MAKE | BUILD | HACK | CREATE

HackSpace

TECHNOLOGY IN YOUR HANDS

hsmag.CC | July 2024 | Issue #80

BEST 3D PRINTER 2024

Prusa vs Bambu

We put the XL, MK4, and Mini up against the X1C, A1 Mini, and P1P to find the ultimate modern replicator

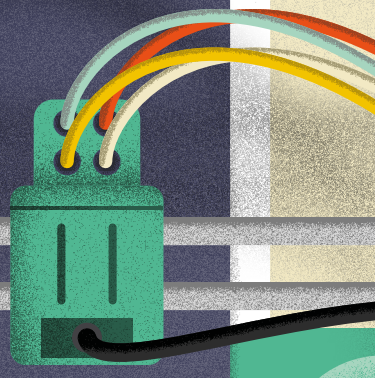

- 3D SCANNING WITH YOUR PHONE
- BUILD A TREASURE HUNT WITH RFID
- • MESHTASTIC • A DECENTRALISED RADIO NETWORK
- RASPBERRY PI AI KIT ON TEST

July. 2024
Issue #80 £6
Raspberry Pi PRESS
80
9 772515 514006

EMF CAMP NIGHT LIGHT **3D PEN** USEFUL SENSORS

The parts we sell help lives become richer

Imagine hearing aids that let a child experience her parents' voice clearly for the first time.

At DigiKey, the parts we sell help companies turn innovative, game-changing ideas into real-world solutions that change lives.

Find your part at digikey.co.uk

DigiKey

we get technical

Welcome to HackSpace magazine

I spend a lot of time – probably too much – thinking about what makes a particular piece of technology good or bad. There can be a temptation in the technical press to get obsessed over numbers. If the numbers go up, that must mean that something is fundamentally better. However, that's wrong. All technology is about tools, and a tool is better if it makes it easier to achieve your aim – this month, the tools in question are 3D printers.

It is, of course, possible that your sole aim is producing a Benchy in the shortest possible time, but it's far more likely that you have something else in mind when you get a 3D printer. Perhaps you want to expand your range of gaming miniatures, or you want to make enclosures for your project, or you want to learn how to design mechanisms.

Speed is obviously part of the equation, and so is quality. However, software, documentation, support are key. Hackability and upgradability are also important factors to consider, as you may find that you want to improve your printer over the coming months and years. This month, we'll be trying to judge 3D printers by these standards to see what comes out on top.

BEN EVERARD
Editor ben.everard@raspberrypi.com

Got a comment, question, or thought about HackSpace magazine?

get in touch at **hsmag.cc/hello**

GET IN TOUCH

hackspace@raspberrypi.com

hackspacemag

ONLINE

hsmag.cc

Available on the App Store

GET IT ON Google Play

EDITORIAL
Editor
Ben Everard
ben.everard@raspberrypi.com
Features Editor
Andrew Gregory
andrew.gregory@raspberrypi.com
Sub Editors
David Higgs, Nicola King

ADVERTISING
Charlotte Milligan
charlotte.milligan@raspberrypi.com
+44 (0)7725 368887

DESIGN
Head of Design
Jack Willis
Designers
Sara Parodi, Natalie Turner
Illustrator
Sam Alder
Photographer
Brian O Halloran

CONTRIBUTORS
Marc de Vinck, Jo Hinchliffe, Rob Miles, Nicola King, Andrew Lewis, Robert Hamilton-MacLaren

PUBLISHING
Publishing Director
Brian Jepson
brian.jepson@raspberrypi.com
Director of Communications
Liz Upton
CEO
Eben Upton

DISTRIBUTION
Seymour Distribution Ltd
2 East Poultry Ave,
London EC1A 9PT
+44 (0)207 429 4000

SUBSCRIPTIONS
Unit 6 The Enterprise Centre
Kelvin Lane, Manor Royal,
Crawley, West Sussex, RH10 9PE
+44 (0)1293 312193
hsmag.cc/subscribe
hackspace@subscriptionhelpline.co.uk

recycle

This magazine is printed on paper sourced from sustainable forests and the printer operates an environmental management system which has been assessed as conforming to ISO 14001.

HackSpace magazine is published by Raspberry Pi Ltd,
194 Cambridge Science Park, Milton Road, Cambridge, England, CB4 0AB.

ISSN: 2515-5148.

Contents

Cover Feature

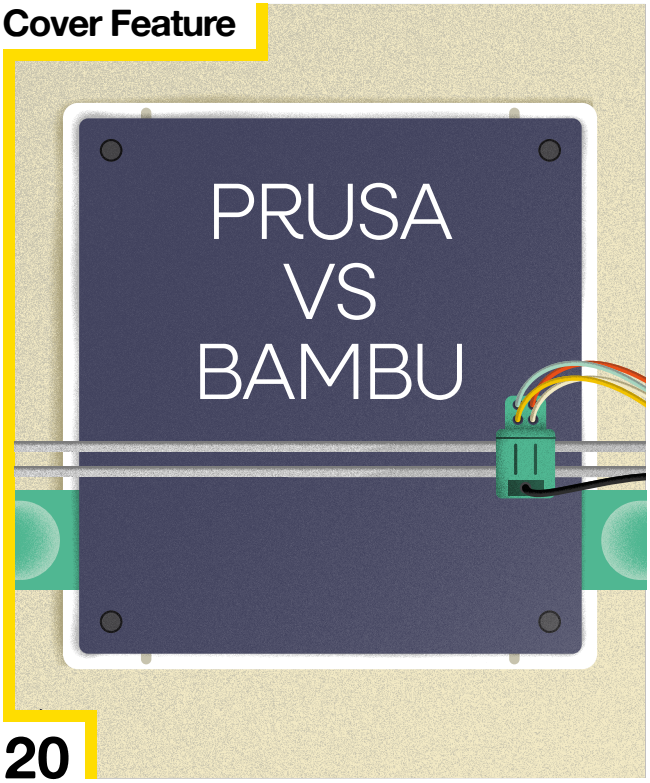

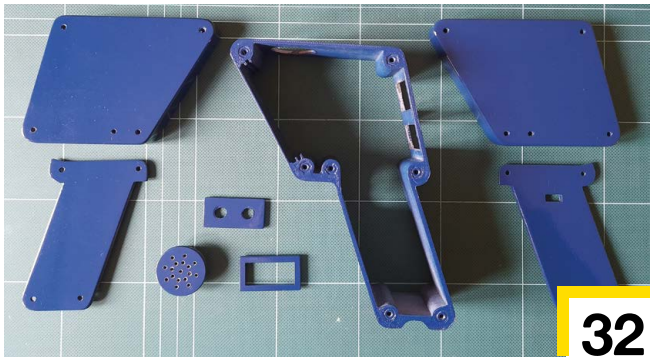

Tutorial

3D printing pen

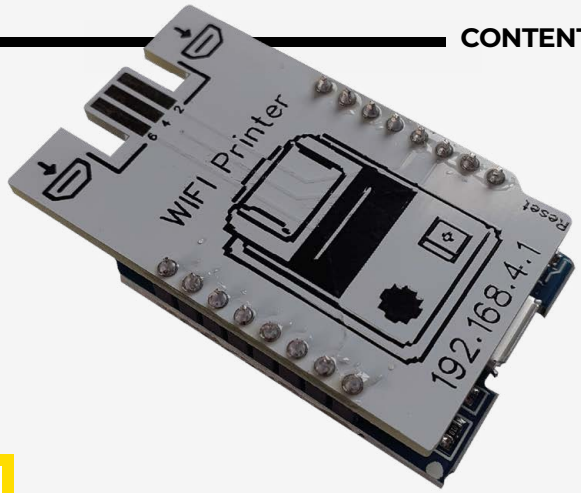

Interview
Pete Warden

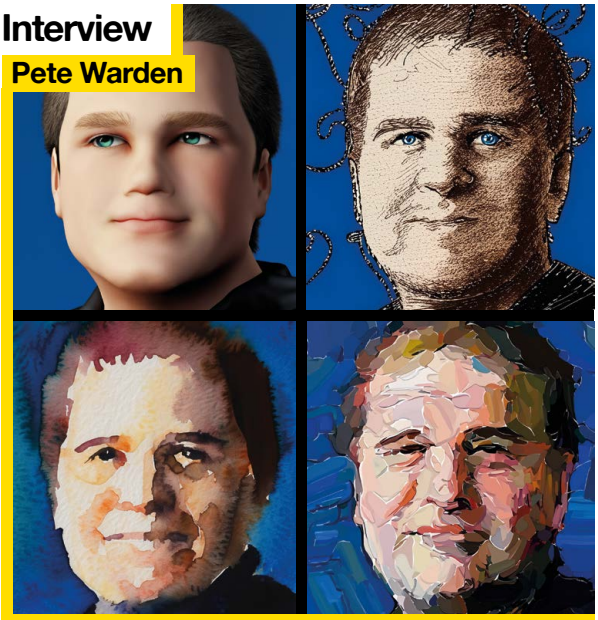

Review
Raspberry Pi AI Kit

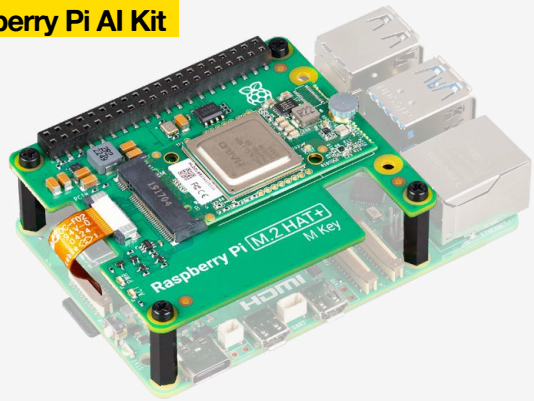

53 FORGE

83 FIELD TEST

REGULAR

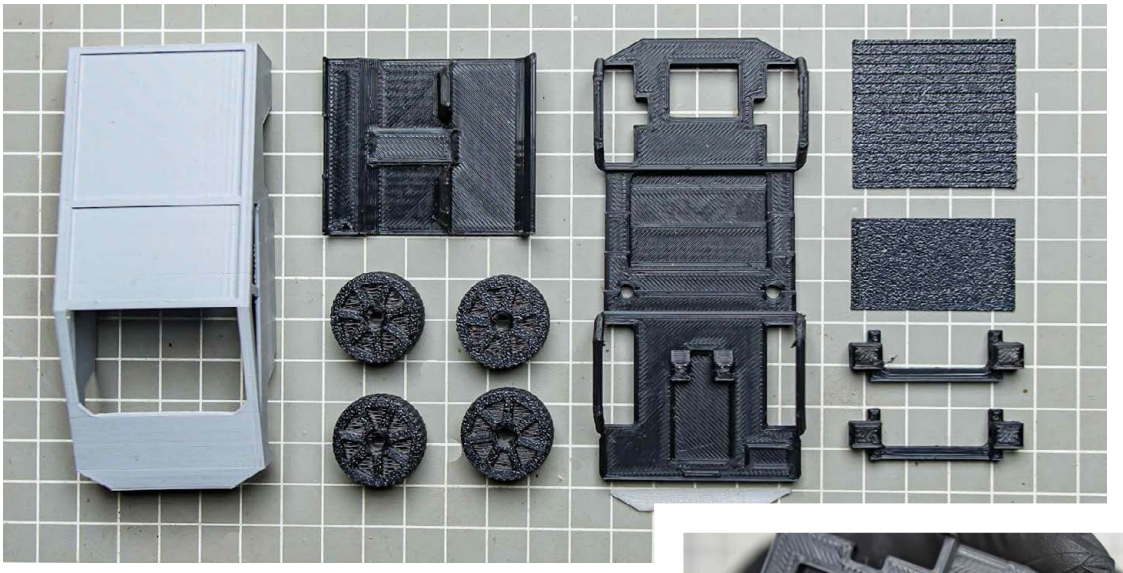

Cyber Clock

By **gokux** hsmag.cc/CybertruckClock

Ah, the Tesla Cyber Truck. Its angular, low-poly bodywork means that it's unlikely ever to pass the Euro NCAP (European New Car Assessment Programme) safety standards that govern what we can drive in the UK. Apparently Europeans have soft, fleshy bodies made out of meat, which deform unfortunately when wrapped around angular stainless steel objects build without crumple zones. What these faceless Brussels Eurocrats don't understand is that human life isn't that important compared with looking really cool. The maker of this project, gokux, reckoned the Cyber Truck's flat planes would make the perfect case for a desk clock, and by golly, they were right.

The Cyber Clock can not only tell you the time, but also the humidity and temperature thanks to its onboard DHT11 sensor, displayed on a 1.5-inch transparent OLED screen via a Speed Studio XIAO ESP32C3 .

This is such a polished build – the maker has had an idea, and executed it perfectly. Even the wheels go round on the 3D printed Cyber Truck's body. Add a 500mAh battery and USB-C charging, and you've got the perfect executive toy for the billionaire who has everything. ◻

Left
That OLED screen reminds us of a head-up display, as seen on jet fighters (but not cars — yet!)

CH32V003 FM Radio Receiver

By **Stefan Wagner** hsmag.cc/FMRadio

PCB design has come a long way since everything was green. This FM radio is pretty much just a PCB, with the controls and even the speaker mounted directly to the one printed circuit board. This device is designed not just to be simple, but to be affordable too: the components are all easily obtained, and include a 32-bit RISC-V microcontroller, a 128 x 32 pixel OLED display, and a Li-Ion battery charger. The 3W speaker isn't going to give you the best sound quality, but is just the job if you want to listen to the cricket commentary on Radio 4 (or even baseball, if that's what floats your boat).

This design is barely more than a single PCB — there is a 3D printed case, but to hide those electronics would be a crime. ◻

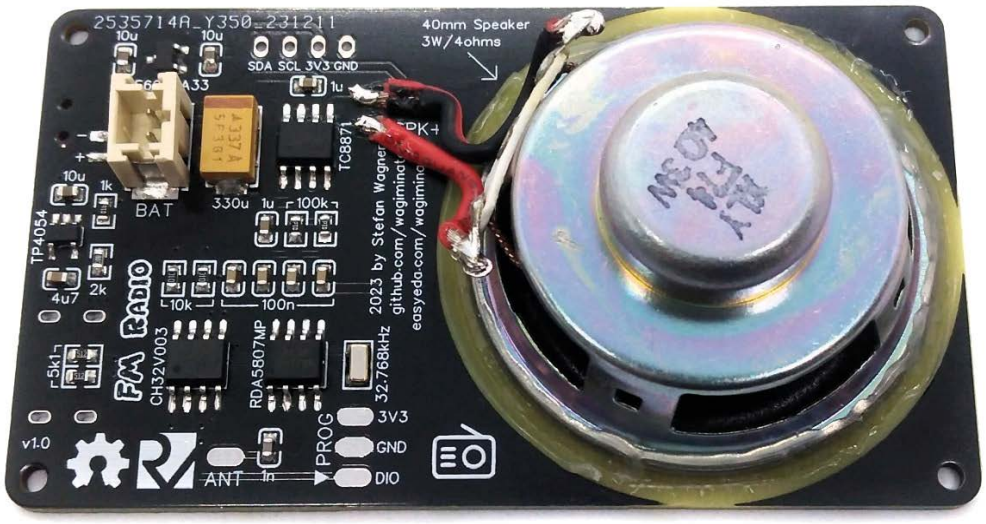

Right
The attention the detail is obvious from the back of the PCB, where the components are arranged to make this radio as easy to build as possible.


GND VDD SCK SDA
FM Radio
104.40
MHz
Volume:
ON OFF
CHARGE
VOL-
CH+
VOL+

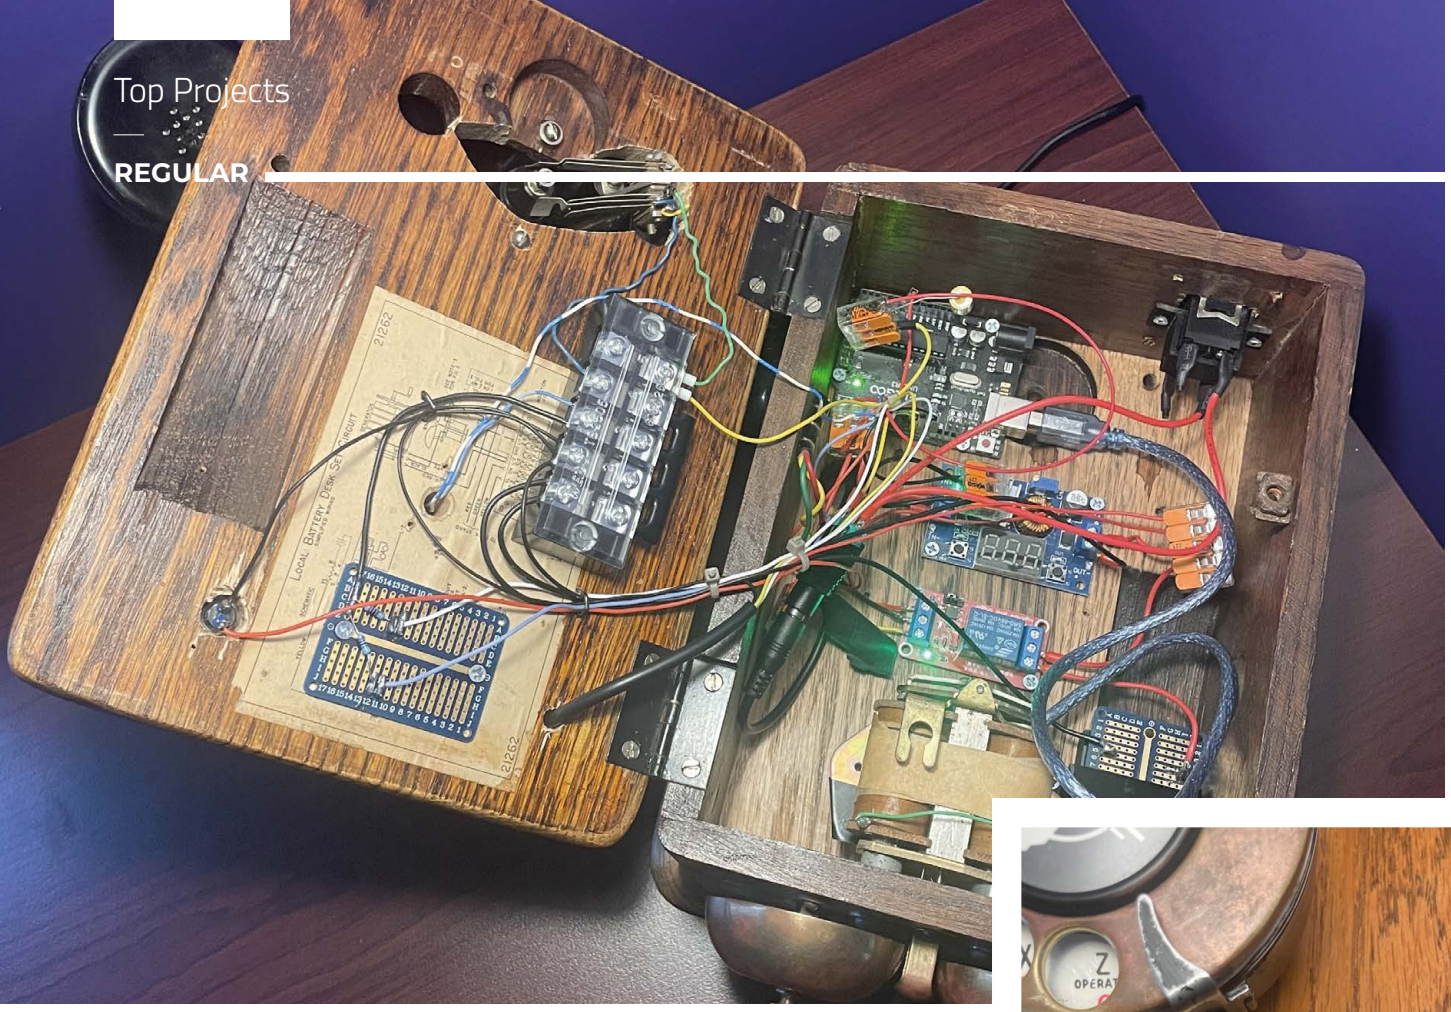

Rotary Phone timer

By **Scott-28**

hsmag.cc/RotaryPhoneTimer

Speak to anyone who's had a go at up cycling old electronics and they'll tell you that older is better. The devices were built to last in the days before planned obsolesce, so if you decide to reincorporate original components, there's a good chance that they'll still work. The cases are put together with screws, nuts and bolts, meaning you can get into the innards without destroying anything, as you would if they were glued together. And anything built before the 70, when microchips first made their way into consumer electronics, will necessarily be bigger in order to house the chunky old electronics components used at the time.

