



**Xylo™ Audio 3**  
**Development Kit**  
**Manual**  
**March 2026**

# Content

- 1. Product definition ..... 1
  - 1.1. Xylo™Audio 3 ..... 1
  - 1.2. Xylo™Audio 3 Development Kit ..... 2
- 2. Features ..... 3
  - 2.1. Xylo™Audio 3 ..... 3
  - 2.2. What is on Xylo™Audio 3 ..... 3
  - 2.3. About the Xylo™ family ..... 4
- 3. Getting started ..... 5
  - 3.1. Installation ..... 5
  - 3.2. Quick start with Xylo™ ..... 5
  - 3.3. Troubleshooting ..... 6
  - 3.4. Application pipeline for Xylo™Audio 3 ..... 7
  - 3.5. Rockpool documentation for Xylo™Audio 3 ..... 8
  - 3.6. Next steps ..... 8
- 4. Mechanical specifications ..... 9
  - 4.1. Mother board ..... 9
  - 4.2. Daughter board ..... 12
- 5. Change log ..... 14





# 1. Product definition

## 1.1. Xylo™ Audio 3

Xylo™ Audio 3 is fully digital ultra-low-power AI device, which accepts audio input and performs real-time inference. It includes a dedicated audio encoding block compatible with PDM-encoded digital microphone input, and a low-power inference core (the Xylo™ Core), using sparse quantised spiking neural networks (SNNs) for real-time sensory inference.

Xylo Core is a highly configurable SNN processing core, supporting a wide range of flexible network architectures (feed-forward, recurrent, residual, etc.).

The inference core in Xylo™ Audio 3 supports up to 16 spiking audio encoded input channels; input projection weight dimensions of up to 16x128; 992 configurable hidden neurons; output projection weight dimensions of up to 128x32; and up to 32 readout channels for classification purposes. Xylo™ Audio provides computationally powerful current-based Leaky Integrate and Fire (CUBA-LIF) spiking neurons, with fully-configurable per-synapse leak, bias and threshold parameters.




Figure 1 . Xylo™ Audio 3 chip

## 1.2. Xylo™ Audio 3 Development Kit

The Xylo™ Audio 3 Dev Kit is powered by the SynSense Xylo™ Audio 3 chip, which brings the flexibility of highly configurable SNN processing to microwatt energy budgets. With appropriate network design and training, you can use Xylo™ Audio 3 to verify a range of applications with the on board analog microphone or digital microphone, for example:

- Key word spotting (KWS)
- Ambient sound detection, e.g. baby crying in noisy environments
- Voice activated intelligent personal assistants

Development of these application prototypes is straightforward, using the SynSense python-based toolchain [Rockpool](#).



Figure 2 . Xylo™ Audio 3 Development Kit



## 2. Features

### 2.1. Xylo™ Audio 3

- 2-wire I2C slave (up to Fast-mode Plus, 1MHz)
- PDM interface for digital microphone
- Supports up to 992 hidden neurons and up to 32 classification output channels
- 1-bit configurable interrupt indicating classification / detection done or 3 bit direct output indicating triggered classification type
- Up to 50 MHz internal operating frequency
- Extremely low memory footprint (~124 KB), memory power control granularity down to 2 KB
- On-chip one-time programming block

