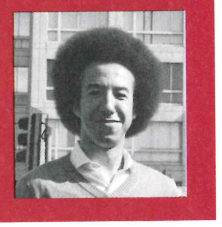

Calvin Robinson

Calvin is Director of Computing & IT Strategy at an all-through school in northwest London.

Resources

- Mycroft https://mycroft.ai/get-mycroft
- Raspberry Pi 3
- microSD card, 8GB or larger
- USB microphone
- Speakers
- Etcher https://etcher.io

Make an open source voice assistant with Mycroft

Forget Cortana, Alexa, Google Home and Siri – we're going open source and creating our own voice assistant

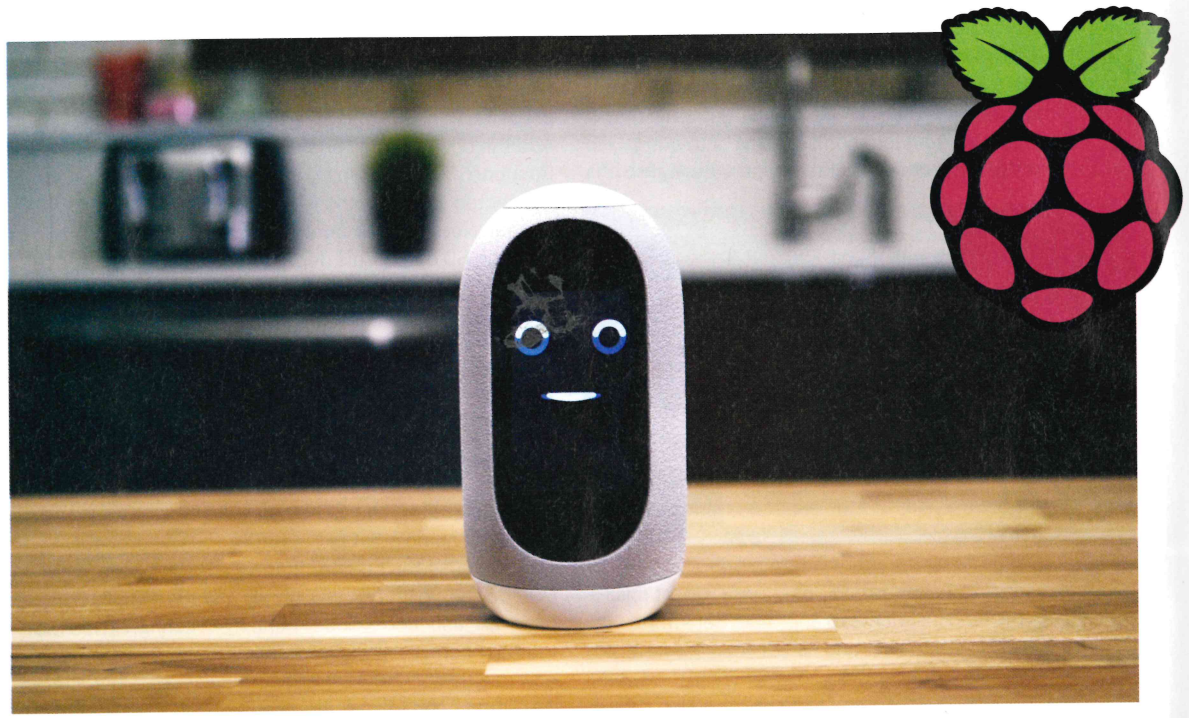

Above Mycroft Mark II, expected in December this year, looks like being an impressive piece of hardware

Voice assistants are all the rage at the moment, what with Microsoft's Cortana, Apple's Siri, Google's Home and Amazon's Alexa all entering the market. Users are becoming more comfortable talking to a device and receiving audible instructions, in a way that's not too dissimilar from the computer in the *Star Trek* franchise. However, with current concerns regarding privacy, it's important to know what data is collected, where it's going, and who could potentially be eavesdropping on your conversations.

We don't mean to sound paranoid, but if you've got an open mic in your environment it's pretty important to know where any data might be heading; some of the larger corporations collect information about users to better target advertisements towards them. That's why users are turning to open source alternatives. Last issue we interviewed John Montgomery, the CEO of Mycroft AI, who has set out to address this problem. This issue we're having a go at building one of these units for ourselves, armed only with a Raspberry Pi, a USB microphone and some speakers.

01 Download and flash

There are Linux (Arch, Fedora and Ubuntu/Debian) and Android versions of Mycroft available, but for this tutorial we're sticking with the Raspberry Pi flavour. We recommend you use a Pi 3.