This old phone had already been upcycled by a previous owner before Scott-28 got his hands on it, converting a hand-crank phone into a rotary phone. Now Scott-28 has put an Arduino in it and turned it into a kitchen timer. He didn't want to mess too much with the vintage aesthetic, so opted for a red 7-segment display, a big, heavy power switch, red status lights and hand-stamped brass labels.

Above
Scott had to break out his oscilloscope to work out the pulses that come from the rotary dial so he could convert them into digital data useful for the Arduino.

Post Apocalyptic terminal

By **Rick Pannen** hsmag.cc/ApocalypticTerminal

We're not quite at the level where we can walk into the house and control our home appliances with voice commands – though thanks to Pete Warden and Useful Sensors, who we speak to this issue on page 38, that day is getting closer. Instead, the tech enthusiast has to come up with their own bespoke way to control their lights, music, heating and suchlike. One such device is this wall-mounted terminal by Rick Pannen.

We can almost hear your incredulity through the page: "Wall mounted! But that's a CRT screen. It'll weight a ton!". Although this looks it uses a heavy CRT screen, Rick's achieves a similar effect by using a 7-inch LCD screen in the back with a fresnel lens in front. It's a brilliant visual trick, as is the way he's dirtied down the 3D printed frame. First, some context. Own the days before ubiquitous CGI, film and television used models to depict space ships, space stations, and other non-real objects. To show the effects of being exposes to space dust, model makers would deliberately damage the surface of the models, making them look more realistic – a process known as dirtying down. Think of how lived-in the sets and models used in, say Red Dwarf or Alien look, and compare them with those of Flash Gordon, whose director expressly forbade anything to be dirtied down, instead opting for a brand-new, pristine look that translates onto the screen as looking fake.

Anyway, apart from the Raspberry Pi 400, the screen and the lens, the rest of this build is 3D printed in PETG on a Bambu A1. It was designed in Fusion 360 and took 75 hours of printing time and 2.2kg of filament. To achieve the weathered look Rick used a mixture of acrylic floor polish and solid from his garden, adding years of wear and tear in just a few minutes. ◻

Left
The case was printed in two halves to fit on Rick's 3D printer, then glued together using epoxy.

IV-12 tube calculator

By **Oskar 1527** hsmag.cc/VFDCalculator

If you love the look of Nixie tubes, but are scared by high voltages, VFD tubes could be the retro display device for you. They work using the same logic as nixies, with a connecting pin for each segment of the display plus three pins to provide power, and you can get hold of them at reasonable prices from online auction sites.

Maker Oskar wanted to build a project incorporating IV-12 VFD tubes as a display, and chose to build this beautiful calculator. It uses an Arduino Nano, Cherry MX Brown switches, key caps bought off Ali Express, and a custom PCB to allow the 1970s Soviet technology of the VFDs to work with an Arduino Nano.

Programming in Arduino's custom version of the C programming language can be scary for anyone used to working with the relatively more user-friendly Python language, but a calculator is a relatively simple concept that gives you an ideal reference point if you want to make the jump. And if you don't want to code, but you do like the look of a retro Soviet-looking desktop calculator, Oskar has made the code available on his Github page.

This project has been so beautifully realised that you may at first glance think the case has been up cycled from an existing build. It hasn't: the case is made from 3mm plywood, designed in the OnShape CAD software, and cut out by a laser cutting company, finished off with a walnut veneer and a few coats of linseed oil. Gorgeous.

Right
This beautiful calculator even emits a slight hum when it's turned on, which only adds to its already substantial charm.

REGULAR

Linamp

By **Rodmg** hsmag.cc/Linamp

For a long time our favourite word was "skeuomorphic" Examples of skeuomorphism include concrete given a texture to make it look like wood; or plastics made shiny and fitted with fake rivets to look like metal. It's a way of making one material look like another. You can see it all the time in the design of graphical interfaces for computers:" we've left 3.5 inch floppy disks behind a long time ago, but they linger on as the Save icon in computer operating systems to this day.

One great example of skeuomorphism is the Winamp music player, originally released in 1997, which was designed to look like the control panel of a Hi-Fi. Now the circle os complete, with software engineer and electronics tinkerer Rodmg bringing the aesthetic of Winamp into the physical world with this build: Linamp.

In the maker's own words, Winamp is 'a stand alone music player that aims to be easy to use, have great sound quality and look good." It can play music files, manage playlists, display a spectrum analyser and track information, control playback and more, with even more features than that on the way in coming versions.

To make this physical version of Winamp, Rodmg used 3D printing and sheet metal parts, designed by the maker and fabricated by PCBWay; he tore down an AIWA radio to take measurements of the metal parts, thin created his own using the Onshape online CAD software. The hardware includes a Raspberry Pi 4 and a 7.9-inch touchscreen. ◻

Left
The front panel was 3D printed in nylon by PCBWay.

HiPi.io

HIGHPI PRO

The new case from the HiPi.io team

- Rapid tool-free assembly and disassembly
- Large internal volume for HATs
- Compatible with Pi 2/3/4
- Multiple lid options
- Passive & Active Cooling options
- Secure microSD card cover
- VESA mount support
- Molding-configurable output ports customizable for volume orders
- Printed logo for your branding

Available at these great Pi stores:

PiShop.us PiShop.ca PiShop.mx ThePiHut Pi-Shop.ch Welectron.

Contact your favorite Pi store if it's not listed here

HackSpace
TECHNOLOGY IN YOUR HANDS

LENS

HACK | MAKE | BUILD | CREATE

Uncover the technology that's powering the future

FEATURE

BEST 3D PRINTER 2024

After years of lots of companies putting out almost identical 3D printers, it feels like the 3D printing community world is starting to centre around just two brands: Prusa and Bambu. Yes, there are still plenty more on the market, but if you're after a replicator in 2024, the chances are you're going to start with one of these two brands. We're going to look at three printers from each…

Left
When it was first released, the P1P was the first printer that reliably produced high-speed, high-quality prints easily

Above
The Prusa MK4 brought high-performance printing to the "bedslinger" layout

BAMBU

A1 mini (£269) – The cheapest printer on test, and it's often reduced from this price. As the name suggests, its print volume is smaller than most 3D printers, at 18×18×18cm.

P1P (£679) – the printer that made Bambu famous. It's a CoreXY machine that really brought this kinematics model to the mass market. There is also a P1S which adds an enclosure.

X1C (£1179) – A fully enclosed printer with higher-grade parts than the P1P.

And from Prusa, we're looking at:

PRUSA

MK4 (£791 kit / £1055 built) – the latest incarnation of the long line of bed-slinger printers that has its roots in the early days of RepRap and the hobbyist 3D printing movement.

Mini (£383 kit / £407 partially built) – An 18×18×18cm print volume machine that's recently been upgraded to include input shaping and, therefore, has had a big speed boost.

XL (from £1799 to £3791) – a big machine (print volume 36×36×36cm) with a tool changer and up to five tool heads. While only one tool head can be engaged at any one time, they can all be loaded with filament and ready to squirt.

There are some very distinct machines in here that have very different pay-offs for different purposes.

We're going to put them through their paces to find out which is really the fastest, which works best with multiple filaments, which is kindest to the maker community, and ultimately help you decide which (if any) you should get.

Our methodology is informed by years as 3D printing hobbyists, as well as the dozens of different machines that have come through our test lab over the years. We have attempted to focus on the things that matter to us in our printing rather than obsessing over artificial benchmarks. Turn the page to get started and find out how they stack up! →

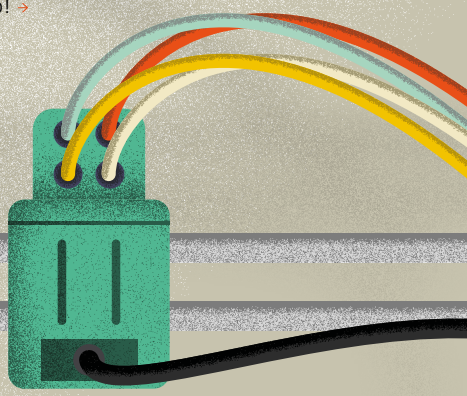

SPEED

NO MORE WAITING AROUND

Let's start in the traditional place with a baseline for how fast these printers can print a Benchy test. Take a look at the table and you'll see that the P1P and X1C are the fastest printers, capable of spitting out a tiny (and highly detailed) boat in just 39 minutes using their standard 0.2mm layer height profile.

We're going to be testing all the printers using their default profiles. Yes, there are probably some speed gains to be had by tweaking profiles, but then it's very hard to make it a fair test, and we suspect that few people these days print with anything other than the defaults.

The others range from a respectable 42 minutes for Prusa's MK4 to a relatively sluggish 51 minutes for the Prusa Mini.

3D printer	Benchy time (minutes) 0.2mm print time only
P1P	39
X1C	39
A1 mini	46
MK4	42
XL	48
Mini	51

Below
The fastest Benchy we produced in our testing was this one from the X1C

Let's take a look at another option: a lion skull. Scale this to 70% (so that it fits on the print bed of most of the printers), and we get a slightly different set of results. Namely, the Bambu X1C and P1P can print it in 13h 19m (using 394g of filament), while the Prusa MK4 does it in 16h 53m using 364g of filament. This shows an even bigger margin for the Bambu.

However, this isn't actually how we'd print this skull, and this is where the slicer makes all the difference. In both cases, we're using the recommended slicer for the hardware – PrusaSlicer for the Prusa printers and Bambu Studio for the Bambu printers.

In PrusaSlicer, we can use Automatic Support Painting to give us a minimal set of working

Below
(MK4) Every printer in this test could produce a good-looking Benchy using the default settings

supports. This reduces the time to 13h 28m (using 311g of filament). That's still slightly slower, but saves 20% of the filament.

If we enable lightning infill, this goes down to 10h 04m and uses 168g of filament. Unfortunately, neither of these options are available in Bambu Studio. We could reduce the infill percentage further, to 10%, and we get a speed of 12h 47m using 359g of filament. That's slower and uses over twice the filament of the Prusa MK4.

Furthermore, we could speed things up. Both printers have official support for up to 0.8mm nozzles. With these on and everything turned up as fast as we can make it, the Prusa MK4 comes in at 6h 47m, while the Bambu comes in at 12h 37m.

At the smaller scale, things play out similarly between the A1 mini and the Prusa Mini. While the Bambu

There's nothing to directly compare the XL to here, but it's worth noting that while it's slower in absolute print speed than the MK4 or the X1C, it can print larger things in a single go (or more small things in a single go), so if you need the space, it's massively faster.

Speed is a complex topic, and simply going by the speed of a Benchy is highly misleading for real-life use. The speed you'll get will depend a lot on the sort of models you're printing and the slicer settings you want to use. For the models we print, the Prusa MK4, when used with PrusaSlicer, is the fastest printer. You can use other slicers for your printer, but then you miss out on first part support and the finer points of network connection. →

Above
(X1C) 3D printed vases can look stunning, but aren't usually water tight (design by Mensandesign)

Left
(MK4) A few weeds can brighten up your table (design by Pachek)

MULTI-MATERIAL

MAKING PRINTS LOOK SPECTACULAR STRAIGHT OFF THE BED

If there's a hot topic in 3D printing at the moment, it's multifilament printing. This is an optional add-on to all the printers we're looking at, except the Prusa Mini. For all of them, except the Prusa XL, this means that there's an additional unit that feeds filament into the printer, and can swap this filament in and out of the hot end as you go through the print.

There are two big problems with this as a basic technology that affects both the Prusa and Bambu versions. Firstly, it's very slow. Secondly, it uses a lot of filament. Both of these things are to do with the need to print layers sequentially. If you use, for example, three colours on a layer, the printer needs to load, unload, and prime the nozzle three times. It depends on the model you're using, but typically, this means the printer takes far more time and uses far more filament in the process of swapping filament than it does in the rest of the print.

There are a number of smaller drawbacks – it's hard to print flexible filaments in this system, and it's hard to print different types of filament this way since the nozzle temperatures don't match. Neither of these are necessarily impossible, but don't expect it to be easy.

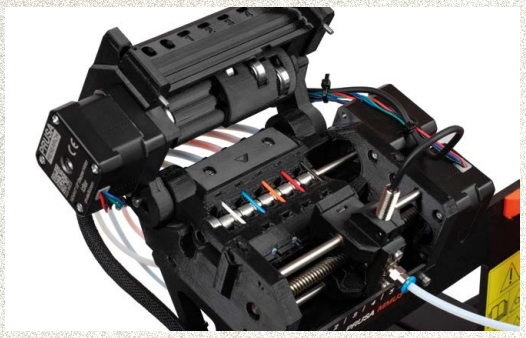

Right
Prusa's Multi Material Unit 3 takes up to five filaments and works with Prusa MK3 and MK4

Below
The Prusa XL's tool changer allows you to switch filaments far more efficiently than any other system

While they use the same basic technique, there are some differences.

Bambu's Automatic Material System (£309 for the AMS compatible with the P1P and X1C, or £229 for the AMS lite compatible with the A1 mini) lets you add four spools of filament. You can connect up to four AMSs to a single P1P or X1C printer, allowing for up to 16 colours, or a single AMS-lite to an A1, allowing up to four colours.

Prusa Multi Material Upgrade (MMU) allows you to connect up to five filaments to a Prusa MK4.

The different units have very different performances. For example, a two-colour Benchy takes 6h 47m and uses 85g of filament on an X1C, while the same model takes 3h 9m on a MK4 and uses just 42g of filament. Models with more colours make the difference even starker.

The significant difference between them is the 'flushed' filament – sometimes called 'poops' where the AMS deposits large quantities of filament out of a hole in the back of the printer. This is necessary because it cuts the filament on unload. On this model, 29g of filament is 'pooped' out.

The Prusa XL prints different filaments in a completely different way. It has up to five

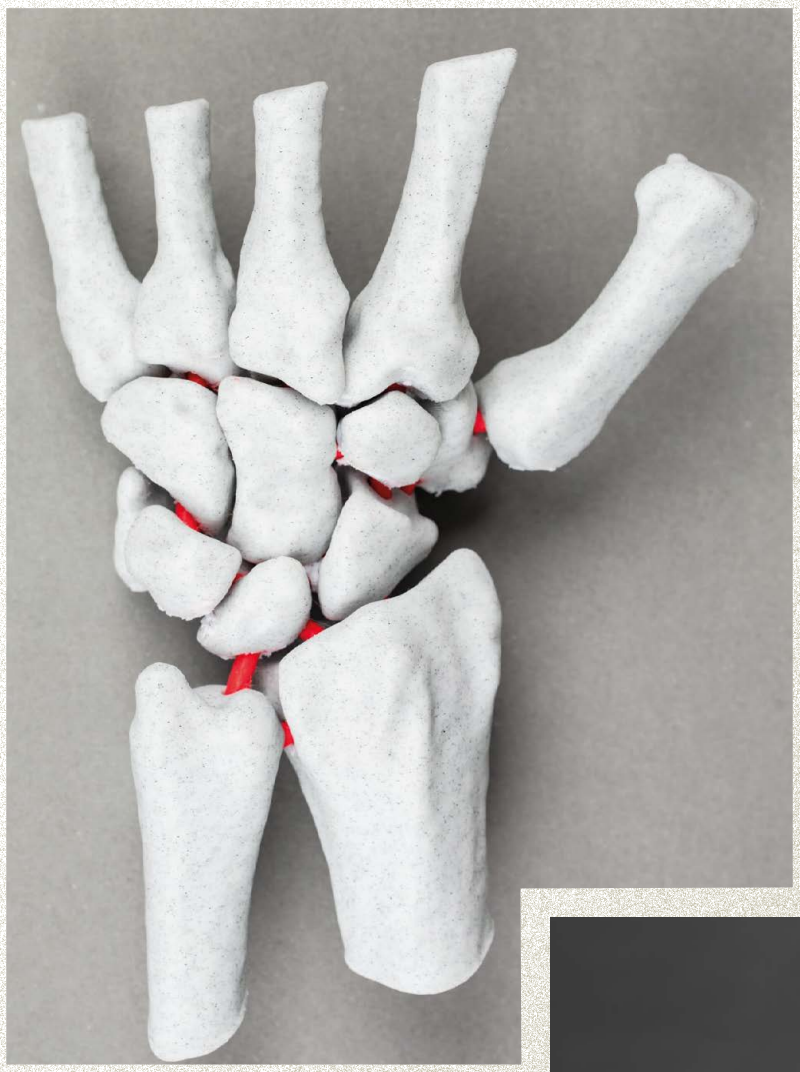

Left
(Prusa XL) This print combining flexible and ridged filaments to make a realistic model is particularly relevant to this review as the author is writing this while waiting for an MRI scan of the same portion of his wrist

Above
(Prusa XL) Multicolour prints can look excellent, but can waste a lot of filament

independent tool heads (though only one can be attached to the movement unit at a time). Each tool head can be loaded with different filament. This means that it can swap filament without loading and unloading a tool head. This means it's monumentally faster and uses significantly less filament. If your filament is fresh and very dry, you can print without a purge tower, which means that you use basically the same amount of filament as you do for a single-filament print. You can also print different types of plastic – this has a few uses, such as using dissimilar plastics for support and the main print so they don't bond together (PLA and PETG work well), or embedding flexible parts in an otherwise rigid frame.

The XL is the only printer on test that does multi-material prints well, in our opinion. It can combine plastics, doesn't waste huge amounts of filament, and it doesn't take forever. However, that obviously comes at a price.

The X1C and P1P get some points for being able to print with a massive 16 colours, but we do question whether this is something anyone wants. We deduct points for the colossal amount of wasted filament and time they have.

The A1 mini gets points for being the cheapest system on offer, but loses points – like the X1C and P1P – for the amount of waste.

The MK4 is the best of the single extruder multi-filament printers by a significant margin because of its speed and filament saving.