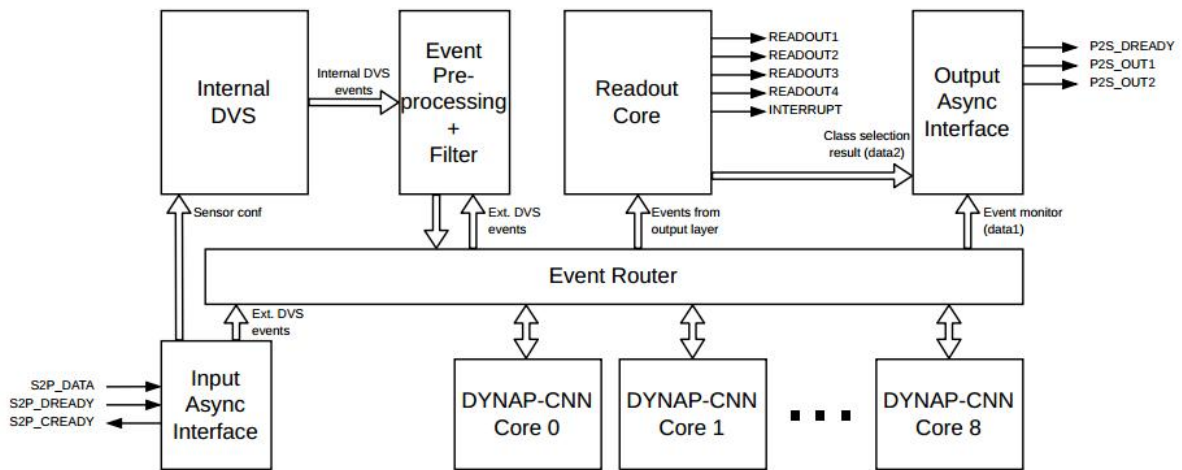
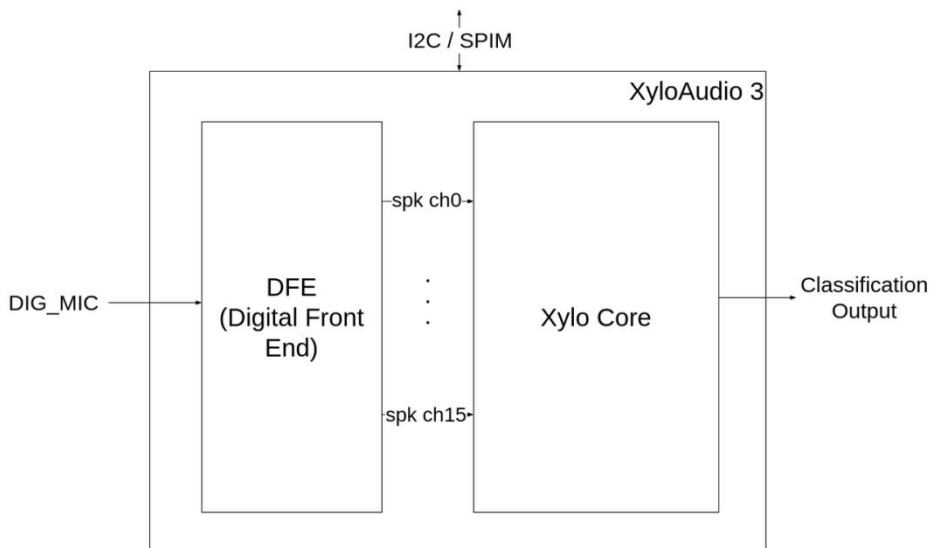


Figure 3: Top Level Chip Diagram

### 2.2. What is on Xylo™ Audio 3?

Xylo™ Audio 3 consists of a PDM interface which receives digital audio data, e.g. from a digital microphone, and the Xylo™ core for inference.



## Xylo™ Spiking Neural Network (SNN) inference core

The Xylo™ SNN core performs inference on a stream of sensory data. It provides up to 16 channels of quantised event-encoded input, and 32 channels of binary event-encoded output.

- Convenient training, quantisation and configuration with rockpool.ai
- 1024 leaky integrate-and-fire (LIF) model spiking neurons
  - 992 LIF hidden neurons
  - 32 LIF output neurons
- 15 events per time-step for input channels and hidden neurons
- 1 event per time-step for output neurons / output channels
- Maximum fan-out of 1024 per hidden neuron
- Mean fan-in of 32 per hidden neuron
- Maximum input projection to 128 hidden neurons
- Maximum output projection from 128 hidden neurons

## 2.3. About the Xylo™ family

Xylo™ is a family of ultra-low-power sensory processing and classification devices, providing efficient sensor interfaces for a range of modalities, and relying on low-bit-depth, sparse neural networks for efficient inference.

Xylo™ devices comprise an efficient sensory encoding front-end, a Xylo™ SNN inference core with on-board memory, and control logic.

Xylo™ is currently available in several versions with varying HW support and front-ends. The various devices in the family interface directly with analog and digital audio sensors (Xylo™ Audio); IMU motion devices (Xylo™ IMU); and other sensory input modalities.



## 3. Getting started

### 3.1. Installation

Xylo™ Audio 3 is accessed via the Python library Rockpool. This has several dependencies, which are easiest to manage if you use a Python environment tool such as [conda](#), [pyenv](#) or similar.

