DOI: 10.61782/fa.2025.0605

NEW RELEASE OF THE BINAURAL RENDERING TOOLBOX: INTRODUCING VERSION 2.0

Daniel González-Toledo[1] **María Cuevas-Rodríguez**[1] **Luis Molina-Tanco**[1]
Lorenzo Picinali[2] **Arcadio Reyes-Lecuona**[1*]
[1] Telecommunication Research Institute - TELMA, University of Málaga, Spain
[2] Imperial College, London, United Kingdom

ABSTRACT

The Binaural Rendering Toolbox (BRT) is an open-source software suite for psychoacoustic research and headphone-based immersive audio applications. It includes the BRT C++ library, which provides a modular framework for building dynamic and reproducible virtual acoustic scenarios, and BeRTA, a standalone renderer controllable via Open Sound Control (OSC). Researchers can use the BRT for psychoacoustic and auditory experiments without requiring custom implementations. This paper presents Version 2.0, now released under the GPL license.
Key updates in Version 2.0 include support for sound sources directivity, a hybrid rendering mechanism combining BRIRs convolution (also with Ambisonics encoding options), and simulation of free-field propagation, including attenuation and delay. Furthermore, the BRT introduces the ability to apply filters to the binaural signal, enabling headphone compensation and the simulation of hearing protection devices. By providing a comprehensive overview of the features and capabilities of BRT Version 2.0, this paper highlights its potential to drive innovation and reproducibility in auditory research and immersive audio applications.

Keywords: *binaural rendering, spatial audio, psychoacoustic, immersive audio.*

[*] ***Corresponding author***: *areyes@uma.es*

1. INTRODUCTION

Binaural rendering plays a fundamental role in spatial audio research, in virtual reality, as well as in psychoacoustic experiments, allowing for realistic auditory scene simulation over headphones. However, many existing binaural rendering solutions operate as black-box systems with limited customization, making it difficult to ensure full control, transparency, and reproducibility in scientific applications.

The Binaural Rendering Toolbox (BRT) Library Version 2.0 has been developed as an open-source, modular framework for real-time binaural rendering, designed to address these limitations. This library integrates and expands upon the foundations established in the 3D Tune-In Toolkit [1], which served as a precursor and from which some key spatial processing algorithms and functionalities have been inherited and integrated. A description of the early stages of development was presented in [2]. The BRT Library provides a fully configurable and open processing chain, allowing users to assemble custom rendering pipelines by combining source models, listener models, and environment models. By supporting standardized spatial audio formats such as SOFA, the BRT ensures that researchers can precisely document and share their rendering configurations, facilitating the replication and validation of spatial audio experiments.

A key aspect of reproducibility in auditory research is the ability to fully trace and share the entire rendering process, including input parameters, spatialization methods, and output data. The BRT's transparent, open-source structure ensures traceability, allowing users to export standardized audio outputs with detailed annotations for each experiment. Moreover, future developments of BRT aim to strengthen this approach by enhancing interoperability with tools such as the Auditory Modelling Toolbox (AMT) [3], further integrating binaural rendering into computational auditory research.

This paper introduces Version 2.0 of BRT which includes, among other tools, the BRT Library and the BeRTA Renderer, a standalone application which can be remotely controlled via the Open Sound Control (OSC) protocol [4]. The paper is structured as follows: Section 2 describes the overall architecture of BRT Library, followed by Section 3, which details the different models included in the library, such as the source, environment, and listener models, as well as the binaural filtering. Section 4 presents the BeRTA Renderer standalone application and its control interfaces. Finally, Section 5 discusses the broader implications of BRT for reproducible research in binaural audio and spatial hearing experiments.

2. ARCHITECTURE OF THE BRT LIBRARY

The BRT Library is designed as a modular C++ framework which follows a layered architecture:

- **Low-Level Core.** Provides shared data structures, connectivity mechanisms, math utilities, and file I/O (including SOFA support). This layer establishes the foundation on which higher-level modules are built.
- **Processing Modules.** Houses the main signal processing components, like convolution engines, filters (e.g., headphone compensation), near-field effects, and so on. Each module follows a common interface, enabling easy combination or substitution of different implementations.
- **High-Level Models.** Bundles the low-level modules into cohesive "models" (e.g., a Listener Model or Source Model). Listener models might integrate HRTF management, user tracking, and headphone equalization; source models might combine directivity, and environment models include distance attenuation and reverb processing. These models act as building blocks for complex audio scenarios.