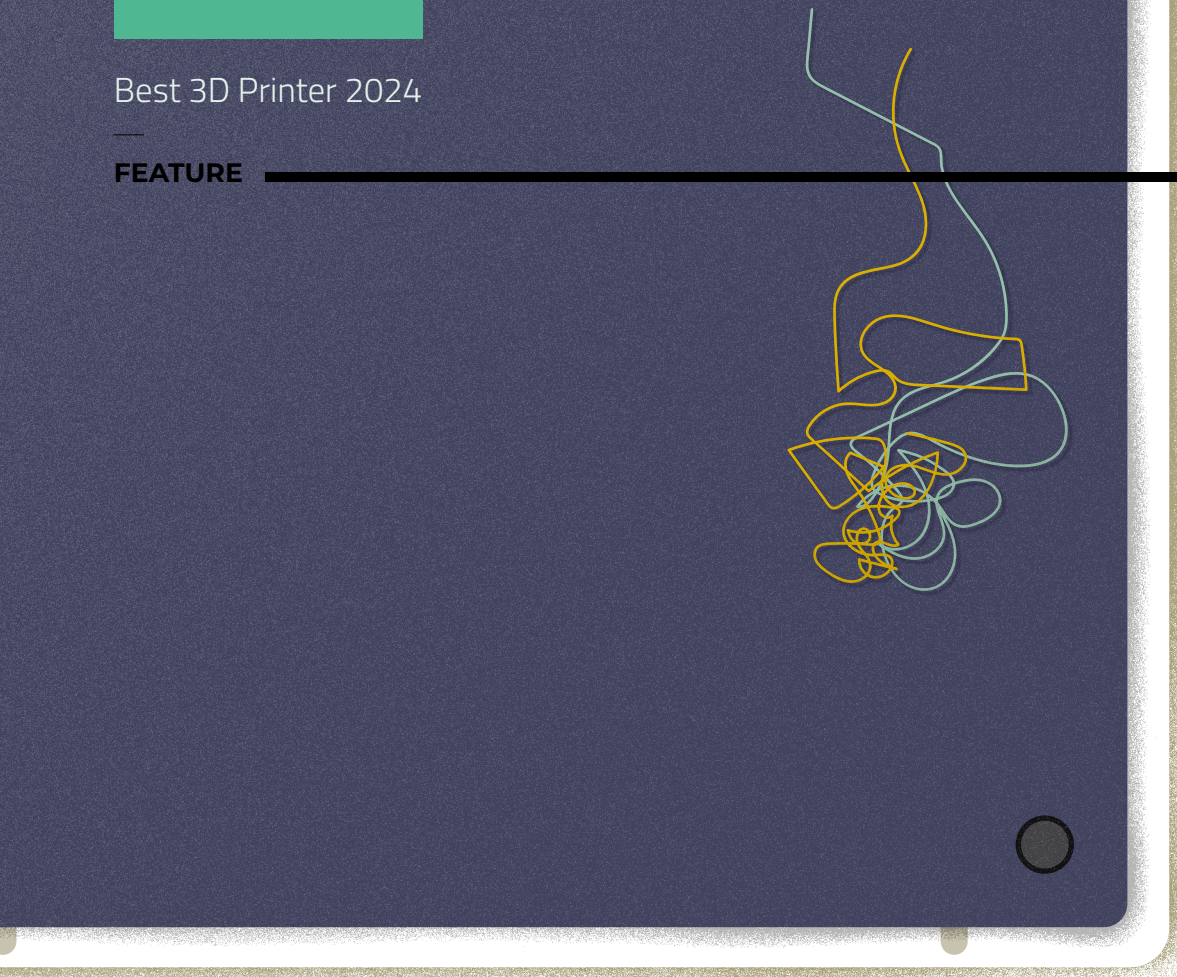

RUNNING COST AND WASTE

LET'S NOT THROW AWAY FILAMENT

The Bambu printers we tested used a lot of filament. Partly this is through their design – they cut the filament and 'poop' out the spare. It's also a software choice – they primed the nozzle with far more filament than any other printer we've used. Bambu Studio also has fewer features for minimising filament usage.

How big an impact this will have depends a lot on what you're doing. In some cases, we've found the Prusa printer using half the filament of Bambu. In some multi-material prints, the differences can be huge when compared to the MK4, and staggering when compared to the XL.

At £20 per roll, you're effectively paying an additional £20 for every roll of filament on the Bambu printers – that means that for every £100 a of filament used in a Prusa printer, a Bambu printer costs up to an additional £100. Admitedly, this is the worst case for single-colour prints, but the price between a Bambu P1P and a Prusa MK4 is made up for in as few as 25 rolls of filament. However, this is highly dependent on the actual models you're printing. Broadly speaking, the more structural filament your model needs, the less you'll save. If your model needs a significant amount of infill and doesn't need many supports, you might find that the difference is minor.

When it comes to multi-material (or multi-colour) prints, the differences can be absolutely staggering. The Prusa MMU is more efficient than Bambu's AMS in terms of amount of filament used. However, it's the Prusa XL that absolutely excels here, and with dry filament and a bit of tweaking, you can print in multi colour with no more filament than the same print in a single colour (though you may well get a better result if you do use a purge tower which will use additional filament). Although it's a significantly more expensive printer, if you plan on printing a lot of multicolour models, you may find that you end up saving money with the XL compared to any of the other printers on test.

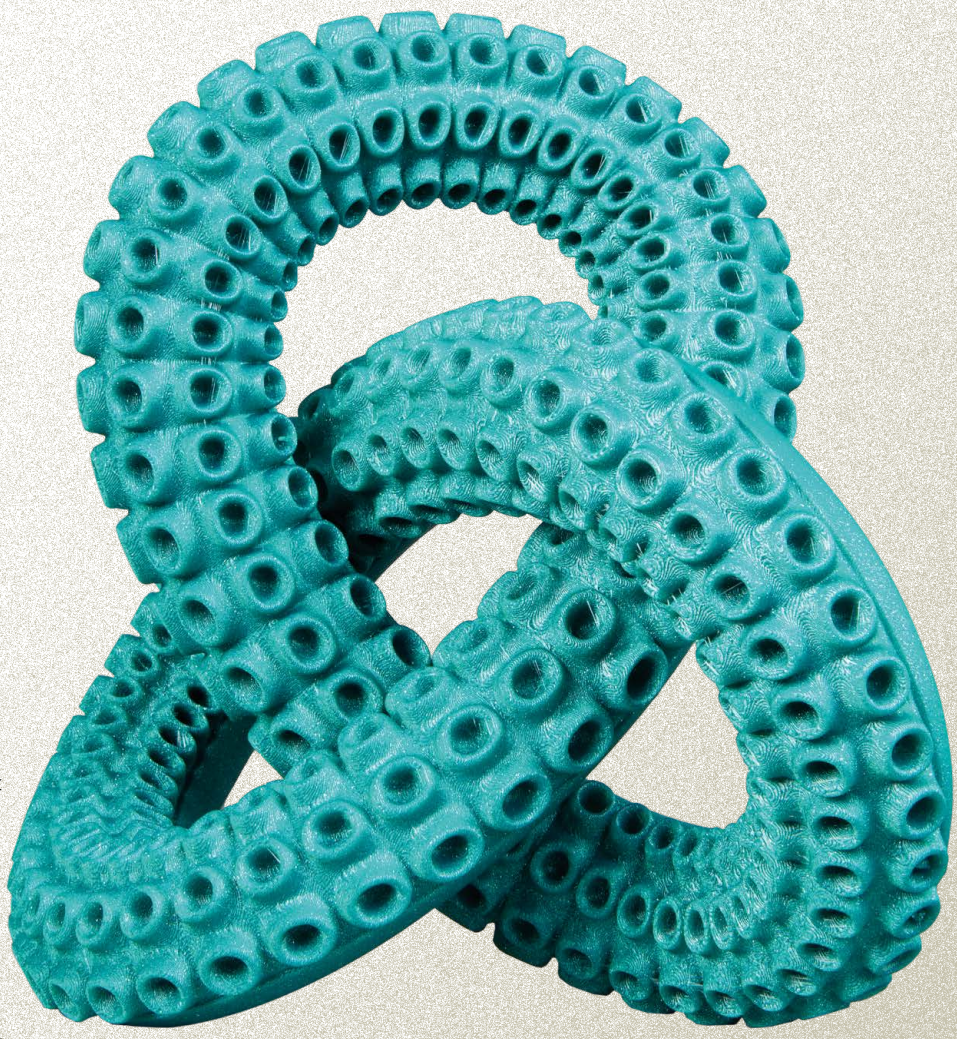

Above
(X1C) Nightmare fuel or fun trinket? You decide

SAVING THE PLANET

You just can't claim, in 2024, ignorance over the damage that plastic does to the environment. Irresponsible plastic usage is one of the biggest causes of damage to the planet, and it's hard to ensure that the entire process (from creation to ultimate destruction) is handled with care. And despite the fact that PLA is a bioplastic, we do not think it's reasonable to call it biodegradable. It's simply impossible for the majority of people to compost PLA. And any PLA that makes it into the environment isn't likely to break down into something safe in a meaningful time frame (though it may break down into microplastics).

It is possible, though not particularly easy, to recycle PLA. Few commercial places accept it, though it is possible to process it at home and we've looked at some of the options in HackSpace magazine.

What does it mean to be someone interested in 3D printing, but also be environmentally responsible:

- Minimise filament plastic. This means using as few supports and as little infill as possible. It also means reducing other plastic waste – filament that comes on cardboard or reusable spools saves a significant amount of plastic.
- Use recycled filament where possible. Yes, it can be a bit more expensive, and the colours aren't always as vibrant, but we have finite resources.
- Ensure you have the most efficient waste filament processing. Unmarked PLA is a bit tricky to dispose of, so it's worth taking time to find out where you should be putting it. We've heard of some people putting it in food waste because they believe it's compostable. DO NOT DO THIS UNLESS YOU HAVE EXPLICIT APPROVAL FROM YOUR WASTE PROCESSOR. Not only is PLA not compostable in many industrial composters, but the waste processor may reject not just your food waste, but all that it has been collected with as it's now contaminated. It might be possible to put it in mixed plastic waste, but it's highly unlikely to be recycled (as it's not marked), so whether or not this is a good place for it depends on what the plastic processor does with rejected plastic. At our workshop, they incinerate discarded plastic to produce power to run the recycling plant. In other places, they landfill discarded plastic.
- Use bioplastics where possible. There are cases where PLA is simply not up to the job, but for most users, those cases are few and far between. We'd encourage people to stick with PLA unless necessary. Fortunately, it's cheaper and available in more colours than any other plastic as well.

In our testing, Prusa printers are better for the planet and are cheaper to run than Bambu, though to what degree this is true depends significantly on the models you're printing. Both Bambu and Prusa use refillable spools, though only Prusa sells recycled filament.

OPEN SOURCE

GIVING BACK TO THE MAKER COMMUNITY

The history of hobby 3D printing is one of open source. It's arguably the big success story of open-source hardware. Prusa (both the man himself and the company he created with his name) has been at the centre of this for almost as long as there's been a hobbyist 3D printing scene. The modern 3D printing landscape would look radically different if open source and Prusa hadn't been part of it.

Prusa has an overview of which of its products are open source (almost all of them, including the MK4 and Prusa XL, though there was a delay in releasing this), and links to all the relevant sources are at **hsmag.cc/PrusaOS**.

Bambu has a very different history. When it first launched its printers, it was tied to a proprietary cloud service. After user backlash, it relented. Bambu has an encrypted NFC system on its filament spools that allows only Bambu filament to be automatically detected (though non-Bambu filament can be entered manually). Bambu Studio is open source, but it was legally obliged to be as it is built on the open-source PrusaSlicer (which, in turn, is built on the open-source Slic3r). None of Bambu's other software is open source (such as the printer firmware), and none of the printer designs are open source.

However, this isn't the full story on intellectual property. Open source works within the framework of copyright; there's also patents to consider.

Bambu is amassing a war chest of patents that take entire concepts away from the rest of the community (whether or not they are actually ever implemented by Bambu). 3D printing was invented in the 1980s, but hobbyist 3D printers didn't come into existence until the early 2000s because the entire field was littered with patents preventing anyone, but a few companies, from making 3D printers, and these companies only wanted to make high-end industrial machines.

Hobbyist 3D printing came out of the maker movement, and if the maker movement is to continue to help it develop and flourish, we have to avoid companies that try to close the hobby away with restrictive intellectual property practices.

Obviously different people will have different feelings about how important open source is to them. It can be hard to feel the benefit of open source to you as a user, and part of this is because the benefits come years down the line. Open source hardware is more likely to be upgradeable and repairable in the future, and it helps not this generation of printers, but the next.

Right (various printers) These stackable cats are a 3D printed game. Take turns to stack a cat and see who knocks down the tower (design by Valeria Momo)

Left
(A1 Mini) Voroni patterns are popular in 3D models and are hard to produce in any other manufacturing technique (design by Marcel Baumgartner)

THE BEST 3D PRINTER

So, having gone through all those categories, which machine are we going to pick as the best 3D printer available to buy at the moment? The Prusa XL.

There are faster printers, there are cheaper printers, but there's no other printer pushing the field of 3D printing forward in the same way as this printer. It can print things no other printer on the market can. It can print multicolour prints at a fraction of the time and cost of any other printer out there. Of course, it's not the right printer for everyone.

You'll have to decide for yourself how to balance the different competing factors like speed, cost, software, and open source-ness for your particular needs. Hopefully we've given you the information you need to make up your own mind about what's right for you.

SUBSCRIBE TODAY

GET SIX ISSUES FOR JUST:

£30 UK / **€43** EU / **$43** USA & Canada

SUBSCRIBER BENEFITS:

- **Get every issue of HackSpace magazine delivered to your door**
- Learn from hackers and makers in our in-depth tutorials
- Early access to the PDF edition
- **Get a free Raspberry Pi Pico W**

hsmag.cc/subscribe

HOW I MADE AN RFID TREASURE-HUNT

By **ROBERT HAMILTON-MACLAREN**

Some time ago, I had an idea to create a treasure-hunt using RFID (radio-frequency identification) cards and an MP3 player module. Answers would be hidden until the cards were scanned and the order of reading could be controlled. While the concept seemed achievable, an early problem was how to create the clues as MP3 files. I experimented with recording myself and tried using Amazon Polly, an online voice synthesis service, but neither was quick or straightforward. I parked the idea until I read *Make interactive greetings cards with NFC tags* in HackSpace magazine, issue 67. My interest in this project was revived. What if I could store the clue on the card as text and use the microcontroller to read using voice synthesis? I stumbled on the Gravity DFRobot Speech Synthesis module (**hsmag.cc/speech-synth**) which seemed suitable for the project, so I ordered one.

Left ←
RFID treasure-hunt reader

Below ↓
Card writer app, created in Visual Studio

The first task was to store clues onto the cards. The original plan was to use a mobile phone as the method for saving the clues onto the cards, however, this turned out to be a non-starter. Although an NFC phone app made it easy to write text to a card, the chosen RC522 card reader module reported the card as corrupted. Undeterred, I set about understanding how the data is stored on the cards. A card is split into blocks of 16 bytes, and every fourth is reserved as a checksum. This made writing a long sentence more complex, so the solution was nested loops. The text is split into blocks of 16 characters ready to write into the first three blocks, skipping the fourth. I decided to allow 144 characters, which meant nine iterations were required.

With the routine to write data sorted, the Serial Monitor in the Arduino IDE would have been the quickest solution to receive user input. However, to make the process more accessible, I wrote a program in Visual Studio to provide the interface for storing the sentences to the cards. This turned out to be quite complex as two-way communication was required between the Arduino and the app. It was worth the effort as the end result makes it much easier to read, edit, and save the data to the cards. Steps were added to the procedure to store a card number (for sequencing) in another block. →

"I EXPERIMENTED WITH RECORDING MYSELF AND TRIED USING AMAZON POLLY"

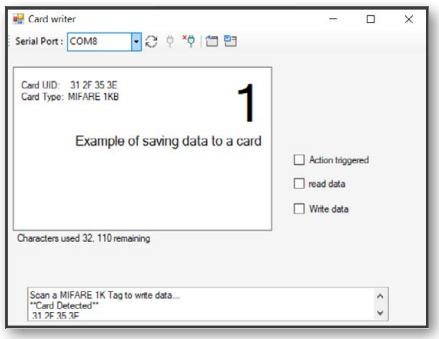

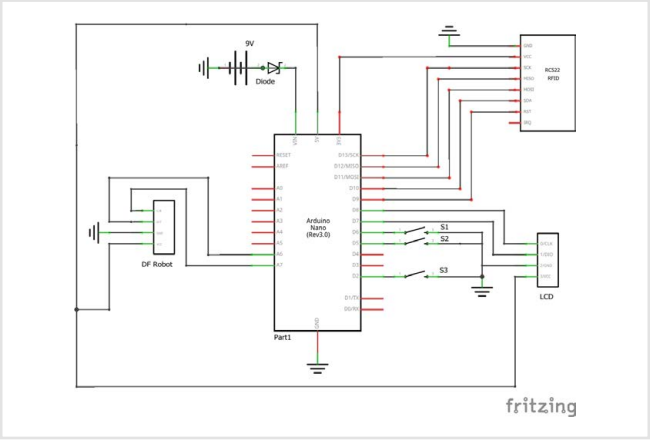

MAD

Keeping the card configuration and the trail reader as separate programs simplified the code and avoided designing a device to do two things. As the code grew, space also became a consideration, with total program memory limited on an Arduino Nano. I started by assembling circuits on a breadboard. To begin with, the wiring was straightforward, being just the Speech Synthesis module and an RC522 card reader.

The DFRobot module links via I2C to receive phrases. There is not much documentation, but the sample code is fairly straightforward to modify. It's worth noting that, unlike the MP3 module, code does pause while speaking (this was beneficial to my application as I didn't want the code to continue until the message had been played).

READER COMPONENTS

- Arduino Nano
- DFRobot Speech Synthesis module
- JST connectors
- RC522 card reader and cards
- 2 × 3 AA battery holder
- Slide switch
- Diode
- Four-digit, seven-segment display
- 3 × 'push-to-make' buttons
- Electrical wire
- Stripboard
- Suitable enclosure

When testing, an intermittent fault occurred whenever the reader was moved away from the card too early. Although there were several instances of the fault reported online, none of the various fixes worked. The solution that finally worked was to reset the buffer size in the code, which lost its value when a signal was interrupted.

Some words don't work well with the synthesis, but you can play around with the spelling to get a better result. One mispronunciation for a pirate treasure-hunt I could not fathom was that 'smuggler' seemed to sound like 'sunflower'!

With the cards able to be read, it was time to test the concept. My daughter was keen to give it a go, so we prepared our first treasure-hunt. As the prototype was on a breadboard and connected to the computer, it was necessary to bring the cards back to the reader each time.