The Xylo tools currently support Python  $\geq 3.8$  and  $\leq 3.11$ , and require Linux or MacOS.

#### Create a Python environment

With conda:

```
conda create --name xyloa3 'python<=3.11'
```

#### Install Rockpool

To facilitate the development of audio processing applications and their deployment into Xylo™ Audio 3 HDK, we offer the open-source Python library “Rockpool”.

Rockpool allows you to build networks, simulate, train, test, and deploy them in simulation or direct to one of our HDKs.

To get started, first install Rockpool by executing:

```
pip install 'rockpool[xylo]'
```

This will install Rockpool, along with all dependencies required to interface with Xylo.

For more information about installing and using Rockpool, see the documentation at <https://rockpool.ai>.

#### Accessing Xylo™ Audio 3 HDK in Linux

Our software tools need specific permissions from Linux computers to recognize and detect Xylo™ Audio 3 devices.

Make sure your Xylo™ Audio 3 device is connected to a computer through a USB 3 port and that udev rules are updated, following this documentation:

<https://synsense-sys-int.gitlab.io/samna/0.48.4/install.html#udev-rules-on-linux-systems>

### 3.2. Quick start with Xylo™

Activate your Python environment, and open a python terminal. Then enter the following



commands:

```
# - Import the Xylo HDK detection function
from rockpool.devices.xylo import find_xylo_hdks

# - Detect a connected HDK and import the required support package
connected_hdks, support_modules, chip_versions = find_xylo_hdks()
```

`connected_hdks` will be a list of device objects, for each Xylo HDK you have connected.

`support_modules` will be a list of python packages, corresponding to each HDK, which provide all the support functions and classes for that connected HDK.

`chip_versions` will be a list of strings, each one describing the version of Xylo on each connected HDK.

If these lists are empty, the Xylo HDK was not found and you need to debug. Use `lsusb` on a linux machine, or the "System Information" app on Max OS, to see if the device was detected by your computer.

### 3.3. Troubleshooting

#### The LED indicator lights don't turn on on the Xylo HDK

- Try a new USB cable, a different USB port, or a different PC
- **Are you using a USB3.1 compatible cable?**
- Xylo HDK requires USB 3.1 to work. Try a known-good cable; ideally the one shipped with Xylo.
- **Are you using a USB 3.1 compatible port?**

Xylo HDK requires USB 3.1. Check that Xylo is not connected via a USB 2 hub or USB 2 port. Use `lsusb` (linux) or System Profiler (Mac OS) to check.

`find_xylo_hdks()` reports an error `ValueError: The firmware of the connected Xylo HDK is unsupported, and must be upgraded.`

- **You need to upgrade the firmware on Xylo**

The firmware for all our development kits is available here:

[https://www.dropbox.com/scl/fo/srjp39ljp4vwhokjm2j68/ACPOOrCj\\_gq6MTCsEoTBfKw?](https://www.dropbox.com/scl/fo/srjp39ljp4vwhokjm2j68/ACPOOrCj_gq6MTCsEoTBfKw?)

[rkey=ey5mfefd3u3ybjvxwu2qcux&e=1&st=xo8v4k88&dl=0](https://github.com/synsense/synsense-int.gitlab.io/samna/0.48.4/flasher.html)

For Xylo™ Audio 3, the firmware is under the directory “Xylo™ Audio 3 Development Kit”

Please check the samna documentation for updating firmware on your device:

<https://synsense-sys-int.gitlab.io/samna/0.48.4/flasher.html>

**Xylo doesn't show up in Rockpool using `find_xylo_hdks()`**

- **If on linux, have you set up your Udev rules correctly?**

Check the samna installation documentation.