A user who does not need to create new models or algorithms can easily employ the high-level models, connecting them in a straightforward manner to customize a binaural renderer according to their application needs. In this sense, a key aspect of BRT's flexibility is its observer design pattern, which propagates updates (e.g., new listener positions) to all dependent modules in real time. This architecture ensures consistent synchronization across the entire rendering chain without requiring manual reconfiguration.

3. MODELS IN THE BRT LIBRARY

As outlined in the previous section, the BRT Library follows a modular architecture, allowing users to construct fully customizable audio rendering chains by combining different building-block models. Each model encapsulates a specific aspect of the binaural simulation, such as sound sources, environments, or listeners. This enables the creation of highly adaptable simulators, tailored to different research and application needs. Users can concatenate, replace, or modify models within the processing chain, defining their own configurations for spatial audio rendering. Whether integrating room acoustics, head-tracking mechanisms, or real-time HRTF switching, the BRT Library provides the necessary tools to assemble custom binaural renderers.

Figure 1 shows a diagram that illustrates the different kind of models and its connections. Each model contained in the BRT Library is described below.

3.1 Sound Source Models

In the BRT Library, each monaural sound source in a virtual acoustic scene requires an instance of a sound source model. These manage the position and orientation of each source, ensuring that their spatial transformation is properly integrated into the rendering pipeline. The modular architecture of BRT allows each sound source model to dynamically communicate changes in its state to all relevant processing components, facilitating seamless updates in real-time simulations. Two sound source models are currently implemented in BRT:

- **Omnidirectional Source Model.** This is the simplest source model, which directly connects the input signal to the rest of the processing chain.

- **Directional Source Model.** This model implements source directivity rendering, applying directional filtering based on a predefined SOFA file that encodes directivity information. During rendering, the input signal is convolved with the appropriate filter corresponding to the direction of emission relative to the listener's position. This requires the source model to continuously track

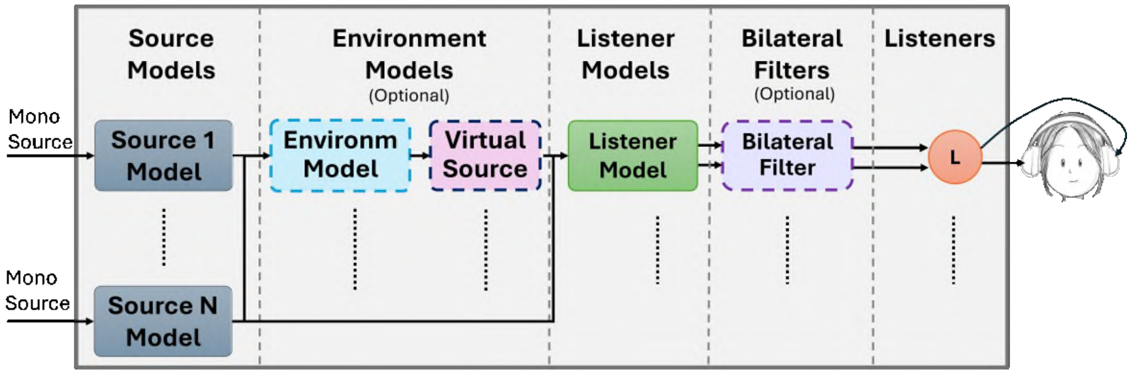

Figure 1. Generic connection diagram of the BRT Library models.

the listener's location, ensuring the correct filter is applied at each frame. The directivity rendering can be enabled or disabled dynamically.

Users integrating BRT into their applications are responsible for updating the position and orientation of each source as needed, and for providing the frame-by-frame audio stream to be rendered. Thanks to the subscription-based architecture, each model seamlessly reacts to spatial updates, ensuring accurate binaural reproduction in dynamic environments.