During testing, it became apparent that speech wasn't always sufficient to tell us what was going on, so a four-digit, seven-segment display was added. Although four characters is limited, it was enough to display words like 'SCAn', 'rEAd', and 'PLAy'. In general, the module was easy to use, but when displaying text, repeated characters would only be displayed once. I accepted this when displaying words, but it was more of a problem when displaying time, as the colon was in the wrong place. A solution was found in a split-screen example, allowing minutes and seconds to be displayed separately. I don't know if the

Left
Reader schematic

Right
3D-printed parts, before sanding and painting

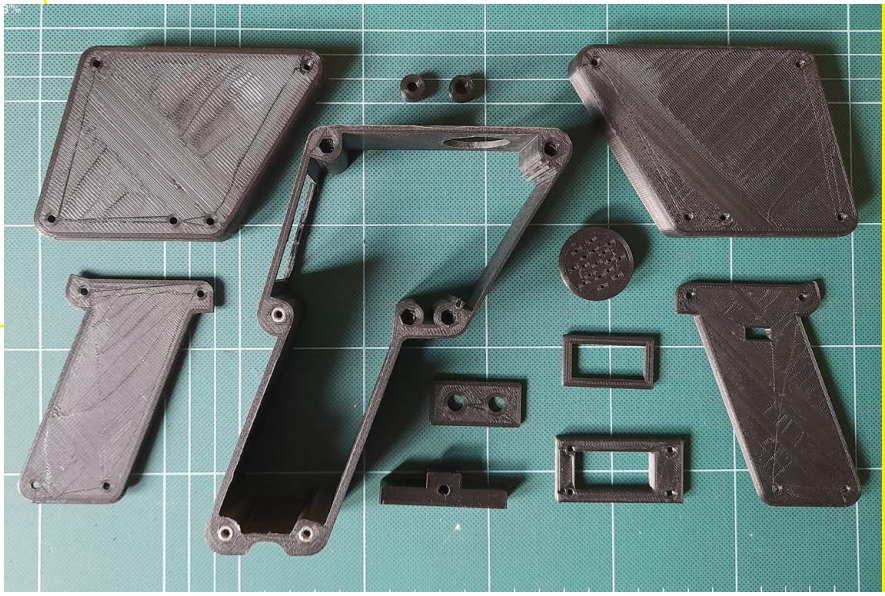

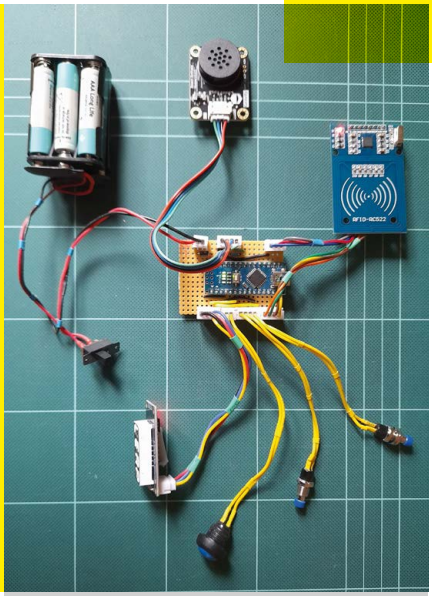

Above
Final circuit ready for enclosure

Below
3D parts after the topcoat of paint

issue was with the library or the module itself. There are a number of alternative libraries that could have been tested.

In addition to the reader, sound module, and display, the design consists of three buttons to scan and navigate around a menu. Transferring the circuit to a stripboard was easy as it only contained the Arduino, connectors for the modules, and a diode for battery protection. Using JST connectors meant fitting the parts did not result in a tangle of cables.

Flexible silicone wires make fitting leads into the case easy. However, this prevents wires from being twisted together into a group. There wasn't space in this build for cable ties, so I came up with the idea of using short heat shrink bands as small collars to keep the wires grouped together.

After the code had been written and the circuit built, it was time to create an enclosure. The modules made it quick to build a prototype but took up more space than an integrated circuit would. The 3D-printed case was a bespoke CAD design. All of the parts were created so that they could be oriented requiring minimum support. The case was fitted using a method that I had trialled before. Recesses for M3 nuts were incorporated to accept screws which pass through other components and hold them together. Most nuts will push fit, but some require gluing. As the heaviest components to place were the batteries, these were put in the handle – this achieved a nice balance. Each module required its own mounting method. These were printed separately, making it easier to refine positioning and avoid reprinting the whole case.

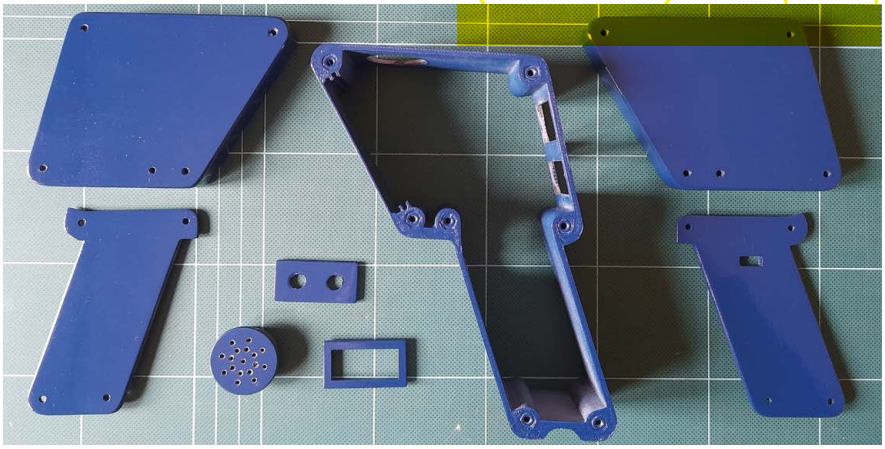

"THE PROCESS WAS TIME-CONSUMING, BUT A GOOD RESULT WAS ACHIEVED"

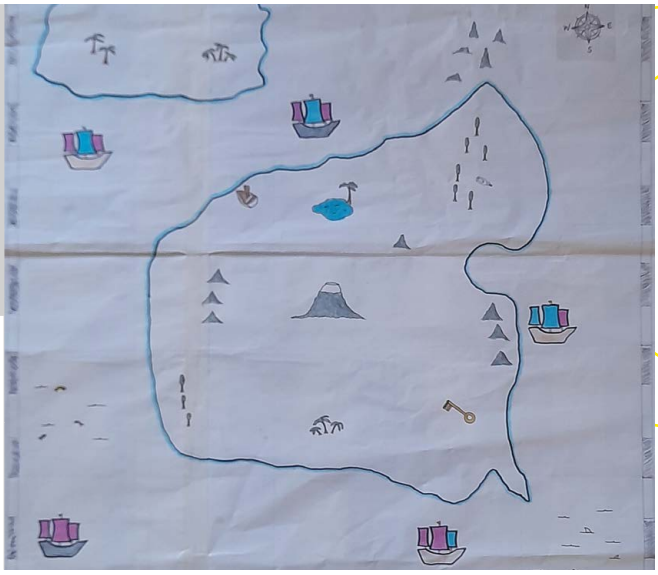

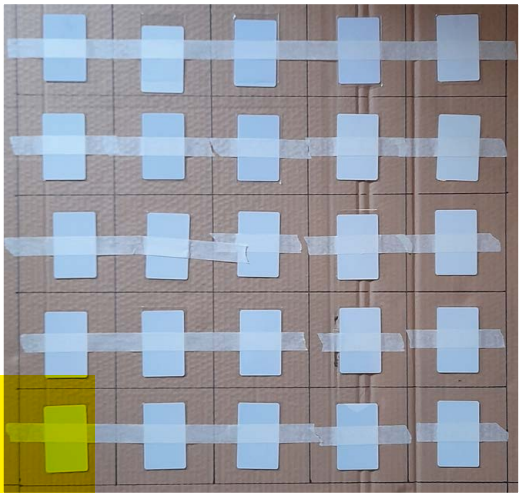

What started as a trial print for the main part of the case, with a little filing, ended up being the finished article. Having selected print settings for speed over quality, the finish was not great. I decided to paint the components rather than reprint them. This was my first attempt at painting a 3D print, so I watched some tutorials to determine an approach. The chosen process involved repeated sanding with increasingly fine-grade paper, priming and filling to create a surface ready for a topcoat. Using a spray paint that dried quickly allowed layers to be slowly built up to avoid runs. The process was time-consuming, but a good result was achieved.

Once the paint had hardened, the module mounts were glued in place. The same brand of primer and topcoat paints were used to ensure compatibility, but I missed checking the glue used to fix the bezels in place. Although minor, the glue reacted to the paint which meant a couple of spots needed to be sanded and sprayed again. With the mounts in place, the modules were easily fitted. As the main module was in the centre of the case fixed at one end, I covered the back of the stripboard with insulation tape as an extra precaution should it come loose from its mount.

Now, the enclosure had been assembled and the code had been further refined. I created two games, the first being the original concept, which was a treasure-hunt. In its basic form, you follow clues and riddles until you get to the end. Code prevents cards from being skipped, which means the hunt can't be done in the wrong order. To make the hunt more interactive, it's possible to build in some challenges (scan the wrong tag and points are lost; scan the right one to reveal the next clue). The code also looks for special tags in the card sentences, `{T3}` adding a three-minute timer after the message, and `{E}` signifying the end of the game. One variation tested was a 5×5 grid of cards hidden behind a treasure map (**Figure 1**), combined with breakout mini puzzles, which worked well for a smaller space. **Figure 2** shows a mini puzzle where a fish matching a

Figure 1
The treasure-hunt doesn't need to spread over a large area – cards are positioned in a grid under the map

Below
All parts fitted into the case, with batteries in the handle

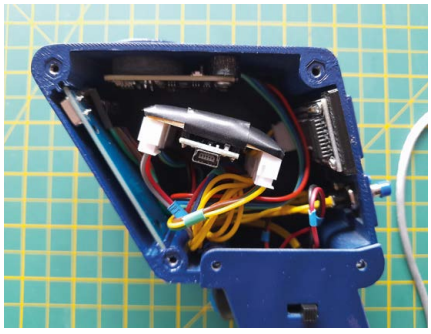

OTHER APPLICATIONS

Time trial – See how quickly you can complete a trial between a series of waypoints – no clues required, just numbered cards. This would require a minor modification to the 'Point to Point' game code.

Memory chain – A chain of numbers is called out, and the player scans the cards to repeat the sequence. Each time, the chain increases. The aim is to see how long a chain can be remembered without making a mistake.

Maths grid – With a grid of numbers, for example, 1 to 20, maths questions would be called out and the player has to scan the right answer. Different overlays could contain different sets of numbers.

Tour guide – Applications are not limited to games. Adding tags to objects could allow facts to be discovered while an attendee explores a museum.

Figure 2
One of the mini challenges: catch the matching fish within the time limit

Figure 3
'Point to Point': race against the clock to scan as many numbers in two minutes

specific description must be found, caught, and scanned within a time limit. Equally, a treasure-hunt could be over a much larger outside area. Murder mysteries and escape room puzzles are all possible with the creative use of cards and sequences. There is space on the cards to hold multiple sentences; other games might make use of random messages, or returning to a card a second time may result in a different message.

The second game, 'Point to Point', is much simpler and requires minimal setup (**Figure 3**). A series of numbered waypoints – one to five – are set out in an open space. The player starts the game by scanning card one, then another number is randomly called; the player runs between the numbers with the aim of seeing how many points can be collected in two minutes. It is possible to make the game harder by increasing the distance between the numbers or by adding obstacles.

This, of course, is the first iteration. Many improvements could be made to the design. The four-digit, seven-segment display is very limited – an OLED screen would provide much greater flexibility.

Opting for the Speech Synthesis module provides something different, but a better screen which could display the card text is an alternative, as long as reading isn't a barrier for users. The DFRobot module selected is only capable of speech – others also produce sound effects. This would offer more flexibility as to how the device responds to actions; however, these appear to have less documentation and examples so may be more challenging to use. The six AAA batteries take up quite a lot of space, meaning the handle grip is on the large size for smaller hands. An alternative battery would allow us to resize the case.

I've enjoyed building this project. It is fun to use and it challenged my creativity, requiring regular problem-solving to overcome various obstacles along the way. When a solution was not obvious, pausing the project allowed time to consider a resolution. As with previous projects, I have been able to further develop skills including programming, electronics, and 3D printing, and also new skills such as painting. The end result is a configurable RFID reader with uses only limited by imagination and time. □

HackSpace magazine meets...

Pete Warden, Useful Sensors

AI for the masses is coming to a tiny, cheap sensor near you

Pete Warden used to be the technical lead of the mobile and embedded TensorFlow group, in Google's Brain Team. You might have heard of TensorFlow – it's a machine learning platform, and it's open source, which means anyone with the right hardware can download it and get to work on building it into their own gadgets. And the right hardware is now cheaper than ever.

Pete's been banging the drum for embedded AI for years now, arguing that the right place for your data is on a device that you own, rather than a data centre somewhere. And, with his company Useful Sensors, he's ready to bring AI – secure, easy to use, and yes, useful – into the home. →

Above
We've used AI to make Pete Warden look more like a Pixar, oil painting, watercolour, line drawing kind of guy. His own flavour of artificial intelligence is a lot more useful

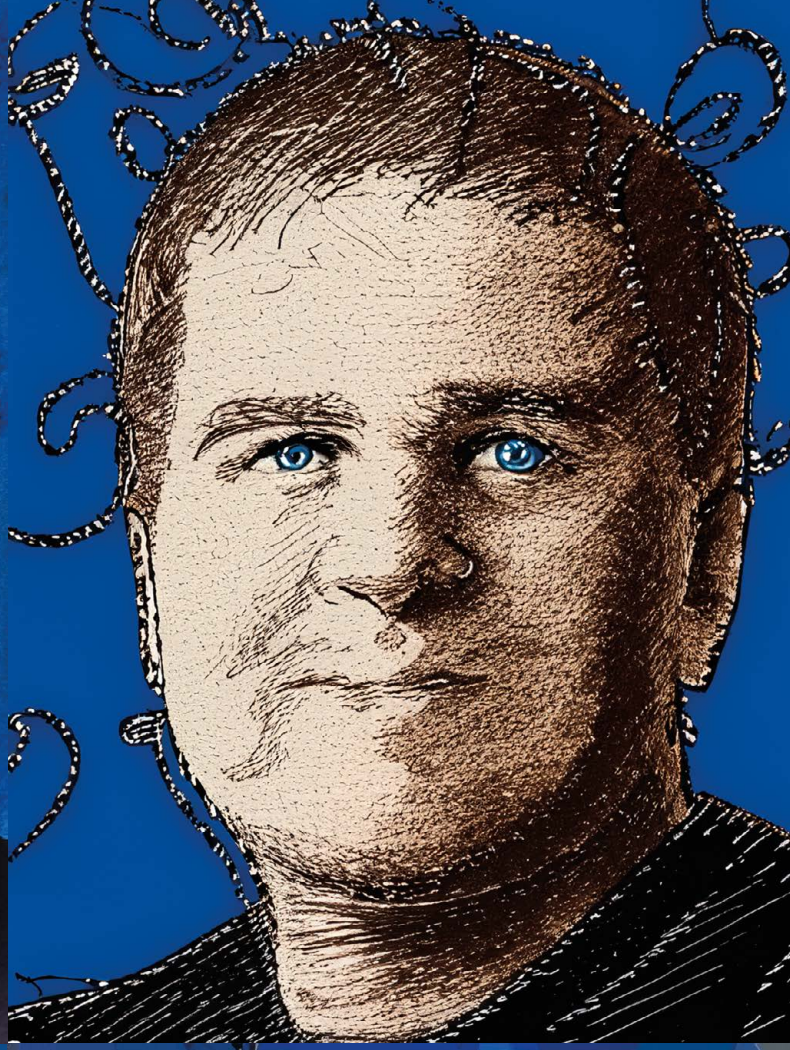

HackSpace Introduce us to Useful Sensors; what do you do?

Pete Warden We make everyday objects as easy to interact with as another person. That's our goal. That means having your TV pause when you get up to make a cup of tea without having to find the remote, for example. I would love to be able to look at a lamp, say "on", and have the lamp come on.

And be able to have voice interfaces that are really intuitive and easy to use. And private, where none of your audio data ever leaves the device. Everything stays local. That's the high-level vision of what we're moving towards. I love the idea of a Disney house, where you walk into it, everything's talking to you. You're talking to everything.

And obviously, that vision can be super-creepy, so we're trying to actually build things in a way where they're not only secure, and private, but where people can actually confirm that security and privacy.

HS AI still feels really new to me, so I guess the phrase 'traditional AI' is kind of meaningless. But with traditional AI, I guess information gets sent back to the cloud or to a big megacorporation somewhere.

PW Exactly. It goes into somebody else's data centre.

When I worked at Google, on anything to do with speech, friends and family would always ask me something like: "I was just talking with some friends about – say – Morocco. And then the next day, I got all these ads for holidays in Morocco. Is Google spying on me?" And I know the code and I was able to say hand on my heart, no. I've worked on that code. We just, we just don't, we don't have the capability. We don't do that.

But the thing is that I can't prove that – I can't have a third party verify that claim. Because everything's happening internal to Google's data centre. And it's extremely difficult to audit the flow of information. Once something goes behind a corporate firewall, it's very hard to keep track of all the ways that could be copied or moved around the system.

I really wanted to build systems where you could break them up into little modules that do a particular job. With the person sensor, for example, nobody can access the camera that's in there; nobody can access the frame data – all they can get is the metadata about, hey, there's a person here, they're looking at the camera, and they're standing, or sitting, or whatever.

We've actually commissioned a third party to verify what we're doing, and verify our claims. And because it's a small enough system, it's actually possible to do that.

> We've actually commissioned a third party to verify what we're doing, and verify our claims

HS I'm so disappointed that you said that about Google not listening to everything we say, because I'm certain that that's happened to me.

PW People don't like to hear this, but you're very predictable. As far as I can tell, it's that there's usually a bunch of other signals – maybe somebody on your same IP address did some searches, or looked at some articles. Or, maybe there was a TV show the night before that talked about Morocco that a lot of people saw. The eye-opening thing is that people don't need to listen to your conversations to figure out a lot of this stuff. Maybe that's even scarier.

I think I saw a survey that said a majority of Americans think that their voice assistant is spying on them, but they still use it. So, there's some deep stuff to untangle there.

HS OK. So, if you're not sending data anywhere, that must mean that all the processing is being done on the sensor itself, or rather, on the board that the sensor is on. That feels like the stuff of science fiction. How can you do that on a board that size and that price?

PW One thing I like to say is that AI wants to be on the edge. I did some calculations a few months ago. And by my estimates, there's about twice as much compute [power] available in all of the embedded microcontrollers out in the world as there is in all of the GPUs and all of the data centres in the world.

I like to think of that as 'dark compute', because it's kind of just sitting there. You know, things like an Arm Cortex-A, like the Raspberry Pi has, is ridiculously over-specced for most of the jobs it does.

And the same goes even for things like the RP2040 chip on the Pico, which most of the time is just blinking an LED or something like that. The rest of the time, it's just kind of sitting back not doing anything.

So there's all of this dark compute out there that's just not being used. And the nice thing about AI is that you can fit it into a couple of hundred kilobytes to do something like a computer vision algorithm to detect people. And then all you need is compute.

And so, we're able to run these AI vision and audio algorithms on these tiny, low-cost devices, because compared to 10–20 years ago, they actually have a lot of compute.

HS You mentioned RP2040. Are you working on anything using that chip?

