



No, thanks [Learn more](#)



Research Our research

All Microsoft

Spatial Audio

Established: December 1, 2015

[Overview](#) [People](#) [Publications](#) [Videos](#) [Groups](#) [News & features](#)

Spatial sound is perceived by a listener as emanating from a certain location in space, due to temporal and spectral cues that inform the auditory system about the sound's direction of arrival and distance. Rendering sound spatially, by encoding these localization cues and delivering them to the listener via headphones, allows placing virtual audio sources arbitrarily in the listener's environment. Spatial sound is an integral part of mixed reality applications, including telepresence, gaming, and entertainment.

Head-related transfer functions

The temporal and spectral cues used by our auditory system to determine the direction of arrival of a sound source can be expressed in head-related transfer functions (HRTFs). HRTFs are measurements that capture the directivity patterns of human ears, that is, the way sound, arriving from a certain direction, reaches the left and right ear. HRTFs are a function of source azimuth and elevation, distance, and frequency. Figure 2 illustrates the right HRTF of a subject for the horizontal plane at a distance of one meter. Figure 1 shows the sensitivity of the right ear at 1000 Hz as function of azimuth and elevation at a distance of one meter.




Figure 1




Figure 2


Early applications and their challenges

HRTFs were first used in the 1950s in binaural recordings created by recording sound, e.g., a concert, via two microphones near the ears of a mannequin. Listening to these recordings via headphones creates the illusion of being acoustically present at the recorded event. Technical challenges of these early applications included a smeared acoustical image if the listener's head did not match the mannequin's. In addition, head movements by the listener would move the entire audio scene, which does not happen when listening to a real sound scene and thus may break the binaural illusion.

HRTF personalization



The HRTF of a person can be measured by using a setup similar to the one shown on the left. A set of loudspeakers is rotated around a person wearing small microphones in their left and right ears. Test signals are recorded from each loudspeaker location to measure the spatial directivity patterns of the ears, that is, the person's HRTFs.



This measurement process requires specialized equipment and is time-consuming and cumbersome. To reduce the measurement time or the amount or type of information needed about a subject, personalized HRTFs can be synthesized via acoustic models or machine learning, e.g., from anthropometric features (head width, height, length; ear entrance locations; etc.) or even a crude head scan, as shown on the left. There exists a trade-off between the accuracy of the personalized HRTF and the amount and quality of information known about the user. The challenge in practical applications is to synthesize good-enough HRTFs while not unnecessarily burdening the user with data collection.

Applications for spatial audio

Gaming

Gaming is an ideal application for HRTFs since the 3-D coordinates of individual sound sources are typically available, allowing to collocate visual and auditory sources.

Virtual surround sound

Rendering 5.1 (six channel) or 7.1 (eight channel) surround sound spatially creates a similar audio experience as listening to an actual loudspeaker system. Virtual surround sound can enhance the acoustic experience of games or movies even when using regular headphones.

Mixed reality

Spatial sound is a key feature of many mixed reality applications, as it can enhance the sense of presence and immersion, or create a more realistic experience of virtual content.

Stereo music rendering

Stereo music is intended to be listened to through two loudspeakers in front of the listener. Listening to it with regular headphones places the audio scene between the two ears, inside the listener's head. With spatial audio, the two loudspeakers can be rendered in front of the listener, placing the audio scene in front, where it is supposed to be.

Technology transfers

The Audio and Acoustics Research Group worked closely with our partners in the engineering teams to convert spatial audio research projects to shippable code in various Microsoft products:

- Virtual surround sound support in Windows 10 and in Xbox One.
- The 3D audio rendering engine in [Microsoft Soundscape](#).
- And, of course, the spatial audio engine in [HoloLens](#) – Microsoft's augmented and virtual reality wearable device.



Follow us: [Twitter](#) [Facebook](#) [LinkedIn](#) [YouTube](#) [Instagram](#) [RSS](#)

Share this page: [Twitter](#) [Facebook](#) [LinkedIn](#) [Reddit](#)

What's new

Surface Pro 9

Surface Laptop 5

Microsoft Store

Account profile

Download Center

Education

Microsoft in education

Devices for education

Business

Microsoft Cloud

Microsoft Security

Developer & IT

Azure

Developer Center

Company

Careers

About Microsoft

Surface Studio 2+	Microsoft Store support	Microsoft Teams for Education	Dynamics 365	Documentation	Company news
Surface Laptop Go 2	Returns	Microsoft 365 Education	Microsoft 365	Microsoft Learn	Privacy at Microsoft
Surface Laptop Studio	Order tracking	How to buy for your school	Microsoft Power Platform	Microsoft Tech Community	Investors
Surface Go 3	Virtual workshops and training	Educator training and development	Microsoft Teams	Azure Marketplace	Diversity and inclusion
Microsoft 365	Microsoft Store Promise	Deals for students and parents	Microsoft Industry	AppSource	Accessibility
Windows 11 apps	Flexible Payments	Azure for students	Small Business	Visual Studio	Sustainability

[Sitemap](#)
[Contact Microsoft](#)
[Privacy](#)
[Terms of use](#)
[Trademarks](#)
[Safety & eco](#)
[Recycling](#)
[About our ads](#)
[© Microsoft 2023](#)