3.2 Environment Models

Environment models in the BRT Library are optional components that simulate sound propagation in a given acoustic environment. When no environment model is used, sound sources are directly connected to the listener without any additional processing. However, when an environment model is present, the audio stream from the source is processed before reaching the listener model, enabling realistic spatialization effects such as delay, attenuation, and reverberation. Currently, the library includes two environment models:

- **Free-Field Propagation Model.** This model simulates sound propagation in an open-space environment, incorporating both distance-based attenuation and propagation delay. A distance dependent attenuation is applied to the source based on the acoustic power law. The propagation delay is dynamically adjusted, meaning that if either the source or listener move, the delay compresses or stretches smoothly using a Doppler effect simulation. This dynamic delay system ensures that coherent sound sources remain correctly aligned, allowing for accurate simulation of image sources and their respective propagation delays.

- **Scattering Delay Network (SDN) Model.** Based on the SDN algorithm proposed by De Sena [5], this model simulates room acoustics by placing virtual sources on the walls of the simulated room, representing the first-order reflections of the original sound source. Additionally, the reflections are interconnected using a waveguide-based feedback delay network, modelling the late reverberation of the space. The current implementation supports rectangular (shoebox-shaped) rooms, with absorption properties defined for each wall.

As a roadmap for future environment models, an Image Source Method (ISM) model is being actively developed within the BRT to simulate early reflections in arbitrarily shaped rooms [6]. This model will leverage the propagation delay system used in the free-field model and, in future iterations, will be combined with other reverberation models to create hybrid approaches for late reverberation rendering.

3.3 Listener Models

Listener models in the BRT Library define how sound is perceived at the listener's position, processing spatial cues through binaural filtering and dynamic adaptation. These models provide a flexible framework for applications to update the listener's position and orientation dynamically. BRT currently supports four main approaches for binaural rendering within its listener models:

- **Direct HRTF Convolution.** Each sound source is individually convolved with the corresponding Head-Related Transfer Function (HRTF) based on its relative position to the listener.
- **Direct BRIR Convolution.** Instead of using HRTFs, this approach convolves each source with a Binaural Room Impulse Response (BRIR), incorporating the effects of room acoustics into the binaural rendering.
- **Ambisonic Virtual Loudspeakers.** Sources are first encoded into a first-, second-, or third-order Ambisonic sound field, then decoded on a virtual loudspeaker grid, convolving each loudspeaker with the corresponding HRTF. This potentially simplifies rotations, and allows a reduction of the computational cost when a large number of sources are used.
- **Reverberant Virtual Loudspeakers (RVL) [7] .** Similar to the Ambisonic Virtual Loudspeakers method, this model encodes sources into an Ambisonic field, but the conversion to binaural is done convolving each virtual loudspeaker channel with a BRIR, enabling efficient rendering for multiple sources within complex room acoustic conditions.

A key feature of the BRIR-based models is the ability to apply customizable windowing (fade-in and fade-out) to the impulse response. This windowing is implemented as a raised-cosine window with adjustable start time and slope, allowing precise control over which portions of the BRIR contribute to the rendering.

It is important to note that BRIR-based listener models are not purely listener models in the strict sense, as they inherently include environment effects. However, within the BRT Library we consider them as "extended listener models with environment", since their structure and functionality closely resemble those of standard HRTF-based listener models. Their classification as listener models is primarily an architectural decision, allowing seamless combination of multiple listener models to handle different parts of the impulse response. This flexibility makes it possible to employ hybrid models, such as mixing direct path rendering with distinct processing for early reflections and late reverberation, ensuring a modular and scalable approach to binaural rendering.

Examples might include combining the direct path, processed using HRTF convolution, while the room response is rendered using first-order Reverberant Virtual Loudspeakers with BRIR convolution. Another possibility is a hybrid multi-stage reverberation, where separate models could handle the direct path, early reflections, and late reverberation with different spatial resolutions, optimizing realism and computational efficiency.

Regardless of the selected approach, BRT's listener models based on HRTF convolution include additional processing capabilities that enhance spatial realism:

- **Near-Field Simulation**. Distance-dependent filtering is applied to enhance the perception of sources in close proximity, adapting HRTF and BRIR characteristics dynamically [8].
- **Acoustic Parallax Correction.** The rendering pipeline accounts for differences in the position of the ears while the HRTF is referred to the centre of the head. This correction is based on calculating the projection of the vector from the ear to the source on the HRTF sphere (i.e. the sphere on the surface of which the HRTF was measured), giving a more accurate rendering, especially for near-field and far-field sound sources [8].
- **Independent ITD Processing**. BRT allows separate processing of Interaural Time Difference (ITD), either extracted from the SOFA HRTF file or adjusted independently, for example according to the listener's head circumference, enabling more flexible spatialization models.