PW Yes, we have a tiny QR code reader. It's so small that, when I ordered the packaging for it, I thought I'd made a →

Above
AI In A Box can do many things – such as generate real-time translations of spoken language

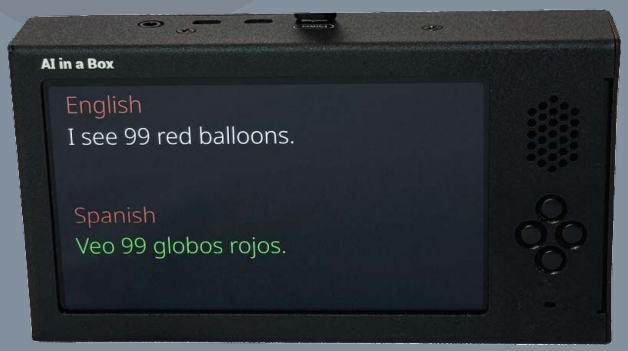

mistake with the dimensions, but no, it is that small [19 × 16.3 × 10.7 mm, smaller than a quarter coin]. You know the way some restaurants use QR codes for menus now? Somebody put together this box that uses one of our QR code readers. You point it at the QR code on the table, and it prints out a paper menu. That was really nicely put together.

It doesn't work offline, because you have to fetch the menu, but turning the QR code into a URL happens on the sensor.

HS OK, that's pretty cool. Other than extremely savvy tech people with lots of time on their hands, who is the market for your Useful Sensors?

PW We're working with a bunch of consumer electronics companies, things like audio, speakers, TVs, lighting. For example, we have outdoor security lights that only go off when there's a person around. So they don't trigger when there's animals or trees.

And the same for indoors. You know, if you've ever been in a really boring meeting where nobody's moving, the lights go off, and you have to stand up and wave your arms.

We have the solution to that. And that brings us back to privacy: if you've got a TV in your bedroom, you really want to make sure it's private, and nobody can turn the webcam on.

HS How can an AI model like this actually learn things on the fly when it's not connected to the internet?

PW It gets confusing when you talk about machine learning. Because these models, they have already been taught how to recognise people. And they're not updating themselves on the fly. They're like a compiled program. They're there as an executable; they're fixed. And it doesn't do any updates or any learning; it's already learned, by looking at millions of photos, how to tell if there's a person there or not. And it just keeps doing that job without doing any updates or changes.

HS So you've got to be pretty confident then, when you ship a person sensor with the code baked-in, that you've not just trained it to identify white men between the ages of 30 and 50.

PW Exactly. There's a lot of work and a lot of testing that goes in. … Does it recognise people who are lying down, because there aren't that many photos out there of people lying down? Obviously skin tone and gender are a big thing too. But also the different kinds of clothing that people wear in different countries. I'm not claiming we've got all of that perfected yet, but it's something that we're very conscious of and we're constantly trying to improve.

HS A slightly facetious question that I have written down here is have you trained the model to recognise men in kilts?

PW I lived in Scotland for a few years and did Scottish country dancing, and I've worn a kilt to a ceilidh, so I'm personally motivated to make sure that works. But I can't claim that I've dug out a kilt and tried it out myself.

And actually, on that topic, there was a fascinating research paper put out by the US Army where they challenged soldiers to fool a general kind of AI that was using computer vision to detect people.

And they came up with some amazing tactics, like putting themselves underneath a cardboard box. I think one team did cartwheels, because the AI wasn't trained to recognise people cartwheeling.

HS I presume that if you've bought a TV that recognises faces, you won't try to fool it. That seems counterproductive.

Your website says that step one in using a Useful Sensors device is to plug it in. Step two: there is no step two. How on earth does that work?

PW One important thing is that this is very different from most smart appliances. We hear from manufacturers that less than 30% of smart appliances ever get connected to the internet. Because, if you think about connecting a fridge, you have to download a special app, you have to connect to Bluetooth, connect to your fridge, you have to type in the network name and the password…

You have to create an account on the manufacturer's app. And you know, most people just never make it. Privacy is important, but I don't think that people are gonna buy purely based on privacy. Like they want to have something that works better.

HS Another claim on your website is that you're making AI less biased. How? What does that look like?

PW That's a good question. One of the things that I'm working on is a dataset called Visual Wake Words. And this is used in academic contexts, to try to train the best model to recognise people.

We've really been trying to make sure that that has as representative a set of images across cultures and genders and skin tones and everything else we can think of. So, I think that that's one practical way we're trying to do that.

We're also trying to capture and fine-tune data on low-resource languages for speech to text, you know, languages that don't have that many speakers, so there aren't that many audio files available of people speaking those languages that we can use to train the AI. →

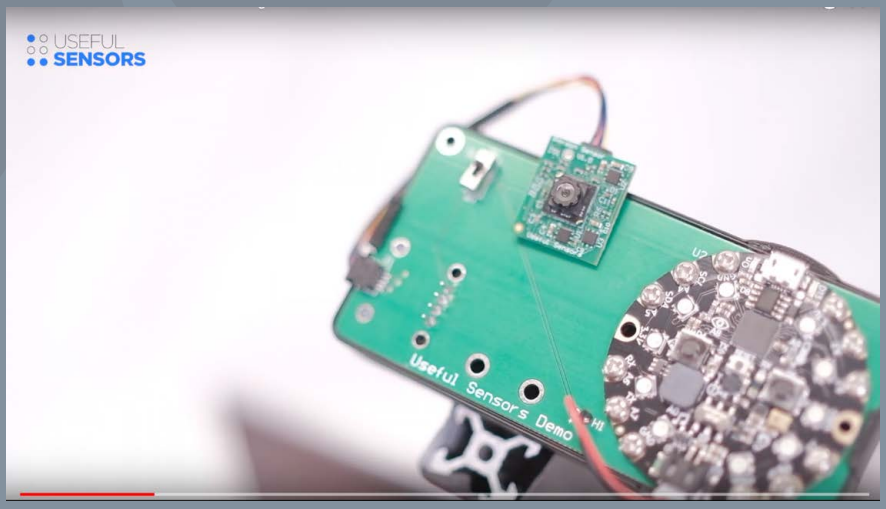

Above
Real-time transcription, without having a third party listening to what you're saying

HS The digital bill of health is another fascinating idea that's on your website. Can you tell us a bit more about that?

PW It's actually come from academia. It's something that consumer reports over here have been pushing.

And it came out of some work, I think it was at CMU [Carnegie Mellon University]. It's an idea that's been kicking around for a few years. And I really like it, because I think that if we can tell people what data is actually leaving their device, they can make informed decisions about what stuff they actually want in their homes.

The tricky part is figuring out what information you need to get across.

You don't want these to become like privacy policies where everybody basically copies and pastes a privacy policy that says, 'Oh, we can do anything with everything'.

So you can imagine a data nutrition label that just says, 'Oh, yeah, we're just sending everything up to the cloud'. And that would be kind of useless as a consumer tool. So, trying to figure out how we can not only verify that the nutrition labels are correct, with a third-party check. But also, actually try to encourage people to send less data, to go on a data diet.

HS How do you connect these sensors if you want to use one in your projects?

PW They use I2C, so they're really easy for the average person with a soldering iron to just slot them in. Actually, you don't even need a soldering iron, because I'm terrible at soldering.

You just buy a cable; it's known as Qwiic in the SparkFun world, and I think Adafruit has a different name for it [STEMMA QT], but it's the same thing. So you don't even have to do any soldering. You just plug in the cable at both ends. It'll only go in one way. So you can't get it wrong.

HS All part of the process of making things easy to use, right?

PW Exactly. And I've got examples for all the major popular boards up as code, circuit diagrams, and step-by-step instructions for how to get things running on Arduinos and Raspberry Pis and Picos, and all sorts of things.

HS You've also got something called AI in a Box. What does that do?

PW Pretty much what it says on the tin. We've taken a bunch of AI capabilities, everything from doing transcription – so giving you live subtitles for real life – to doing translation, so translating between different languages.

And we also have a large language model that runs locally on that box, so you could ask it questions and you can have a conversation with it. Also speech, so you're actually talking to it, and it's talking back, and everything's running with no network connection. And it's all entirely local.

Even though we talk about large language models, they're only like billions of parameters. So we actually chose one that was – I think – around 3 billion parameters.

It's definitely capable, and has learned a lot of general knowledge. You can chat to it about a bunch of different subjects and it won't embarrass itself. ◻

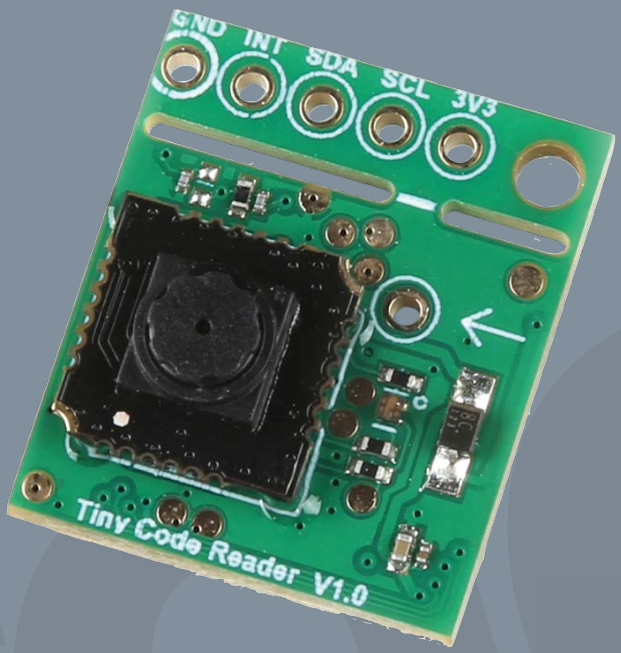

Left
The Tiny Code Reader uses a camera, a RP2040 chip and a fully realised AI model to turn QR codes into URLs. Yours for £6.71

Left
Useful Sensors' Person Sensor recognises people, so you can easily build projects that interact with human beings

Objet 3d'art

3D-printed artwork to bring more beauty into your life

One of the joys/curses of being a maker is that when you set out to make a 'thing', you often end up making tools instead. We've got loads of 3D-printed jigs clotting up our workspace, which were useful once but, alas, are now just taking up space. One 3D-printed tool that we know will see repeated use is this beautiful multi-axis vice by Printables user sdaendi.

It's fair to say that this vice won't replace our big, heavy woodworking vice, with its ability to hold large chunks of wood while we turn them into smaller chunks of wood; however, for electronics, for your Airfix Messerschmitt Bf 109 that you've been working on, or for anything small that you need to drill a hole into at a precise angle, this is perfect.

It's adjustable in two angles, and there are a few different attachments for different use cases, including one for holding PCBs.

Altogether, this print requires about 210g of filament to make all the attachments, as well as a few M3 10mm screws and two M4 20mm screws if you want to print the attachment that'll hold this vice to a table. □

hsmag.cc/MultiAxisVice

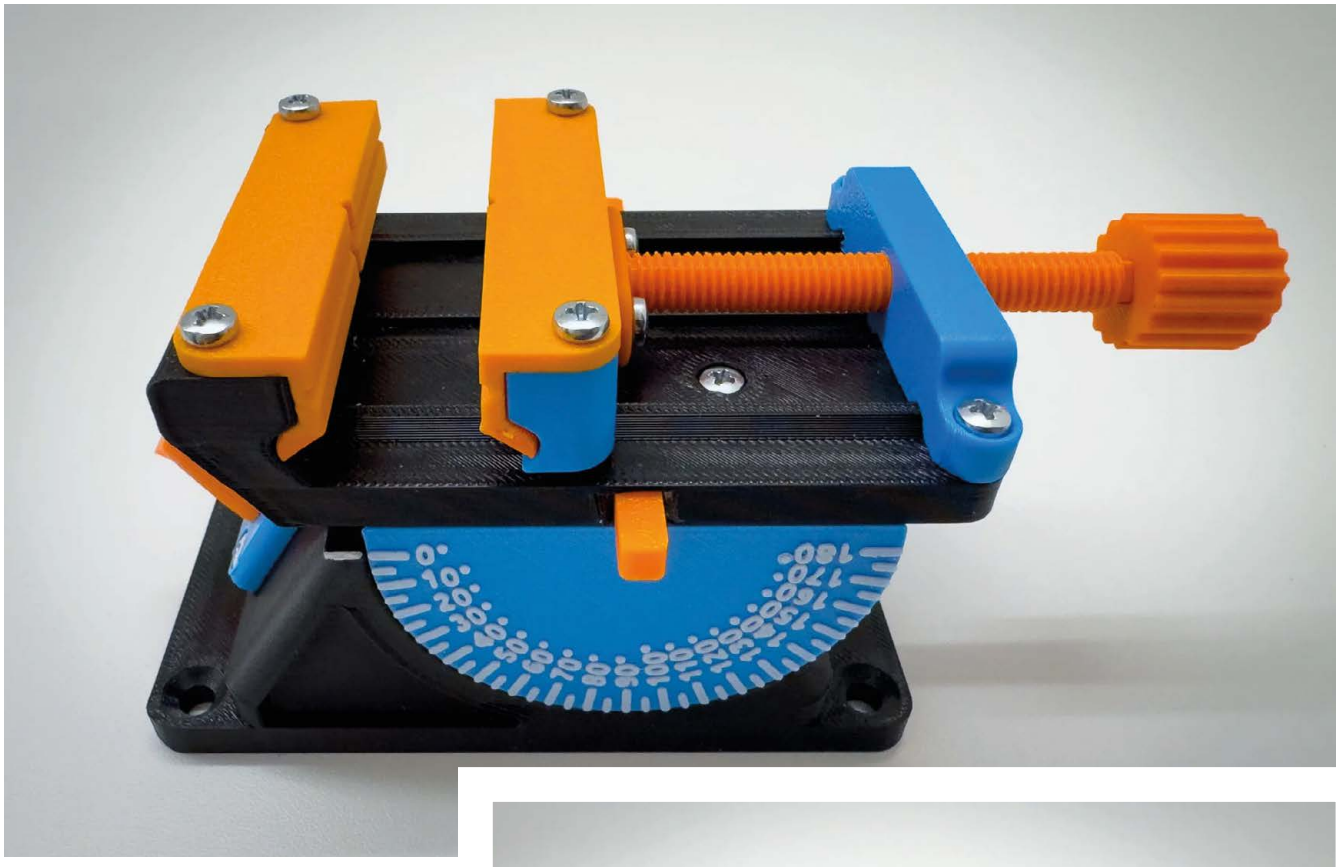
0
10
20
30
40
50
60
70
80
90
100
110
120
130
140
150
160
170
180

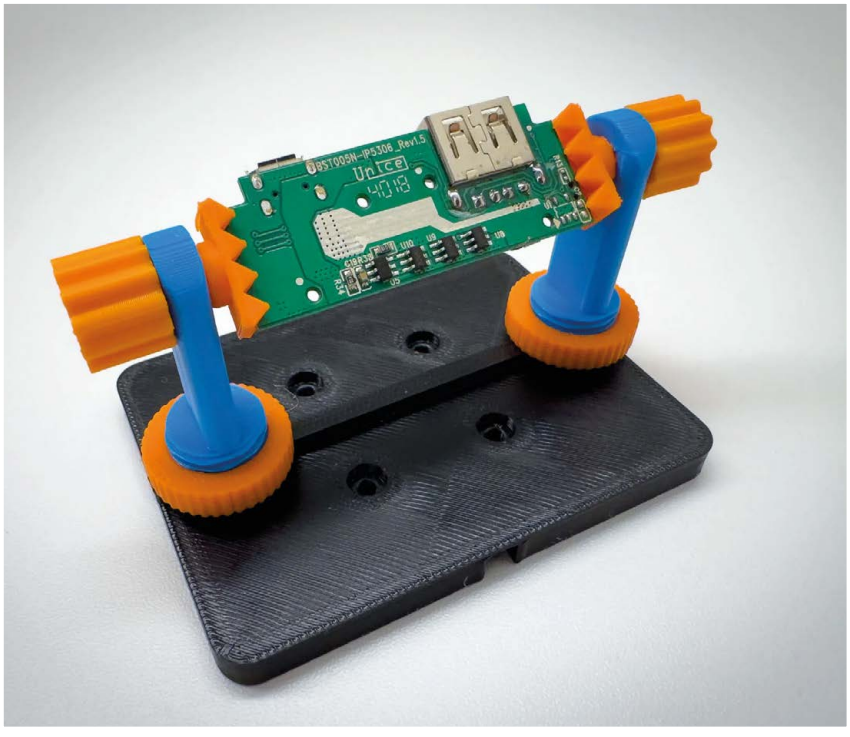
Unice

Letters

ATTENTION ALL MAKERS!

If you have something you'd like to get off your chest (or even throw a word of praise in our direction) let us know at **hsmag.cc/hello**

YOUTUBE

Youtube makes me sad sometimes. Like many people, I watch quite a few videos, and many of them are about 3D printing, but I'd never come across the work of Jón Schone before. I think it's because he takes his time and produces good content rather than constantly feeding the algorithm with huge amounts of mediocre tat. I wish the AI that decides what to feed into my eyeballs took this into account and gave me his content.

George
Glasgow

Ben says: The algorithm isn't aiming to show you what you want to watch; it's aiming to show you what will make you keep watching. This subtle difference is what means, trash, hate, and misinformation flourish while nuance and education can struggle. I wish I had an answer to this, but I don't. Just keep your eyes open for when interesting, thoughtful content does come up and do what you can to support those creating it.

EDGE CASE

Andrew, you coward, do a dovetail on your Raspberry Pi case.

Ben
Bristol

Andrew says: The dovetail is the show-off's joint. I'm secure enough in my craftsmanship that I don't need to show off. A rabbet joint is more than strong enough for these purposes and shows off the wood. Simplicity is sometimes best.

SOLAR

I occasionally drive past whole fields full of solar panels that are diligently converting the sun's radiation into power, and each time I do, I'm amazed that they're locked into a stationary position. I'm sure the companies that run them have done all sorts of cost-benefit analyses that say how to control them that says this is optimal, but that surprises me. At the moment, the limit on the amount of power we can get from the sun is limited by how fast we can produce solar panels, how much land we have, and how well we can connect solar farms to the grid. Even an increase of a few percent efficiency – if it didn't impact any of these factors – could have a noticeable impact on our fossil fuel usage, so surely it's worth it? Why don't they copy Tim Ritson's setup from HackSpace 79?

Freddy
Frome

Ben says: I know very little about the specifics of running a solar farm, but I suspect it may be to do with the difference between running one solar panel and running lots. If you tilt one panel to its optimum angle in the morning or evening (when tilting has the biggest effect), it will cast a shadow over its neighbour, so across the solar farm, you won't get a very big effect.

For an individual using a small setup, though, tilting can have a big effect – as it does for Tim Ritson.

PLOOPY

Those headphones look great! I sit at my computer for eight hours a day while I work, and listen to music (or my colleagues when in an online meeting) for most of them. I'd love to be able to support an open-source hardware company and also nerd out about the settings on my headphones. I can't wait to try them out.

Darren
Wells

Ben says: I had great fun with them, and still use them quite a bit. For cases such as yours (wearing them sitting down and plugged in), they're great – you can nerd out as much as you want to on the settings. They went far further than I understood how to use them. I too am a big fan of the open-source nature of them, and although I haven't yet done any 3D-printed additions, I do have some plans for it.

