain-Computer Interfaces (BCI) are just the next step towards more direct interaction with computers. We started with punch cards, then went through the teletype, and stayed a long time with the electronic keyboard and display. Today, we have touch screens on our tablets, phones, and even wristwatches.

We use speech to ask computers and receive answers via voice. Gesture recognition entered the world through devices such as Kinect and HoloLens, the latter using gaze tracking as an integral part of the user interface. Regardless of how natural and intuitive HMIs become, they still have one thing in common: the chain is brain-muscle-sensor-computer. With BCI, we eliminate muscles from the chain and provide a direct interface with the computer without any moving parts.

In my Audio and Acoustics group, we are working on audio signal processing using machine learning and artificial intelligence approaches. Several people have a background in bio-signal processing, including myself. In 2018, we decided to host the BCI project in MSR because we thought that with our signal processing and ML/AI background, we could contribute. We were at the right time. Three factors enabled the process of moving BCI from research labs to the general population: developments in microelectronics (which allowed building miniature battery-operated blocks for collecting 32 EEG channels with 24-bit accuracy and wireless connection to the computer), new technologies in material science (electrodes with better contact with the skin, glues, coatings, 3D printing), and last but not least, the revolution in signal processing caused by new ML/AI approaches. Today, we have commercial products for the general population that help with meditation, reducing stress, and even active BCI interfaces for people with disabilities.

3. Within the field of HMI (or BMI), what are some open research problems that new students could focus on and try to help solve?

We are in the middle of rapid technological advancement in BCIs. We see on the horizon brain-controlled bi-directional motor prosthetics, using biomarkers for detection and recommendation for personalized treatment of neurological diseases.

On the consumer side, we expect to see BCI components in the HMI of everyday devices, such as phones, tablets, and laptops. These devices have relatively good computing power and memory, connection to the internet, and reasonable battery life. Connected to them are peripheral sensors that do not have enough computing power and battery life for direct connection to the internet. They provide additional information that allows building personalized AI agents to serve better and be more efficient in helping humans. A good example here is smartwatches – they provide information about the health and wellbeing of the wearer, this happens through a more powerful computational device, usually a smartphone. I see BCI sensors entering the same ecosystem. The form factor can be something we wear on our heads (headphones, earbuds, glasses), or something that is unobtrusive and easy to put on/take off – a headband, for example.

4. What role do you see academia playing in trying to help address these concerns? Is there room for academia-industry partnerships in this domain?

Academia is where most of the basic research has been done, and where future researchers for academia and industry are trained. From this standpoint, the role of academic research and PhD programs is critical to advancing further research in BCI. As long as we stay in blue-sky research, collaboration between research groups in academia and industry is possible and happens now. There are many papers in the field published by mixed groups of researchers from academia and industry.

Things change when the technology matures and is subject to integration into a product. Then, usually, the university spins off a startup, or the industrial research lab ceases the collaboration with academia. In both cases, the goal is to have clean IP ownership of the technology for integration into a commercial product.

5. What advice would you offer young researchers entering the field of HMI/BMI, from the latest technical expertise that is/will be required, to the non-technical skills sought today by hiring managers at MSR and elsewhere?

PhD studies are generally training on how to learn and present novel research results. BCI and computer science are fast-paced areas of science, and we are lifelong students until the end of our active careers. Constant learning should be accepted as part of the profession. We report our research results in scientific papers. The publication record and the number of citations is the best representation of researchers and their achievements; this is one of the first places research hiring managers look at.

I mentioned above the three areas which are the big enablers of BCI: microelectronics, material science, and ML/AI. It seems to me that the latter will play a major role in future BCI technologies. For a career in BCI and neuroscience, knowledge, experience, and expertise in machine learning, neural networks, and artificial intelligence are a must.

6. Any last words of advice?

For more than 26 years at Microsoft and over a 40-year-long career, I have learned that one of the most important things is to stay focused on what you are currently working on. Know what you are doing, define the expected outcome, and pursue it with all your energy. Then you can expect a successful outcome for this project, and the next, and the next, and so on. Doing what you enjoy is a critical component of a successful career, not only in research but pretty much everywhere else.