3.4 Bilateral Filters

The binaural filter stage in the BRT Library is an optional processing module applied to the output of the listener models. If no filters are applied, the signals from the listener models are directly delivered to the listener. However, when enabled, this module processes the left and right channels of the binaural signal, allowing for additional auditory transformations.

Binaural filters can serve a variety of purposes, depending on the application. For instance, by applying specific attenuation profiles, this module can simulate the effect of protective hearing devices [9]. Another example is

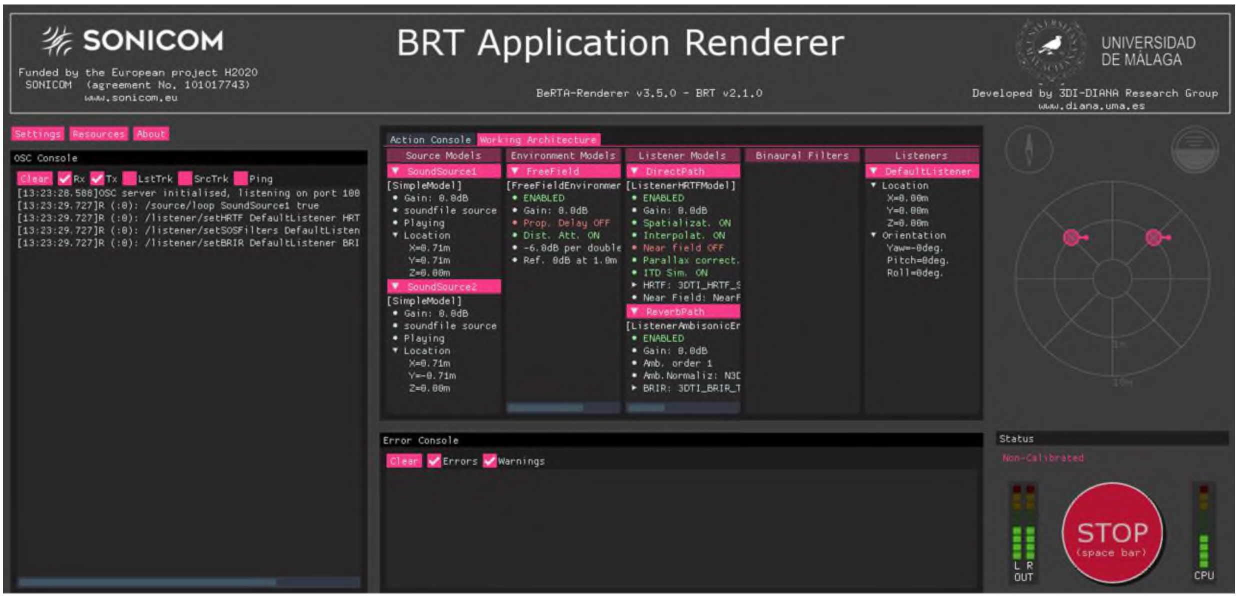

Figure 2. BeRTA Renderer screenshot.

headphone compensation, applying equalization filters to correct for the frequency response of specific headphones. In future developments, more advanced filters, including nonlinear processing, could be used to model different types of hearing impairment, replicating perceptual effects for research or accessibility applications.

4. APPLICATIONS DISTRIBUTED WITH BRT

The BRT Library is a C++ header-only library designed for real-time binaural rendering. It is intended for developers creating audio applications for managing sound sources, either from external devices or files, handling audio interfaces and generating audio streams. These are processed by BRT Library to produce spatialized binaural output. However, since this library is not a standalone audio engine, it does not provide a ready-to-use solution for users who do not wish to develop their own applications.

To address this limitation, the Binaural Rendering Toolbox Application (BeRTA Renderer in short) provides a standalone solution for binaural rendering. It embeds the BRT Library and is capable of processing incoming audio streams to perform real-time binaural rendering, responding to control commands via OSC. BeRTA manages audio input from both line-in and audio files, processing monophonic sources and generating spatialized binaural output.