Electromagnetic Field 2024

The UK's premier geek get-together

Figure 1
An amazing tracked trike project, just one of many DIY vehicles on-site

Figure 2
The ever-favourite Hacky Racers had a busy weekend with lots of teams taking part

Figure 3
The stages and lighting are just one element of the amazing infrastructure EMF creates

Figure 4
The Tildagon, with its hexpansion slots, was often heavily modded and tinkered with

Figure 5
JTRUK live coding onto a large screen in the Field-FX tent using TIC-80

Figure 6
The night market in Null Sector was busy with people selling all manner of interesting items
Photo Credit: Laurence Stant.

Figure 7
Jarkman's popular 'World of Techno', which plays a different tune for every square foot of the planet, encourages you to take art for a walk

Once again, the rolling fields and grounds of Eastnor Castle have been the epicentre of the UK hacker, maker, tech, and geek scenes, playing host to EMF Camp.

For those who haven't been, it's an amazing volunteer-run camp with an astonishing amount of talks, workshops, events, raves, art installations, performances, and more. Around 3000 attendees all camped on-site, and generally hacked and relaxed in this vibrant pop-up community.

It's quite hard to write up as, from every square centimetre of the site, you can see something that's probably worthy of an entire article for HackSpace magazine. Let's start with talks. Probably the most eagerly anticipated and discussed was Tim Hunkin's 'A Short History of Electric Shocks', which was indeed excellent, and lovely to see such a pivotal UK elder statesperson of the hardware hacking scene taking to the stage. It's fair to say though, EMF talks have variety! We particularly enjoyed Dr Footleg's 'How to Map Caves', which wonderfully showed cave mapping techniques ranging from which type of propelling pencil you should use, through to 3D custom maps laid into and over OpenStreetMap and more. Truly wonderful to see such interesting, yet niche, techniques on a stage. Another highlight was Dr Laurence Stant's 'CERN Technology: Following the White Rabbit down a 27km hole', which gave an overview of CERN history with a focus on CERN technologies and innovations that have impacted the core of everyday life. We all know about the WWW and Tim Berners-Lee, but perhaps we don't know that CERN birthed capacitive touchscreen technology?

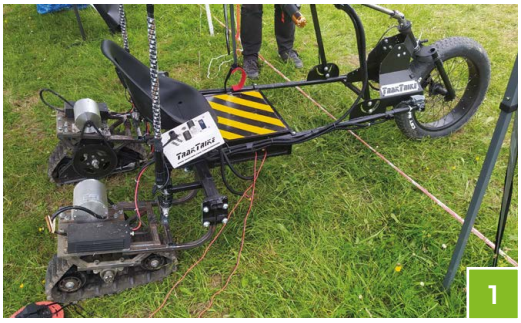

1

Beyond talks, there were lots of amazing projects and techie art installations every way you turned. Lots of interactive sitewide games and puzzles pop up, which keep the young and the young at heart occupied. EMF Camp has the concept of 'villages', where people can create camping areas folk can join that are linked to a theme. This might be simply a particular hack/hackerspace and its members, or it might be around a particular project or subject area. The villages are all friendly and welcoming and are great places to see and meet others. Milton Keynes MakerSpace built an excellent Tea Tent where anyone could drop in for a cuppa, and the Guild of Makers village was welcoming anyone with their bread making machines churning out delicious breads. At the more technical end, 'Odins Village' was full of all things high voltage, with some amazing Tesla coil performances. The spirit of support and cooperation was very much alive in this village, with members pulling great chunks of their machines apart to help others to create a massive Wimshurst generator, the discs of which came from two matching glass garden tables!


4

5


6


2

3

Hacky Racers ran events throughout the long weekend, and it's excellent to see that community grow and prosper. It was fab to see the race vehicles also being used as transport around the site. Speaking of which, it was joyful to see an amazing range of DIY vehicles at EMF, including an excellent trike 'TrakTrike', which consisted of a regular front wheel and two tracked systems from a snow blower providing the drive on the rear. Others of note included a small combustion-engined Henry Hoover, and there were all manner of hacked and modified bikes, buggies, and trailers to be seen.

On the non-hardware side, there was heaps going on. It's amazing to periodically update the 2024 EMF Tildagon Badge firmware and see that, over two days, developing contributors have rolled out update after update. The same can be said for the badge app store which steadily increased in size and technicality, with all manner of apps for the fabulous Tildagon. Live coding could be spotted in a few places, but perhaps was most visible in the Field-FX tent which hosted numerous bytejams and workshops on fun coding environments like TIC-80.

The EMF swap table was constantly well attended, with people bringing and leaving with all manner of things. It's also a nice area to hang in and hear chats about older technology and spot interesting vintage items.

In the evenings, the lasers came on and the rumbling bass of the null sector started to rise. This amazing section of EMF Camp was built from a phalanx of shipping containers and, this year, contained all manner of solar-punk and biomedical-themed installations. A full nightclub in a quarry pumped out tunes; there was archery with clever virtual targets, Laser Quest, and a whole heap of wonderful installations. From old school telephone exchanges to Jarkman's Air Giant inflatables, it was a beguiling sector to wander around, capped off by a busy night market where traders show and sell their wares ranging from robots to cyberpunk club wear.

It's hard to attend EMF and not leave utterly inspired. Look out for it returning in 2026 and join in! □


7

THE *Official* RASPBERRY PI HANDBOOK 2024

200 PAGES OF RASPBERRY PI

- QuickStart guide to setting up your Raspberry Pi computer
- Updated with Raspberry Pi Pico and all the latest kit
- The very best projects built by your Raspberry Pi community
- Discover incredible kit and tutorials for your projects

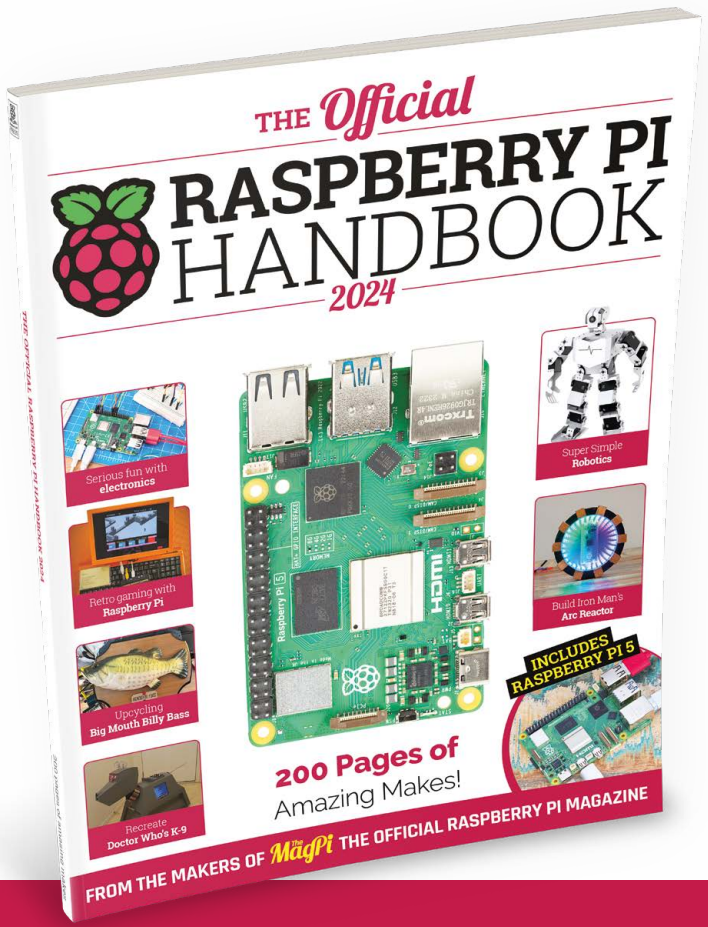

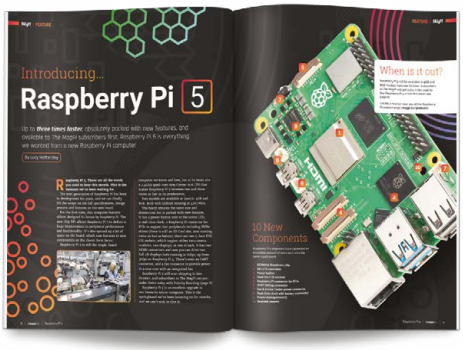

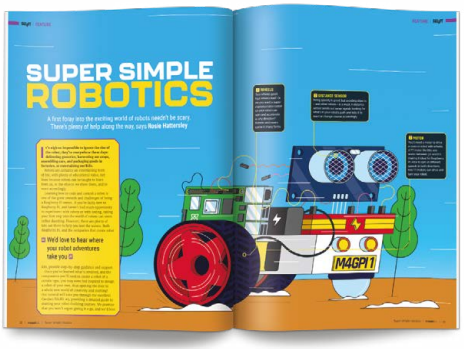

Buy online: magpi.cc/store

HackSpace
TECHNOLOGY IN YOUR HANDS

FORGE

HACK | MAKE | BUILD | CREATE

Improve your skills, learn something new, or just have fun tinkering – we hope you enjoy these hand-picked projects

3D pens

Is a 3D pen just a hot glue gun on steroids? Let's take a closer look

Nicola King

Nicola King is a freelance writer and sub-editor. She tries to produce beautiful things but, as you'll see, sometimes they're not quite as beautiful as she had hoped…all part of the learning process.

Above
A little wonky, and certainly unlike any frame you'd pick up in Specsavers… which is probably just as well or they'd go out of business

Meandering around the supermarket on our weekly shop a couple of months ago, we chanced upon a craft-related gizmo in the 'middle' aisles, priced at the princely sum of £24.99. This author's husband, always one for a new gadget, couldn't resist popping it in the trolley, citing the fact that it would be a great subject for a little tutorial, so we exited the grocery store with our broccoli, cheese, sausages, bin bags, and…oh, a 3D pen set.

Now, it's worth trying to manage some expectations from the get-go on this topic – a 3D pen will never (ever) be able to produce anything close to the quality of output of a 3D printer, but they can be fun to use and are arguably a great way of introducing younger members of the family to the concept of 3D printing. We read one review of a 3D pen online by someone who had tried one out with her young daughter to pique her interest, before going that step further and investing much more heavily in a 3D printer. 3D printers are, let's be honest, on the expensive side and 3D pens are much kinder on your wallet, while still being able to bring your imagination to 3D life. With that in mind, let's investigate what you can expect from such a device.

FREEHAND 3D PRINTING… SORT OF

You might initially think that a 3D pen is something of a gimmick, but they can be used as genuine artistic tools that can help you sculpt in 3D, converting plastic filament into melted material with which you can create objects. So, if you want to create a 3D model by hand, these are worth a try. Here, we are bypassing the software that a 3D printer uses, and controlling it ourselves, with no difficult tech to master. Importantly, before we move on, let's also remember that we are working with molten plastic – it's hot, so keep your digits, and other parts of you, well away from the heated tip. That's especially important to remember when you are pulling any residue off the tip when you have finished, or are changing colours.

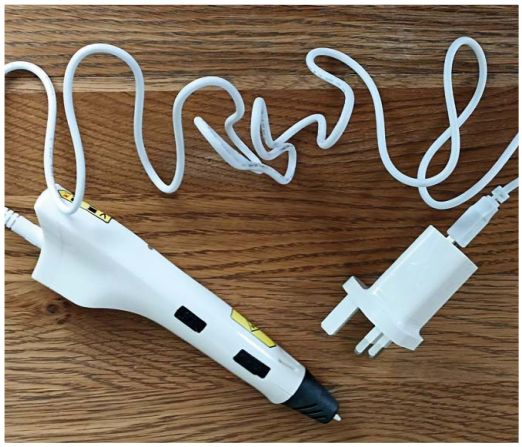

Above
The contents of the kit that we purchased, including an impressive selection of filament colours for the price – ten colours of 5m each in length, 1.75mm in diameter. It incorporated everything that we needed to get underway

Left
The USB cable is a good length, and a plus-point of this device is that it turns itself off when you set it aside after two minutes – always a handy feature. Just press the On/Off button to heat it up again

YOU'LL NEED

- **A 3D pen**
- **Appropriate filament** in a variety of colours
- **Templates to work from** (either bought or drawn by yourself)
- **Power source** (ours came with a USB cable and adapter)
- **A mat to protect your surface**
- **A 'drawing' mat to hold over your template** (clear plastic)
- **Scissors**

So, first things first, we attached the USB cable, and selected the PLA switch on the side of the device (the type of filament provided). Once the On/Off switch was pressed, we waited for the pen to heat up. Interestingly, it was really quick – after about 30 seconds, the pen had reached its required melting temperature (according to the accompanying booklet, around 190°C), the indicator light turned green, and we were ready for the off. Next, we chose the colour of filament that we wanted and fed it into the filament entrance with the straight-cut end (apparently that's important). After a few seconds, melted filament was coming out of the tip and we were able to get creating.

The truth is, at the risk of shattering some illusions, the absolutely best way to start is to begin by creating shapes in 2D on the flat… yes, we know it's a 3D pen but we have to start somewhere, we are beginners, and we are not going to run before we can walk. Plus, that's the advice we've picked up from online, so we decided to follow it.

Once you've got some 2D shapes, you can then weld them together to create a 3D shape. →

WHAT TO LOOK FOR WHEN PERUSING THE 3D PEN MARKET

So, what are the key things that you need to consider before buying?

- Comfort of use – 3D printing projects are not quick, and usually take a significant amount of time to undertake, so how the pen feels in your hand is very important (ours weighs 61 grams). You need it to be lightweight too, as well as ergonomically pleasing. Some pens are larger than others, wider or longer (ours was around 18cm long, with a tip diameter of 0.7mm).
- Reliability – you need a steady, even strand of filament coming out of the nozzle, at a comfortable speed, which is partly down to the dexterity of the user, but is also down to the quality of the pen itself. You don't want your pen prone to filament clogs.
- Does the pen come with everything you need, such as clear instructions and some templates, or even some filament? These are very useful to have for a newbie.
- Cost – what's your budget? Many 3D pens on the market won't break the bank.
- Finally, as with anything, it really comes down to who you are buying the pen for. If you consider yourself a high-end user who plans on using it a lot for design-type projects, and want all the bells and whistles, then you will need to pay for those features. For example, pens with digital displays where you can select the exact temperature you require will come at a more premium price. If you are overseeing younger makers, then you need to be conscious of the age-range specified on the box, the size of the pen, and it goes without saying that you need to be using a non-toxic filament at a lower temperature.

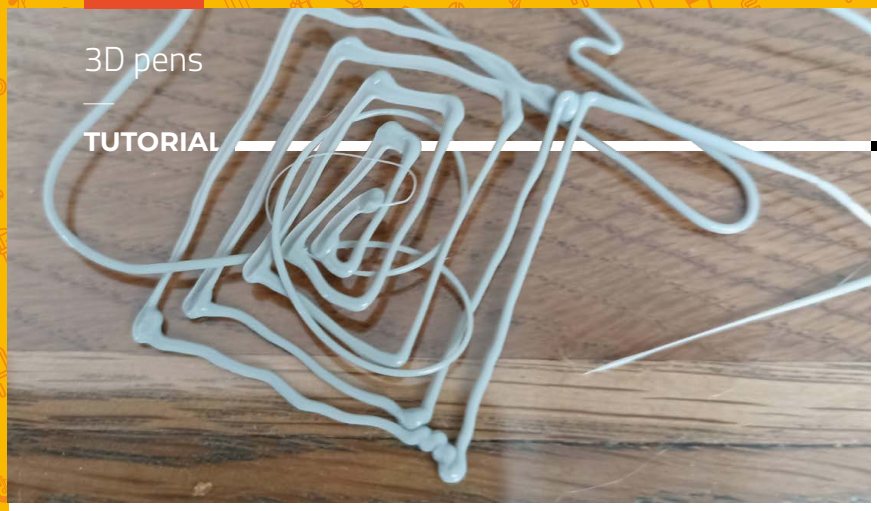

Left
A hot mess of plastic – take your time to practise and get the feel of the pen and the speed of extrusion before you start creating

Below
Place your clear mat over your template before you begin and, to prevent heat marks, hold the pen around 0.5 cm away. It's best not to touch the plastic mat with the tip

QUICK TIP

Protect your work surface from plastic residue and use a mat of some kind. A hobbyist cutting mat will do the job, and makes it easier to move your work around and lift it off.

> **Trying to draw something in the air, so that you really feel** like you are creating a 3D shape, can be very tricky

FILAMENT FAMILIARITY

There are a few options to ponder when it comes to which meltable plastic to select for your 3D pen. Ours could use all three types detailed below. So, let's take a closer look:

- PLA (Polylactic Acid). This is arguably a top choice for use with 3D pens, for a number of worthy reasons. Firstly, it is a natural thermoplastic polyester derived from renewable resources such as soybeans, corn starch, or sugar cane, so it is a bioplastic. It is technically biodegradable (under very specific conditions) and sustainable. Many people find it easy to use PLA with 3D pens due to the low melting temperature, and it does not have a strong smell or give off lots of fumes, something that many makers are keen on avoiding if they can, and great if you are overseeing children using a 3D pen. It's available in many colours, even some that glow in the dark. You can even purchase 'wood filament' for some pens, which is not actually wood, but wood fibres added to a PLA base for a wood effect when extruded. PLA is often used in food packaging and plastic cups, so this underlines how safe it is. It is also rigid and strong; however, poor heat resistance properties make it mainly a hobbyist material – as it is made of organic materials, it has a higher permeability than some other options. If you are looking for longevity in your creations, PLA may also fall a little short.

- ABS (Acrylonitrile Butadiene Styrene). Now this filament is considered more flexible and tougher than PLA, as well as more durable, and is therefore great for 'bridging', i.e. extruding your filament with your pen between two points, with no support from below. Projects are also likely to last longer than PLA creations. However, ABS is not as environmentally friendly, as it is a petroleum-based thermoplastic and non-biodegradable. Because of this, when it's heated ABS can produce an unpleasant smell and any fumes are toxic, so you really need to ensure your ventilation is top-notch. You can use it in your pen; you just need to take the correct precautions. Also, consider that this filament can warp and shrink somewhat when you are working with it.

- PETG (Polyethylene Terephthalate Glycol). Some pens will also take PETG filament, which is safe, strong, and resilient, but it's not biodegradable. That said, it can be recycled.

When choosing your filament, you will also need to think about the filament diameter (many 3D pens use 1.75 mm), the total amount of filament required, tolerance, and compatibility with your pen, as well, of course, as the colours you want to utilise based on the project you are embarking on. Finally, it's worth mentioning that some 3D pen makers want you to use their own proprietary filament, which could be seen as a limitation as it may cost you more – just check this out before you buy.

PLASTIC PRECISION

We wanted some control over the filament flow, as when we first tried to use it, the filament came out way too fast and we had no clue what we were doing. So, we opted to operate the feed manually, which involved keeping the filament feeder button pressed down while we worked. If we wanted to stop the filament feed, we just released the button. There is also a speed switch on this pen, and three options to make the filament flow faster or slower. When you first start out, it's advisable to select a very slow speed.

We chose to create a pair of 'spectacles' as our first foray and, while they are obviously the work of a complete amateur, they held together when we welded the arms to the main body. A butterfly swiftly followed and we were well on a roll.