Download the latest version of Picroft from the link in our Resources section, as well as Etcher. Other than any potential 'Skills' you want to add later on and that should be all you need to download for this tutorial. Etcher is an imaging program which we'll use to burn or flash the downloaded Picroft image to an SD card. Plug your microSD card into your computer, launch Etcher and select the Picroft image. Then flash it!

02 Set it up

Plug your microSD card back into your Raspberry Pi and connect it to a power source. The easiest way to get everything working is to connect your Pi to the local network via the Ethernet port. If you do need to use Wi-Fi, look out for an SSID called MYCROFT; the default password is 12345678.

Once everything is connected, you'll want to either plug in a monitor and keyboard, or connect via SSH to do this headlessly. Whether via Ethernet or Wi-Fi, once your device is connected you'll need to visit **http://home.mycroft.ai** to start the setup process. You'll need to sign in with Google, Facebook or GitHub, or create a new Mycroft account; given that part of the reason for this project is to protect your data from being shared with big corporations, the latter might be advisable!

03 Find your Raspberry Pi

To date, all Raspberry Pi devices start with a MAC address of B8:27:EB, so we can use this to scan our network for the Pi, if we don't have a monitor/keyboard to connect to it. You could use nmap, for example:

```
sudo nmap -sP 192.168.1.0/24 | awk '/^Nmap/
{ip=$NF}/B8:27:EB/{print ip}'
```

You can also use arp:

```
$ arp -na | grep -i b8:27:eb
```

If your home network is not on the 192.168.1.* subnet, change the command line accordingly.

04 Connect to Picroft

SSH into your Picroft and you'll be taken straight into the Mycroft CLI screen. Usually, this is quite useful, but while we get things set up we want to exit that screen using **Ctrl+C** to reach a normal command prompt. Here you'll want to do some basic setting-up. First, change the password: type **passwd** and follow the prompts. Then change the Wi-Fi network settings:

```
sudo nano /etc/wpa_supplicant/wpa_supplicant.
conf
```

Change the network name and/or password in this config file, then press **Ctrl+X** to exit and save the file. Then type **sudo reboot** to reload your Picroft.

05 Add your device to Mycroft.ai

Back in your browser at **home.mycroft.ai**, click 'Add Device' and you'll be asked for a name and location for your Picroft device. You'll also be asked for a registration code; if you turn on the speakers connected to your device you'll notice the Picroft is reading this out to you already, until the device is paired up.

Now that we've paired with Mycroft.ai, we can go to the Settings menu where you can select a male or female voice, the style of measurement units you want to use, and your preferred time/date formats.

If you're concerned about privacy, you may want to keep the Open Dataset box unticked. Keep in mind, though, that selecting this option is a good way of contributing useful data to the open source project and thus improving the performance of Mycroft in the future, assuming your voice assistant isn't in a particularly confidential environment.

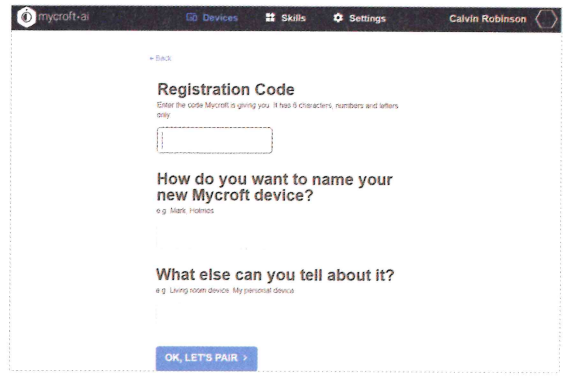

06 Advanced settings

In Advanced Settings, we can really begin to personalise our Mycroft experience. There are a number of pre-programmed wake words, but you can set your own custom version – perhaps 'Computer' a la *Star Trek*, or maybe 'Butler' if you're feeling particularly bourgeoisie.

You'll need to set the phonetic version of your wake word too, so the device understands what it's listening out for. An example would be 'HH EY . B AH T L ER .' for "Hey Butler". You'll probably want to include some kind of exclamation or greeting before your wake word to avoid confusing the Picroft. This is almost certainly why "Hey Google" or "Okay Google" are used on Google Home, rather than just "Google"; it's to avoid said devices picking up on random conversations, something which happened quite a lot in our testing.

You can also switch the text-to-speech engine from Mycroft's Mimic engine to Google's own. This will change the voice you hear to that of Google Home, which is arguably much smoother.