¡Error! No se encuentra el origen de la referencia. illustrates the BeRTA Renderer interface. BeRTA does not include an interactive graphical interface; all user interaction occurs through OSC commands, making it highly adaptable for real-time processing, automation, and experimental setups. The interface consists of multiple panels: the left console shows sent and received OSC commands, the upper central console displays application settings and other data such as the position of the listener or sources, the lower central console reports rendering errors and OSC communication issues, and the right panel features a radar visualization of the rendered sources, their orientation and their spatial positions.

In addition to its core rendering capabilities, the BeRTA Renderer includes several advanced features to ensure accuracy and safety in audio output, such as the calibration of its output levels, ensuring full control over the absolute sound pressure level delivered to the user. This is particularly relevant in research and safety-critical applications where exact auditory exposure must be controlled. Furthermore, to prevent excessively loud sounds, BeRTA provides a configurable signal limit, ensuring that the output remains within safe thresholds and protecting users from unintended high-volume exposure. Finally, another feature of BeRTA Renderer is its ability to generate a log file annotated with the listener's positions and orientations, as well as those of all contributing sources. This annotated audio is stored in a custom SOFA file[1], facilitating the sharing of experimental stimuli.

The following list introduces the different groups of OSC commands that the BeRTA Renderer application interprets to dynamically configure the virtual scene. For a complete reference of all commands, please, refer to the online documentation[2].

- **OSC Control commands** allow to establish and manage connections with BeRTA over the network or locally. Some examples are: `/control/connect` to open an OSC connection to the indicated IP and port to send the information back or `/control/sampleRate` to ask for the sample rate which BeRTA is using.
- **OSC Overall Commands** perform actions which affect the whole audio scene, such as `/play` to reproduce all the sound sources and `/enableModel` to activate or deactivate a specific model.
- **OSC Resources Commands** are responsible for managing the resources used in the virtual scene, such as HRTFs, Directivity TFs, SOS Filters and BRIRs. Two examples of these commands are: `/resources/loadHRTF`, which loads a new HRTF from a specified SOFA file path and `/removeHRTF`, which removes an HRTF from the loaded resources.
- **OSC Source Commands** are responsible for controlling the audio sources within the virtual scene. Two examples of these commands are: `/source/loadSource`, which adds a new source by loading a sound file, and `/source/location`, which sets the location of the source in global x, y, z coordinates.
- **OSC Listener Commands** are used to control the listener, including parameters for position, orientation, and simulation models. Some examples are: `/listener/location`, which sets the listener's location in global x, y, z coordinates and `/listener/enableSpatialization`, which enables or disables the simulation of spatial audio.
- **OSC Environment Commands** are responsible for controlling the environment. Different commands are available depending on the environmental model being used. For example, `/environment/enablePropagationDelay` enables or disables the simulation of propagation delay throughout the entire environment for the free field environment model, and `/environment/setWallAbsorption` sets the absorption coefficients of a wall when the SDN environmental model is being used.
- **OSC Binaural Filter Commands** are responsible for controlling the binaural filters using second order section filters. In this case, `/binauralFilter/setSOSFilter` sets the resource (the SOS filters) to the simulation of the binaural filter model.

In addition to the BeRTA Renderer, a series of supporting applications, components and documentation resources are provided to facilitate interaction with the BRT Library. These tools allow users to access most of BRT functionalities without the need to develop custom audio applications.

- **BeRTA GUI:** a standalone application providing a graphical interface for sending OSC commands controlling the BeRTA Renderer's most important functionalities.
- **Unity 3D OSC Package:** a pre-configured Unity package streamlining the process of sending OSC commands from Unity 3D, making it easier to integrate BeRTA Renderer into interactive and VR-based applications.
- **PureData Patch:** a ready-to-use PureData patch implementing the mechanisms for generating and transmitting OSC commands to BeRTA Renderer, providing an accessible workflow for users familiar with Pure Data.
- **Online Documentation**[3]: a comprehensive resource detailing all models and components of the BRT Library, including configuration and

[1] Currently available as a proposed convention: https://www.sofaconventions.org/mediawiki/index.php/AnnotatedReceiverAudio
[2] https://grupodiana.github.io/BRT-Documentation/osc/
[3] https://grupodiana.github.io/BRT-Documentation/

usage guidelines. It also provides instructions for handling and configuring the different applications, an exhaustive list of OSC commands, and a collection of tutorials and examples for both the library and its applications

5. CONCLUSIONS

The Binaural Rendering Toolbox (BRT) is a suite designed to support scientific research in spatial audio and psychoacoustics. At its core, the BRT Library Version 2.0 provides a modular framework for real-time binaural rendering, enabling researchers to build reproducible virtual acoustic scenarios.