- **Are you using the latest version of samna?**

In your python environment, upgrade samna with the command `pip install --upgrade samna && python -x 'import samna'`

- **Are you using the latest firmware for Xylo?**

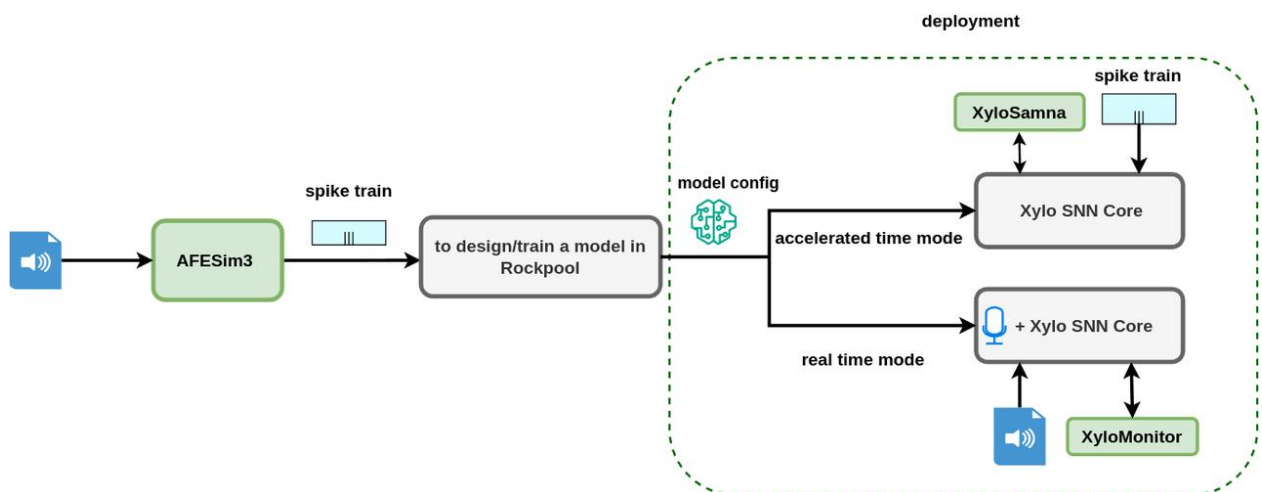
See above for instructions to upgrade the firmware on your device.

### 3.4. Application pipeline for Xylo™ Audio 3

A typical pipeline to build and deploy an audio application for Xylo™ Audio 3 contains the following steps:

1. Preprocessing: converting audio signals to spike trains
2. Designing and training an SNN model with spike-encoded audio data
3. Deploying the trained model in Xylo™ Audio 3 HDK

The following diagram illustrates these steps along with the required tools from Rockpool, highlighted in green:



AFESim3 is a tool in Rockpool that simulates the Audio Front End (AFE) of Xylo™ Audio

3, used as a preprocessing step to convert audio signals to spike trains.

[XyloSamna](#) and [XyloMonitor](#) are Rockpool APIs that interface the user and Xylo™ Audio 3 HDK. These APIs are designed with user-friendliness in mind, making the interaction with Xylo™ Audio 3 a seamless experience.

Please see [https://rockpool.ai/devices/xylo-a3/AFESim3\\_as\\_transform.html](https://rockpool.ai/devices/xylo-a3/AFESim3_as_transform.html), where we elaborated on using AFESim3 as an audio transform to generate spike-encoded samples.

And please check

[https://rockpool.ai/devices/xylo-a3/Using\\_XyloSamna\\_and\\_XyloMonitor.html](https://rockpool.ai/devices/xylo-a3/Using_XyloSamna_and_XyloMonitor.html) for an example of deployment into Xylo™ Audio 3 HDK, using [XyloSamna](#) and [XyloMonitor](#).

### 3.5. Rockpool documentation for Xylo™ Audio 3

For more detailed information and in-depth explanations about Rockpool's features and functionalities for Xylo™ Audio 3, please refer to our online documentation:

<https://rockpool.ai>.

Our documentation is regularly updated and contains step-by-step tutorials and additional tips to help you to make the most of our HDKs.

### 3.6. Next steps

- Quick-start for the Xylo™ SNN core:  
[https://rockpool.ai/devices/quick-xylo/deploy\\_to\\_xylo.html](https://rockpool.ai/devices/quick-xylo/deploy_to_xylo.html)
- Intro to Xylo™ Audio 3: <https://rockpool.ai/devices/xylo-a3/xylo-audio3-intro.html>
- General introduction to Rockpool: [https://rockpool.ai/basics/getting\\_started.html](https://rockpool.ai/basics/getting_started.html)
- Training and deploying a model to Xylo™:  
<https://rockpool.ai/devices/torch-training-spiking-for-xylo.html>