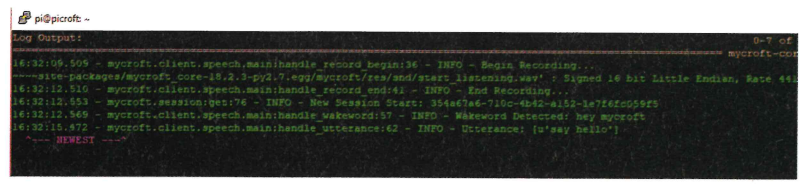

07 Using the Picroft CLI

Now that everything is set up, you should have a basic voice assistant raring to go. Call out the wake word and issue a few commands to get started – Mycroft understands all these examples by default:
Hey Butler, what time is it?

Mycroft, send help

There's an active community on GitHub ready to help with help requests, which will come in handy as Picroft spits out quite a few Python 2.7 errors when Skills refuse to load properly...

Hey Butler, set an alarm for X am.
Hey Butler, record this.
Hey Butler, what is [insert search term].
Hey Butler, tell me a joke.
Hey Butler, go to sleep.
Hey Butler, read the news.
Hey Butler, set a reminder.
Hey Butler, do I have any reminders?
Hey Butler, increase/decrease volume
Hey Butler, what's the weather like?

You can also skip our earlier step of using nmap or arp by asking "Hey Butler, what is my IP address?".

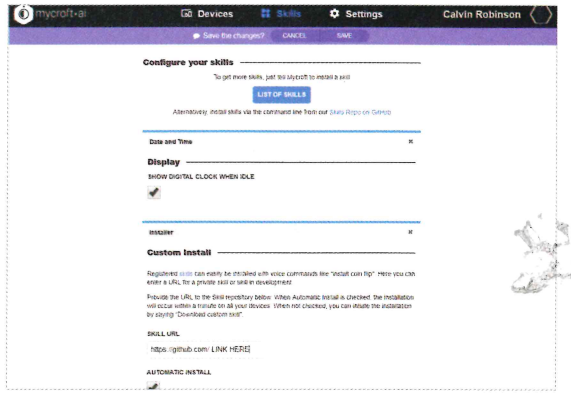

08 Adding Skills

Of course, the default abilities are all well and good, but surely where an open source program comes into its own is with customisation. Mycroft/Picroft is no different in this regard, with a whole range of different voice abilities available. These seem to have been coined 'Skills' – we can thank Amazon for that.

Back at Mycroft.ai, it's time to explore the Skills menu. There's an option to paste a GitHub URL to install a Skill, which is quite useful, but Mycroft does also recognise "install [name of Skill]" as a command. You'll see a link to a list of community-developed Skills, where you can also find the names and command needed to install them. "install YouTube" adds a simple YouTube streaming Skill, for example.

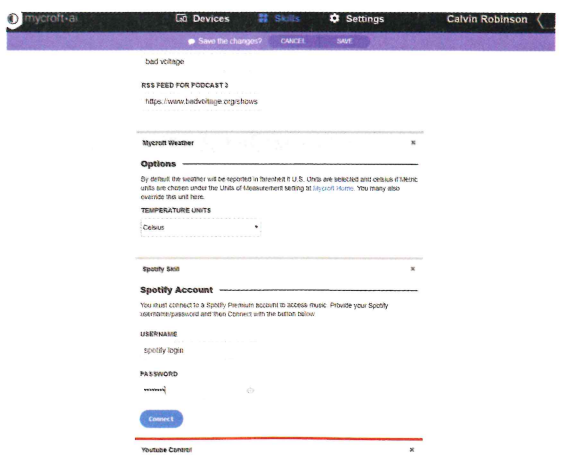

09 Play music

The only officially supported music-playing app seems to be mopidy, which we had great difficulty in getting working. Hours of fiddling with dependencies and an extended deadline later, we still had no luck. However, we did find spotify-skill in the GitHub repository, which works a treat.

Simply by copying the GitHub URL (**https://github.com/forslund/spotify-skill**) into the 'Skill URL' box on Mycroft.ai and ticking 'Automatic Install', moments later we had a new menu option to input our Spotify details. Then "Hey Mycroft, Play Spotify" loads up our most recent playlist. The only problem was that we couldn't figure out a way to stream directly to the Mycroft; spotify-skill only streams music to another Spotify Connect device. It's only speculation on our part, but we assume this is something to do with licensing restrictions for 'official' Spotify devices.

10 Play podcasts and radio

Thankfully, it's much easier to stream podcasts than it is stream music. A quick "install podcast skill" install the necessary Skill, and you'll then have options on Mycroft.ai for your three favourite podcast feeds. Paste in the RSS details and you're good to go. "Hey Mycroft, play x podcast" should then do the trick.