One of the key aspects of the BRT Library is its emphasis on reproducibility. By using standard formats such as SOFA for HRTFs, BRIRs, and directivity functions, it ensures a fully traceable rendering process. Combined with the standalone BeRTA renderer and supporting applications, the BRT enables researchers to document and share their exact rendering configurations, including output signals, spatialization parameters, and listener models, allowing easy replication of experiments across different laboratories.

Furthermore, the OSC-based control system implemented in BeRTA and its supporting applications, allows researchers to automate experiments and maintain scripting capabilities, ensuring consistency across trials and test conditions without requiring custom development. Because the entire processing pipeline is open source, all algorithms used in BRT Library are fully inspectable and modifiable, addressing the challenges of methodological transparency in experimental auditory research.

Future developments will focus on expanding the range of available models, integrating more advanced filtering techniques, and further optimizing computational efficiency for large-scale, real-time applications. Another key area of future development is the integration of BRT with the Auditory Model Toolbox (AMT). In this scenario, BRT could generate binaural outputs that serve as input for auditory models in AMT, allowing the system to simulate perceptual decision-making and dynamically modify the virtual acoustic environment based on the model's response. This would enable the development of interactive auditory simulations where a computational auditory model interacts with a BRT-rendered scene, adjusting conditions in real time. By maintaining its open and modular approach, the BRT suite will continue to evolve as a flexible and reliable tool for the scientific community, fostering new research in binaural perception, hearing augmentation, virtual reality, and immersive audio applications.

6. ACKNOWLEDGMENTS

This research has been supported by the SONICOM project (https://www.sonicom.eu/) under grant 101017743, and the Spanish National Project SONIX (PID2023-152547NB-I00).

7. REFERENCES

[1] Cuevas-Rodríguez, M., Picinali, L., González-Toledo, D., Garre, C., de la Rubia-Cuestas, E., Molina-Tanco, L., & Reyes-Lecuona, A. (2019). 3D Tune-In Toolkit: An open-source library for real-time binaural spatialisation. PLOS ONE, 14(3), e0211899. https://doi.org/10.1371/journal.pone.0211899 .

[2] González-Toledo, Daniel, et al. "The Binaural Rendering Toolbox. A virtual laboratory for reproducible research in psychoacoustics." *Forum Acusticum*. 2023. Turin (Italy).

[3] Majdak, P., Hollomey, C., & Baumgartner, R. (2022). AMT 1. x: A toolbox for reproducible research in auditory modeling. Acta Acustica, 6, 19.

[4] Wright, M., Freed, A., & Momeni, A. (2017). 2003: Opensound control: State of the art 2003. A NIME Reader: Fifteen Years of New Interfaces for Musical Expression, 125-145.

[5] de Sena, E., Hacihabiboglu, H., & Cvetkovic, Z. (2011, February 2). Scattering Delay Network: An Interactive Reverberator for Computer Games. 41st International Conference: Audio for Games.

[6] Gutierrez-Parera, P., Arrebola, F., Cuevas-Rodriguez, M., Gonzalez-Toledo, D., Molina-Tanco, L., & Reyes-Lecuona, A. (2025). *Exploring simple methods for room absorption estimation in early reflections for binaural audio rendering*. In *Proceedings of Forum Acusticum Euronoise 2025*, Malaga, Spain.

[7] Engel, I., Henry, C., Amengual Garí, S. V., Robinson, P. W., & Picinali, L. (2021). Perceptual implications of different Ambisonics-based methods for binaural reverberation. The Journal of the Acoustical Society of America, 149(2), 895-910.

[8] Romblom, D., & Cook, B. (2008). Near-Field Compensation for HRTF Processing. Audio Engineering Society. http://www.aes.org/e-lib/browse.cfm?elib=14762.

[9] Cuevas-Rodriguez, M., Gonzalez-Toledo, D., Lawson, A., Keren, N., & Reyes-Lecuona, A. (2025). Addressing acoustic simulation pitfalls in virtual reality for occupational safety research. In Proceedings of Forum Acusticum Euronoise 2025, Malaga, Spain.