A STEP FURTHER

Now, trying to draw something in the air, so that you really feel like you are creating a 3D shape, can be very tricky and takes practice. For example, you could start off by trying to draw a cube. Begin with a square base, then from each corner draw a line upwards into the air (easier said than done!). Level them off and then you need to practise 'bridging' or connecting two points in space so that you can complete your cube. This is hard if the two points are not stable, and if you're new to using a 3D pen, this may be a little frustrating, which is why it's much easier to start in 2D, welding pieces together and then progressing to more challenging designs.

Another idea is to write a word (joined-up letters so it holds together) and then lean it against a hard surface. Then, with your pen, draw some struts out from either end so it can stand up. Have a play and see where you end up.

Above
This author's husband had a go and produced this very appropriate offering… we did say we need some practice!

Left
Mix your colours to practise changing pen colour filaments. It helps you get a feel for the various buttons on the pen too

EXTRUDING EVALUATION

In conclusion, we enjoyed trying out our very inexpensive 3D pen, but the phrase 'you get what you pay for' definitely comes to mind, and we'd be really interested to know what the slightly more expensive pens on the market are capable of – from the £50 to £100 price mark. Even though this was an inexpensive option, we were still pretty happy with the results and can see that practice and some time investment would definitely improve the finished articles. Be mindful that you will definitely not be able to create an exquisitely detailed model of the Statue of Liberty the first time that you use a 3D pen. It takes time and patience to learn how to keep your hand steady, to manipulate the filament, change colours and so on, but it's definitely worth having a go to see what you can create.

There are some who disparagingly call these cheaper-end pens 'toys', but seriously, what toy has a tip that hot? We feel that description is a little unfair and misleading, and prefer to think of the 3D pen as a creative tool that takes a little practice to master with a steady hand, and that just happens to be a lot of fun.

MENDING MISPRINTS

If you are a stalwart 3D printer user, and have completely discounted ever using a 3D pen in any capacity, do bear in mind a rather handy aspect of owning one – you can use it to smooth out any issues on 3D prints from a 3D printer. Basically, it is very useful for 'gluing' or bonding pieces together, or touching up anything that your machine has printed that did not turn out quite how you hoped it would, so don't underestimate its usefulness. It's an alternative to superglue and worth a try to see how it works for you. We suspect ABS filament would give a stronger adhesion than PLA, but have a go and see how effective it is.

Assuming your filaments in the two devices are the same (which obviously makes colour-matching very easy), you can repair cracks, gaps, and holes in 3D-printed items. For example, two broken parts can be welded back together, as long as you work with a slow extrusion rate to allow more precise work. Any excess when you've finished can be trimmed off with scissors, and then you can sand it down.

Here's a handy link to give more insight: **hsmag.cc/3DPenWelding.**

QUICK TIP

Never leave a 3D pen unattended. Use it on a heat-resistant surface and don't use it near flammable materials – if these pens can melt plastic, they can melt other things too!

QUICK TIP

Don't over-extrude! By that, we mean don't push out too much filament when you are working on a project. Also, practise your hand speed, and don't move in a jerky manner as you work or that will result in unsightly blobs.

QUICK TIP

If you have pieces/bulges on your design that you want to remove, you can use the tip of the nozzle to easily melt away small blips and errors. You can also use scissors for this, but maybe avoid using your best scissors in case they blunt.

Your FREE guide to making a smart TV

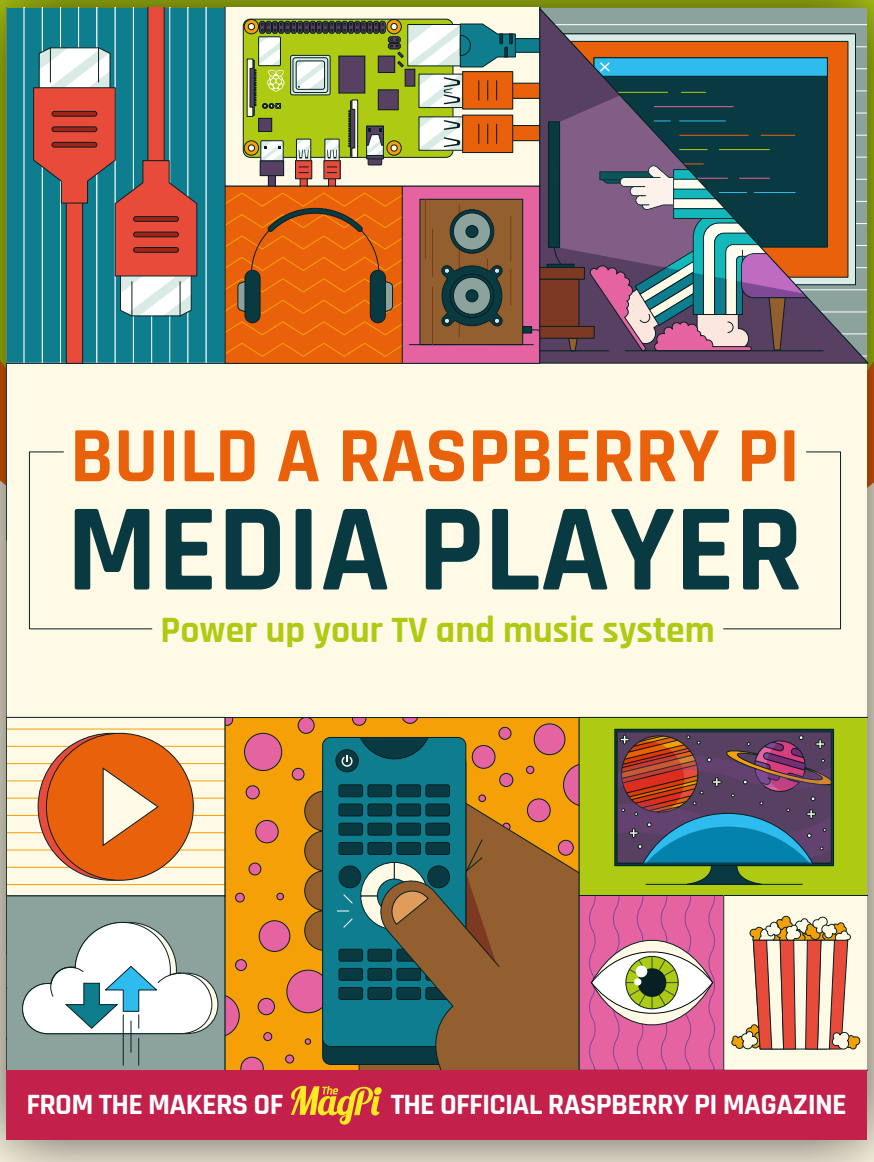

magpi.cc/mediaplayer

Free eBook!

Download your copy from

hsmag.cc/freecadbook

TUTORIAL

Part 02

Upcycle a Sonos Play:1

So far, we have converted a Sonos Play:1 into a Raspberry Pi-driven powerhouse of sound. Now it's time to trick it out with software

MAKER

PJ Evans

PJ is a writer, software engineer and tinkerer. He got fed up listening to Rick Astley (bit of a let down) and now plays The Stick Song on repeat.

mrpjevans.com

You'll Need

- Raspberry Pi Smart Speaker build **magpi.cc/139**
- Installation of Home Assistant (optional)

Whether you're considering converting your Sonos Play:1 or perhaps upcycling a vintage radio, you should have a project in mind by now that'll fill your room with sound, Raspberry Pi-style. Previously, we walked through the hardware build of our super-smart speaker project but so far we have only been able to play 'white noise'; where are the tunes? Now that we know the audio is working, we can focus on what we want to listen to and how to control it. We've selected some common uses for smart speakers, but feel free to explore your own ideas!

01 Wired for sound

We start off with a small stereo problem. The Sonos Play:1 only has a single speaker and it's not a great idea to wire it to both outputs of the amp HAT. Instead, we can reconfigure sound output to mix both channels to the left speaker (which is the one that should be wired up). We can prove this by running:

```
speaker-test -c 2
```

The Play:1 will only play the left-hand white noise signal (**CTRL+C** to stop). To fix this, start by getting the 'number' of your card by running the following command:

```
$ aplay -l
```

The output will probably either start with 'card 0' or 'card 1'. Make a note of it.

02 Two become one

Luckily, we can fix our stereo problem by mixing the two channels in software. To do this, we'll configure the ALSA audio sub-system. Start by creating a global configuration file (it may already exist):

```
$ sudo nano /etc/asound.conf
```

If it already exists, replace all the text with the contents of the code listing. You can also grab the code from **magpi.cc/asoundconf**. If your card number was something other than 0, then make sure you replace all instances of 'card 0' with the correct number (e.g. 'card 1'). Save the file (**CTRL+X** followed by **Y**). No need to reboot, just run `speaker-test -c 2` again and now both channels should be played from the speaker.

03 Pump up the volume

Before tricking out our new speaker, it's a good idea to set the base volume. You'll end up with multiple levels of volume controls depending on your choices. For example, streaming from a phone will have the phone's output volume feeding to Raspberry Pi's volume. If you find you want to set the output of Raspberry Pi's amp, just run the `alsamixer` from the terminal. You'll see a rudimentary volume control where you can use the cursor keys to adjust the master volume output. If you can't see any controls, press **F6** to get a list of available sound devices. When done, **ESC** will return you to the command line.

04 Local hero

Maybe at this point, you just want to be able to play some music or other audio. After all, there's no need to turn this into a

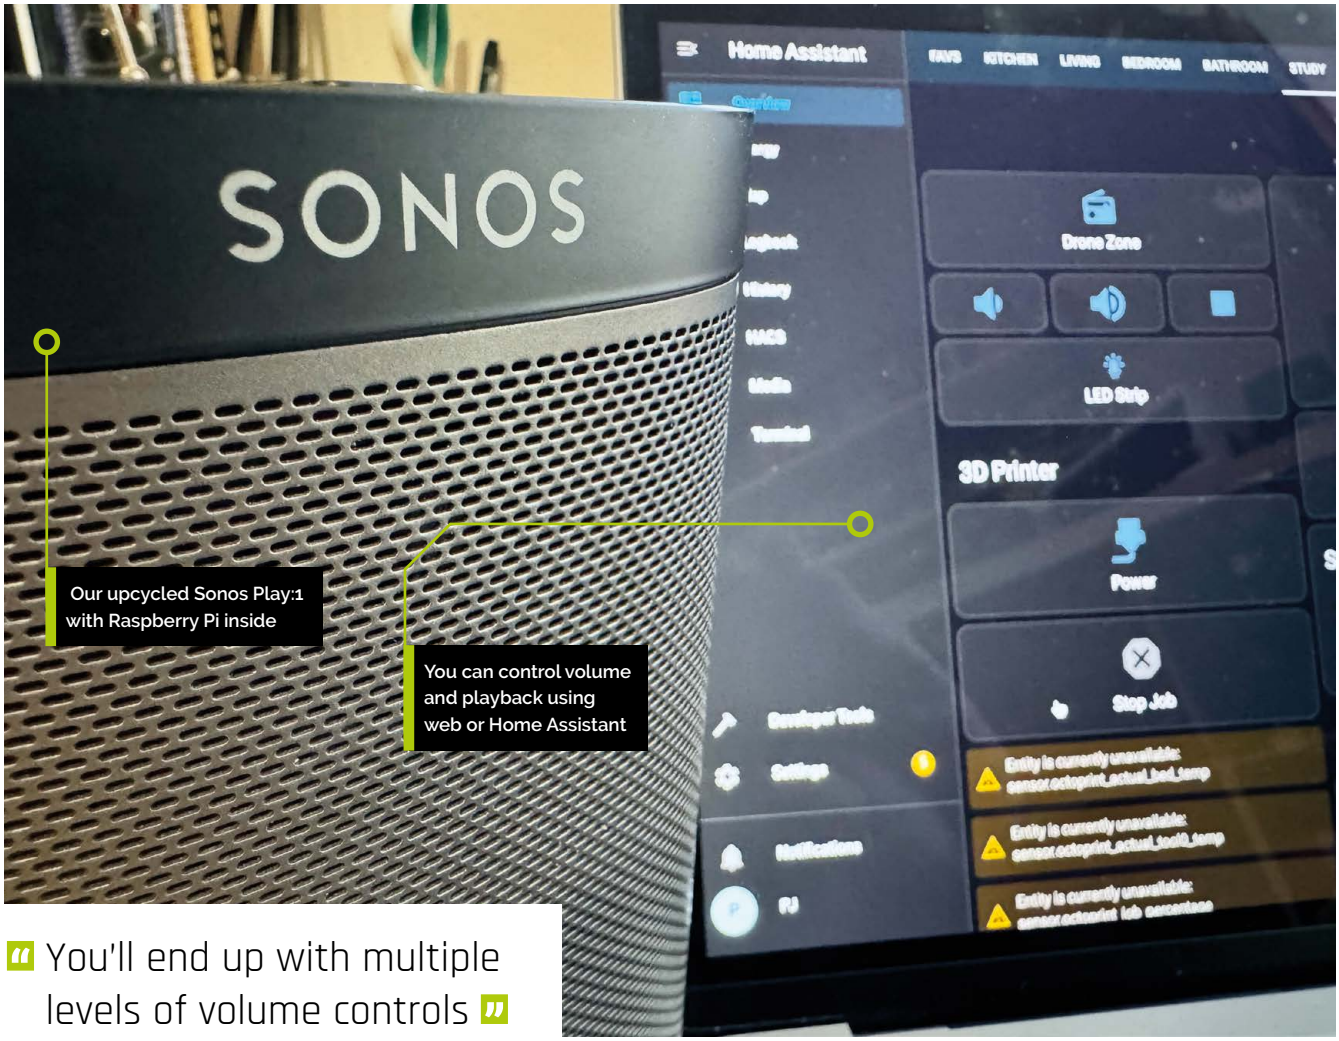

> You'll end up with multiple levels of volume controls

asound.conf

DOWNLOAD THE FULL CODE:
magpi.cc/asoundconf

Language: **BASH**

```
pcm.card0 {
  type hw
  card 0
}

ctl.card0 {
  type hw
  card 0
}

pcm.monocard {
  slave.pcm card0
  slave.channels 2
  type route
  ttable {
    # Copy both input channels to output channel 0 (Left).
    0.0 0.5
    1.0 0.5
    # Send nothing to output channel 1 (Right).
    0.1 0
    1.1 0
  }
}

ctl.monocard {
  type hw
  card 0
}

pcm.!default monocard
ctl.!default monocard
```

streaming device, you've already got enough of the project working to play audio. If you've installed a desktop environment, you could also use it as your computer! For those comfortable with the command line, you can upload your audio files and play them back with mpg321, a handy command line utility. To install: `sudo apt install mpg321`. There are many apps to help you manage and listen to your music collection such as **volumio.com**.

Top Tip

Support SomaFM

SomaFM is free to listen and, bar the odd ident, completely ad-free. If you build this project, consider supporting them: **somafm.com**.

05 Story of the blue(tooth)s

At this point, it would be reasonable to expect a section on converting our speaker into a Bluetooth-capable endpoint. Sadly, Bluetooth is a complex beast and although it is possible to turn the build into a Bluetooth speaker, there are so many different approaches that vary depending on which hardware or operating system you are running, that it would take several tutorials to cover all eventualities. So, to summarise, it is possible, but requires a lot of low-level hacking about. If you still want to give it a try, start with this guide: **magpi.cc/bluealsa**.

06 In the air(play) tonight

Although Bluetooth is a bit of a pain, the good news is that owners of Apple devices can create the same capability using the AirPlay standard. In the terminal use the following commands to build and install the latest version (which supports Raspberry Pi 5):

```
$ sudo apt install autoconf libtool libdaemon-
dev libasound2-dev libpopt-dev libconfig-dev
libssl-dev libavahi-client-dev git
$ cd
$ git clone https://github.com/mikebrady/
shairport-sync.git
$ cd shairport-sync/
$ autoreconf -i -f
$ ./configure --with-alsa --with-avahi --with-
ssl=openssl --with-systemd --with-metadata
$ make
$ sudo make install
$ sudo systemctl enable shairport-sync
$ sudo systemctl start shairport-sync
```

▼ Here's how the AirPlay service appears on your Apple device

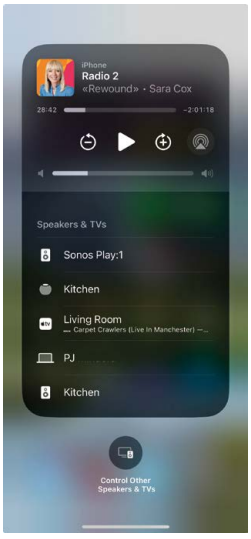

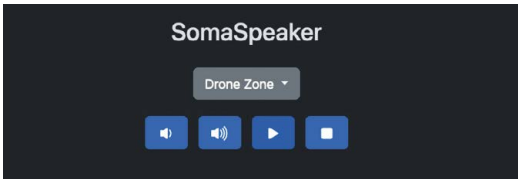

▲ Our basic web interface is a good starting point for your own project

Once installed and running, your iPhone/iPad/Mac should be able to see 'Audio' as a streaming target. Try it out and you should get sweet, sweet music from your speaker. You can change the name by editing the config file: `sudo nano /usr/local/etc/shairpoint-sync.conf`. You'll need to reboot after this.

07 Radio waves

We'll now turn our attention to media streaming. There are countless audio streaming services out there and one of our favourites is SomaFM, a collection of ad-free eclectic stations providing music for all kinds of tastes. We like its ambient channels which are both relaxing and can aid concentration. In the next few steps, we're going to turn our speaker into a SomaFM receiver for Drone Zone, SomaFM's most popular ambient service. Then we'll add MQTT support so we can control the speaker remotely. There's a lot of code involved here, so we have prepared a repo with everything you need. From the terminal:

```
$ cd
$ sudo apt install jq mpv python3-pip
$ sudo pip3 install flask paho-mqtt
$ git clone https://github.com/mrpjevans/
somaspeaker.git
```

08 Station to station

Let's make sure you can listen to SomaFM. Try these commands:

```
$ cd ~/somaspeaker/
$ ./somafm.sh channels
```

...you should have a list of channels to choose from. Now try this:

```
$ ./somafm.sh listen dronezone
```

If you can't hear much, check alsamixer for the volume. Use **CTRL+C** to stop when you're done. Have a listen to the other channels too, you might find a favourite!