1. System Power LED
2. FPGA CFG Done Indicator
3. USB3.0 Controller State Indicator
4. SoC Power Traces State LEDs
5. Debug State Indicator
6. USB3.0 Controller CFG Switch (RSV)
7. System Reset Key
8. USB3.0 Micro-B Port
9. FPGA JTAG (RSV)
10. High Precision Power Monitor
11. Flash
12. USB3.0 Controller
13. FPGA

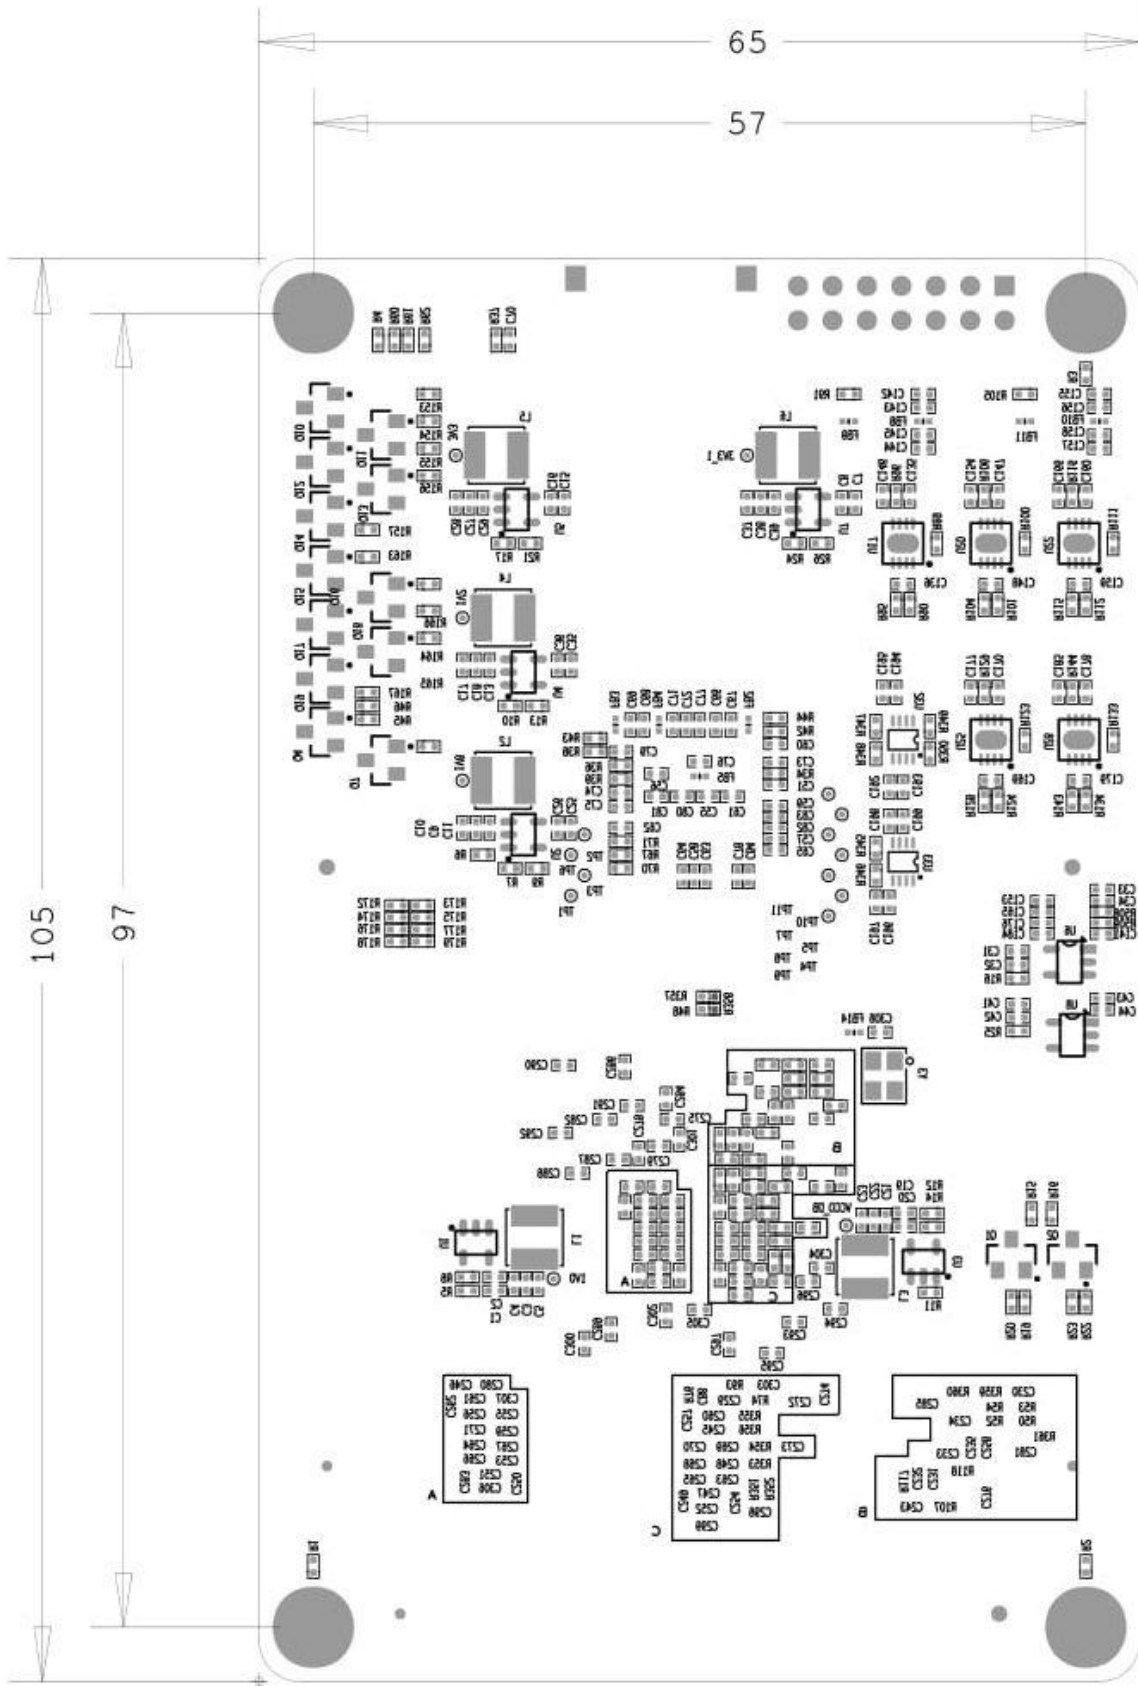


Figure 4 . Back view of the FPGA mother board (mm)

## 4.2. Daughter board




Figure 5 . Front view of the Xylo™Audio 3 daughter board

1. Power ON LED
2. Two channel digital microphones
3. Switch to control use which microphones
4. Debug switch
5. EEPROM
6. Xylo™Audio 3





Figure 6 . Back view of the Xylo™Audio 3 daughter board



## 5. Change log

No.	Version	Date	Editor	Changes
1	V1.0	2024.09		Initial Version
2	V2.0	2024.10	Dylan	Additional quick-start details. Moved Quick-start section up.



**The information contained herein is for informational purposes only, and is subject to change without notice.**

---

### **Intellectual Property Rights**

SynSense owns the copyrights, trademarks and other intellectual property rights and interests in this document. The fact that SynSense provides this document to you does not affect the rights and interests of SynSense as described above.

Brand and product names are trademarks or registered trademarks of their respective owners. No license, including implied or arising by estoppel, to any intellectual property rights is granted by this document.

### **No Warranty**

While every precaution has been taken in the preparation of this document, it may contain technical inaccuracies, omissions and typographical errors, and SynSense is under no obligation to update or otherwise correct this information. SynSense makes no representations or warranties with respect to the accuracy or completeness of the contents of this document, and assumes no liability of any kind, including the implied warranties of noninfringement, merchantability or fitness for particular purposes, with respect to the operation or use of SynSense hardware, software or other products described herein.

### **Disclaimer**

To the extent permitted by applicable law, SynSense shall not be liable for any direct, indirect, incidental, special, incidental or other damages, costs, liabilities or claims of any kind arising out of or in connection with the use of this document, with respect to the operation or use of SynSense hardware, software or other products described herein.

### **Applicable Terms and Conditions for Products**

Terms and limitations applicable to the purchase or use of SynSense's products are as set forth in a signed agreement between you and SynSense or in SynSense's Standard Terms and Conditions.





Make Intelligence Smarter