We didn't have as much luck with the Internet Radio Skill, though. Requesting any Internet Radio stations threw up PHP errors, which are visible in the Picroft command line interface and log viewer. It seems as if Skills are very hit-and-miss at the moment. There is a 'status' column for each one on the community page which is meant to indicate its readiness, but we found the results to be inconsistent.

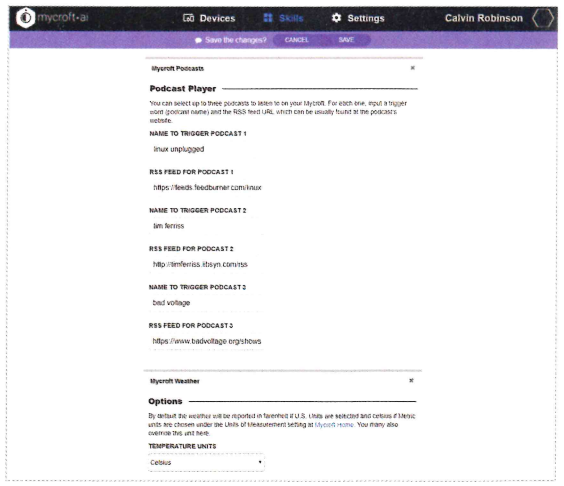

11 Replacing commercial voice assistants

While the Picroft has been a fun experiment to sink (way too many) hours into, do we think it's ready for prime time? In a word: no. While the core experience may be fine, it's extremely limited and the Skills are not yet up to scratch. In our experience, they're just not very likely to work, even after hours of fiddling.

If you're looking for a new hobby and don't mind putting a few days into this, you'll get some enjoyment out of it. However, if you're looking for a new voice assistant to read you the news, wake you up and play your favourite music or radio station, we're still forced to recommend one of the commercial units. Having said that, Mycroft Mark II is available to reserve on IndiGogo right now too.

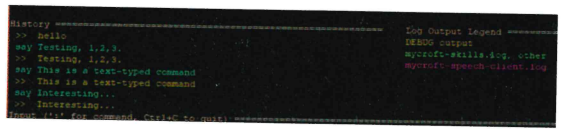

12 Testing

It may be that Mycroft's voice recognition isn't up to scratch, or it may be that the microphone we used for testing was cheap and useless, but constantly issuing commands via voice during the testing process proved to be tiring. Fortunately, Mycroft supports text-based commands, too.

If you SSH into your Picroft you can type text commands directly into the command-line interface. If you exit out of the CLI, there are a number of command prompts available:

`mycroft-cli-client` A command line client, useful for debugging
`msm` Mycroft Skills Manager, to install new Skills
`say_to_mycroft` Use one-shot commands via the command line
`speak` Say something to the user
`test_microphone` Record a sample and then play it back to test your microphone.

13 Becoming a supporter

Mycroft offers an optional subscription service, at $1.99 per month or $19.99 for a year. While the primary purpose of these subscriptions is to support the development team, there are exclusive updates which are made available only to subscribers.

As of May 2017 there's a new female voice available only to supporters. There was also a Group Hangout session with John Montgomery himself in April. The missions statement reads:

"It's hard to overstate how much I value your support Calvin Robinson. It allows my team to make me grow, and become better, faster, stronger. Your contribution takes us all closer to the ultimate goal of creating a general purpose artificial intelligence, which is open for everyone."

14 So, is it worth it?

There are lots of pros and cons to a setup like this. There's the freedom of being able to create your own Skills, or to find them in the brilliant online community.

Above Mycroft Mark I comes with speakers and mic built-in

Other versions of Mycroft

At the moment Mycroft is available in several flavours. The version we're looking at here, technically known as Picroft, consists of the free software only – you'll need to add your own Raspberry Pi to run it, plus speakers and a mic.

If you prefer an off-the-shelf version, you can opt for the Mycroft Mark I ($180), a standalone hardware device which is equally 'hackable' in terms of adding abilities or changing code. Finally, there's Mycroft for Linux, which you need to install using either a shell script or a standalone installer. Mycroft AI describes this as "strictly for geeks".

There are the benefits of being able to jump into the code and have a play-around, or just to check that your data really isn't going anywhere. But you have to balance that against the inability to find a working Skill when, for example, you just want to stream some music. When we did manage to get a stream working, Mycroft would talk all over the audio stream, with false positives of the wake word being picked up.

If you're a hobbyist looking for a new project to sink your teeth into, Mycroft might be right up your street. If you just want a device that you can say "Play the Beatles" to, without getting out of bed, this might not be the right setup for you right now. That's not to say it won't ever be, with the community and Skills growing at a rapid pace – and with the very promising Mark II version on the way, who's to say what Mycroft might be in a years' time? At the moment, though, it's lacking commercial viability. ■