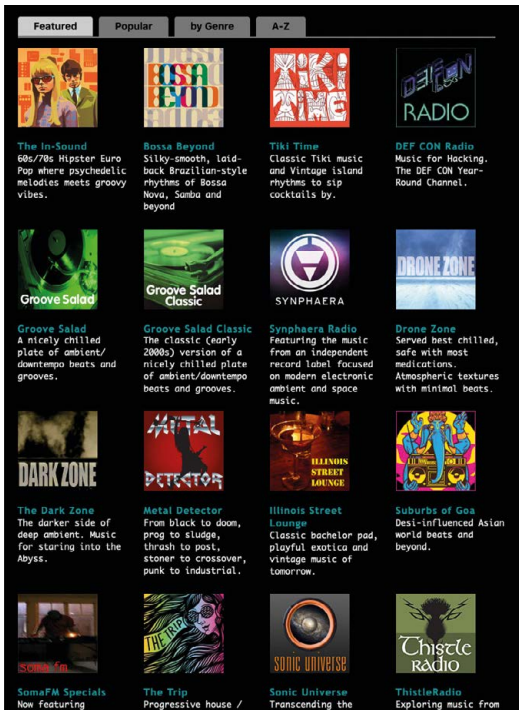

SomaFM is a great collection of music genres

09 Ready to go

Now we have our SomaFM streaming device, we could use a way of being able to remote control both playback and volume control. For this project, we're going to create a small web service which can be accessed on your local network. The code is already in the repo you downloaded in the previous step and we encourage you to investigate it and make changes to suit. To start the service manually, just run this from the directory:

> Listen to the channels, you might find a favourite!

```
$ flask --app somaspeaker run -h 0.0.0.0 -p 3000
```

You should now be able to access the site on http://(your IP address or hostname):3000/. Try starting the station and changing the volume.

10 Start it up

Let's round things off by making sure our web site starts when the Raspberry Pi boots. To do this we need a service file. To save you some typing, we've got one ready to go. Make sure the web server is not running. Now, run the following commands to set everything up:

```
$ cd ~/somaspeaker
$ sudo ./install_web_service.sh
```

You should be able to access the server as before. Reboot your Raspberry Pi and you should find the server is automatically running.

11 Message in a bottle

If you are using Home Assistant, you can embed controls for your speaker right into the dashboard. To do this you'll need an MQTT broker such as Mosquitto. If you have one, edit the `config.py` file in `~/somaspeaker` to enable MQTT support and set the broker's address. Upon restart the server will subscribe to the given topic and listen for instructions. You can then use Home Assistant's MQTT integration to send commands to the speaker. See the provided README.md for examples.

12 More, more, more

You've now got a smarty-pants speaker that can play local files, act as an AirPlay target and stream SomaFM channels. That doesn't have to be everything. You can try adding other streaming services, or create a cool Jukebox app. How about AI-generated playlists, or generative music which can react to light and temperature? Take Home Assistant further by automating playback (e.g. when you switch a light on). Whatever you decide, you have hopefully prevented another piece of e-waste from going into landfill and you've got a great speaker system to go with it.

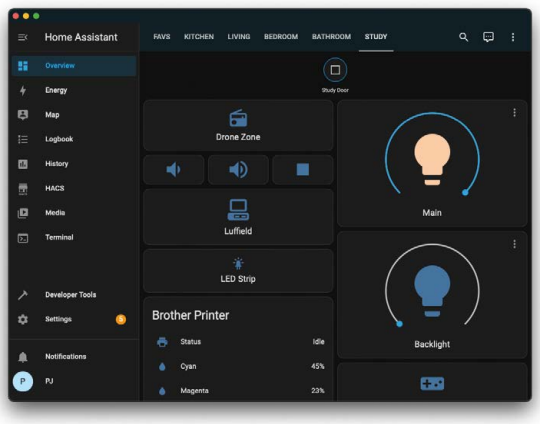

You can customise Home Assistant control playback, even adding automations

Top Tip

Equal rights

If you would like more control over sound quality, equaliser plugins are available for ALSA. Example: **magpi.cc/alsaequal**

Needful things: Hot stapler

Fix broken plastic the right way

Above
Hot staplers are great value and often come with a good selection of staples. In addition to regular straight staples, this kit also comes with staples shaped for fixing inside and outside corners

Dr Andrew Lewis

Dr Andrew Lewis is a specialist fabricator and maker, and is the owner of the Andrew Lewis Workshop.

There are some tools that pretty much everyone has in their toolbox. Hammers, screwdrivers, and spanners are ubiquitous in modern society, and most people know what they are. There are other tools that are relatively specialised in nature, or are just not very well known outside of a certain job. We are not talking about the typical Christmas gift gadgets that get buried in the cupboard; we are talking about tools that are used widely and can be massive time-savers. While you don't need to use them every day, they are valuable tools that can save hours of work when you *do* need them and will leave you wondering how you ever managed before you had them. One of these tools is the hot stapler, and you'll find out all about it here.

Right
Aside from making a good prop for cosplaying, the design of hot staple guns is fairly consistent: a pistol grip with a trigger and a wired cable. Unfortunately, most staple guns are quite large and might be hard to handle in tight spaces, even if they have an angled head

If you've ever cracked a plastic car bumper, suitcase, or car body panel, you'll know that gluing the pieces back together never works well. Even if you build up the cracked area with resin and fibre, there's a good chance that the crack will reappear after a few days of use. This is where a hot stapler comes in. As the name suggests, the hot stapler fixes the pieces of plastic back together using a combination of metal staples and heat. The mechanical support provided by the staples is usually enough to stop the plastic pulling apart, and should allow you to apply whatever cosmetic fixes you need after the crack has been stapled.

Hot staplers look something like a Weller soldering gun, and there's a good reason for that – they both use the same technology to get the job done. While a soldering gun has a fixed metal tip that heats up when electricity is passed through it, the hot stapler uses the staple itself as the heating element.

To get the best out of a hot stapler, you should start by preparing the crack that you're trying to fix. Make sure that the surface of the plastic is clean and free from any contaminants that could react negatively to the application of heat.

Position the plastic so that the crack is closed, and as close to the original position as you can manage. It's worth spending a little bit of time here getting everything aligned properly, as this will give you a much better result in the long run. Stabilise the cracked pieces in whatever way seems best. For a simple crack, you might be able to hold the

> **To get the best out of a hot stapler, you should start by** preparing the crack that you're trying to fix

pieces together with just your fingers. For larger cracks, it's better to use some sort of clamping to hold the pieces together. Thin strips of gaffer tape can be useful for holding the piece together before stapling, or even a couple of dabs of superglue or hot glue, depending on the type of crack that you're dealing with. You won't be relying on the glue to hold things in place in the long term, just while you're making your repair.

With the crack closed and stable, plug in the hot stapler and insert a staple into the two prongs on the front of the staple gun. Pull the trigger to heat the staple for a few seconds, then gently push the staple into the plastic until the zig-zag portion of the staple is buried below the surface, straddling the cracked section.

Release the trigger and let the staple cool for a few seconds, then pull the gun away. The staple should now be embedded in the plastic, with two metal prongs poking out. Repeat this process every few centimetres along the length of the cracked section, paying particular attention to places where the crack forks in another direction or where the plastic is likely to experience the most strain.

Trim the metal prongs flush with the surface of the plastic, and then apply whatever cosmetic effects you want to minimise the visible portions of the crack. That could mean shaving flat with a blade, filling holes with resin, painting, or vinyl coating, or just flashing the piece with a propane torch for a second and texturing it by pushing fabric or sandpaper onto it to set the surface.

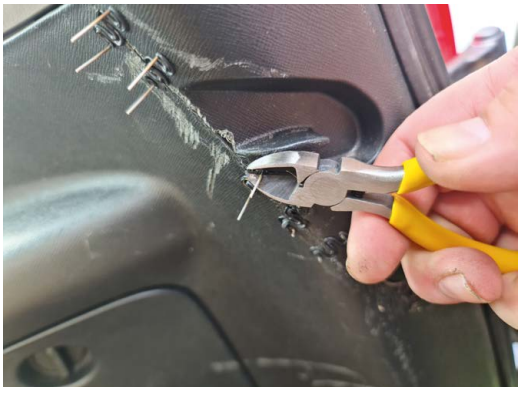

QUICK TIP

A wise engineer once told me, 'You can never hide surface imperfections completely, so make a feature of them instead'. It's great advice. He also said, 'Slap a bit of paint over it and stand in front of it when the Admiral walks past', which is at least honest advice, and also works well.

Above Left
This crack needs to be realigned before it can be stapled back together. Stapling the sides will make a strong mechanical joint, but it won't fix any problems caused by poor alignment

Left
Push the hot staple into the plastic and wiggle it about a little bit to make sure it beds in. Let the plastic cool before you remove the staple gun, or you might pull the staple back out again

Left
The prongs sticking out of the plastic can be cut with a pair of side cutters. The pair of cutters you get free with most stapling kits are usually terrible. A decent pair of flush-cut side cutters will do a much nicer job

3D scanning with your phone

Trap real objects in your own virtual world

Above
This two-colour print was done on the Prusa XL as part of the testing for this month's cover feature

Ben Everard

Ben spends more time making games controllers than he does playing games. It's a strange hobby, but he enjoys it.

We live in a 3D world, and are surrounded by thousands of 3D objects every day. We can use a bit of digital magic to turn these into 3D models. Basically, this involves taking a large number of photographs from every conceivable angle and running these through some clever software that creates a 3D model that we can then use just like any other 3D model. In this article, we will look at printing it out, but you can also use them as assets for games, or other 3D renders. This process is known as photogrammetry. There are some bits of software you can use to do the processing on a home computer. However, they typically require a GPU, and often an NVIDIA one. We don't happen to have one of these in the workshop, and we suspect that many of you don't either. A far easier option is to use an online service which will do the processing for you. We used the KIRI Engine which is available on the web or as an Android or iOS app. There are lots of options available, and they generally do one or more of the same basic processes.

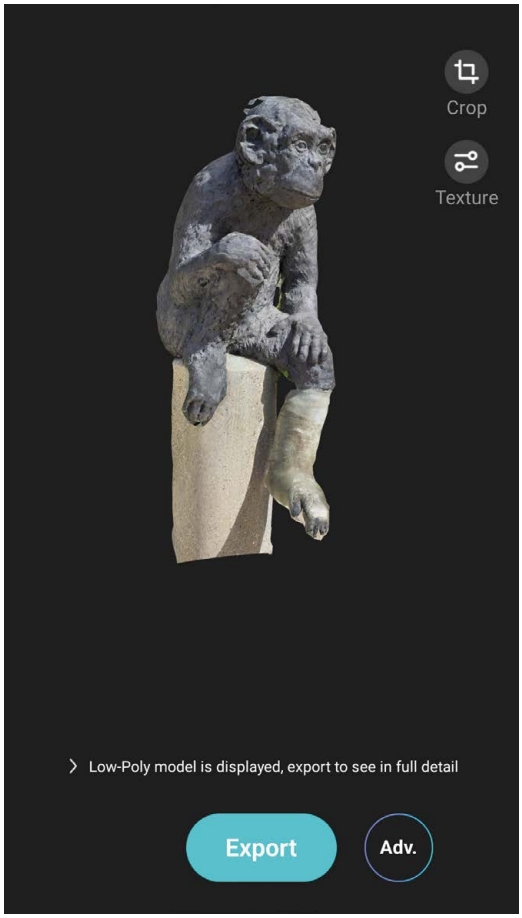

In KIRI, these are known as…

- **Basic Scan:** This is where you take a lot of photographs of an object, and the software attempts to build a 3D model from them. On KIRI, you can add 70 photos per project under the 'Basic' plan, which is free of charge, or 200 under the 'Pro' plan. Our experience was that 70 resulted in fairly poor scans, and you needed to go well over 100 to get good results. You can take the photos in the app either manually (i.e. pressing the shutter button each time) or automatically (where you move your camera around and it takes the photos as needed). Both worked well. Under the Pro plan, you can also upload photos from another camera – this may be particularly useful if using a drone, or a fancy camera.

- **Featureless Object Scan:** With a standard scan, features on the object are used to identify particular points. This works well for many objects, but if your object is a plain colour or has smooth sides, you may find you get better results with a featureless object scan.

Above
Like this lion, this author also has a broken wrist, and took this scan while waiting for a medical appointment

Left
This monkey (like the lion) is found outside Southmead Hospital in Bristol

RUNNING LOCALLY

Cloud services are a great way of getting up and running quickly with minimal hardware or software installation. However, sometimes you want the extra control you can only get from running your own software on your own hardware. There are a few options for this, ranging from the hobbyist to high-end industrial. Here are our picks to help you get started for free:

- RealityCapture. Currently free for businesses or individuals making less than $1m per year as well as all students and universities. It's a powerful piece of software that can render accurate 3D scans quickly.

- Meshroom is an open-source photogrammetry pipeline. It is perhaps not quite as powerful as RealityCapture, but still makes great scans.

Both of these require an NVIDIA graphics card. If you don't have one of these, you can get some of the power of running locally by using a Google Colab sheet, such as the one for Meshroom here: **hsmag.cc/Meshroom**.

- **Gaussian Splat:** Most scans start with a series of photos, but Gaussian splats instead start with a video. To be perfectly honest, we don't understand the maths behind this. Most of the time, Gaussian splats output a rather obscure file format that's a point cloud, and it pretty hard to work with. The KIRI Engine does allow you to output a mesh; however, you have to check the box at the start. You can't change this once you've already created your splat. If you use the in-app camera, it will tell you if you're moving too fast. However, you can also upload videos taken elsewhere. →

Right
You can crop the models in KIRI, but we prefer to do it in Meshmixer

Whichever option you use, you'll typically leave the object still and move the camera. If you want to move the object instead, you need to check 'Auto-mask object'.

If you're a pro user of KIRI (or, if your other service allows it), we'd recommend taking the photos or video outside of the app. While the app is useful, if you have an offline copy of files, you can edit them, remove some, add more, try another method, or anything else you want to do later. If you use the app, you only get one chance to get the photos right and you have to upload straight away.

PICK A METHOD

Gaussian splats are by far the quickest method (at least as far as taking the photos is concerned – processing the data can take a while). If you're on the go, it'll also use far less network bandwidth to upload. The results are really impressive: we found that while they may have missed a bit of detail, they had far fewer anomalies than the photo methods, and as such, often required little more than cropping before they could be printed.

The photo-based scans, on the other hand, were a more mixed affair. They were prone to spurious holes and appendages. The more photos you have (and the better quality the photos), the fewer issues you'll have. Our experience was that if you had enough high-quality photos, the results were better than Gaussian splats, but if you didn't, then they weren't.

TOUCHING UP

There are a few reasons you might want to scan something, and your motivations will inform what you do with the results. Many scans are for scientific or illustrative purposes. In these cases, you probably want to do as little as possible after the scan. If there are errors, you might want to add more photos, or start again entirely.

However, if you're 3D-scanning because you want an interesting scan, or a fun thing to 3D-print, you might be willing to take a heavy hand in editing the resulting model to fix errors, or make it more printable.

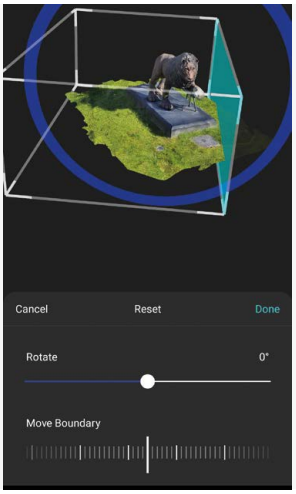

WHAT CAMERA?

You can get started with 3D scanning using a smartphone. However, a better camera will result in a better scan. Low light performance can be an important factor because many things you want to scan will have nooks and crannies that you have to try and photograph.

Image size is obviously important, but bear in mind that you'll be taking lots of photos, so beyond a certain point it can be tricky to manage the data and take a long time to process.

Some photogrammetry software can take RAW format photos, and you may get a better-quality scan with these if your camera supports it.

You can reach for 3D modelling software for this. Blender, for instance, has powerful tools for editing 3D meshes. At least, we believe it does – this author has never mastered this particular bit of software.

Meshmixer is a more straightforward tool for working with 3D models. It's quite limited, but the range of tools works reasonably well for tidying up 3D scans. However, there is a slight problem. Autodesk has bought the software, and stopped giving it updates. It still works, but at some point in the future, it'll probably stop working. For now, though, it's our preferred option.

Our method for touching up 3D scans in Meshmixer is…

- **Step 1:** crop out unwanted parts. The chances are that your scan took in more than just the object you wanted. Not only are these extra bits a distraction you don't need, but they can be a huge number of polygons that your 3D software needs to keep track of. Getting rid of them speeds everything up and generally makes it easier. To do this, go to Edit > Plane Cut. This tool lets you position a plane in the 3D space and discard all the polygons on one side of it. It might take quite a few plane cuts to get rid of all the bits you don't want. If you find that a plane cut isn't working, then you might want to consider the physical setup for the scan. However, there is also a sphere cut in KIRI that you can use if needed.

- **Step 2:** make solid. A 3D scan doesn't create a solid output, it creates a surface. Hopefully, this surface is reasonably complete, but there are probably at least a few holes in it. Before we can use it for 3D printing, we need to complete this mesh and turn it into a solid object. To do this, go to Edit > Make Solid. There are a few options here – we find that Meshmixer is a bit too conservative with the mesh size and so we generally increase it to allow more detail to be retained. Slide the Mesh Density slider up or down and click Update to change the preview. When you've found the right balance between detail and complexity, press Accept. You usually find that at this point the model turns grey. This is because the texture data is lost. In 3D terms, texture isn't texture at all. It makes no difference to the surface of the object – it's just 2D images that are mapped onto the object. This means that as far as 3D printing is concerned, it's irrelevant.

- **Step 3:** fix errors. The sculpting menu gives you some tools to fix any problems you have in the mesh. It takes a bit of practice to work out what tools can solve what problems, so we'd recommend allotting a bit of time to practise, and save a lot as you go along, as it's easy to get into a position where you've almost got it how you want it, then suddenly getting it all wrong.

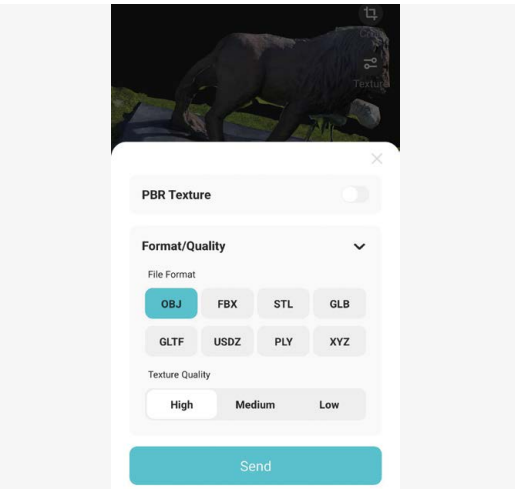

Left
Make sure you export in a format you can import into your slicer – we chose OBJ

SLICING

Once you're happy with the results, it's time to bring it into your slicing software. We use PrusaSlicer for this, as it makes it easy to fix the last few bits. If your scan includes any regular shapes – such as cubes or cylinders – you can replace them with imported versions of the same shapes that have sharper lines.

Highlight your model and right-click, then select Add Part, and then the appropriate shape. You can use the usual move and scale tools to make the part the right size.

You'll probably find that the scan isn't a particularly easy shape for 3D printing. Most real objects have thin bits, overhangs, and all manner of things we're told to avoid in our 3D-printable designs. Fortunately, modern 3D printers and slicers are much better than they used to be, so we don't need to worry too much about this. We have found that our 3D scans print a bit better if we use the automatic support painting feature of PrusaSlicer. To use this, highlight your model, then click on the paintbrush icon (or press **L**), and press Automatic Support Painting. In the Supports section on the right-hand side, select For Support Enforcers Only. This should ensure that enough bits are supported, but there's not too much filament wasted, and it's easy enough to remove the supports from the finished print.

Once you're happy, slice as usual, and send the file to your printer. Photogrammetry is easy to get started with and you can create some really interesting parts for 3D printing.

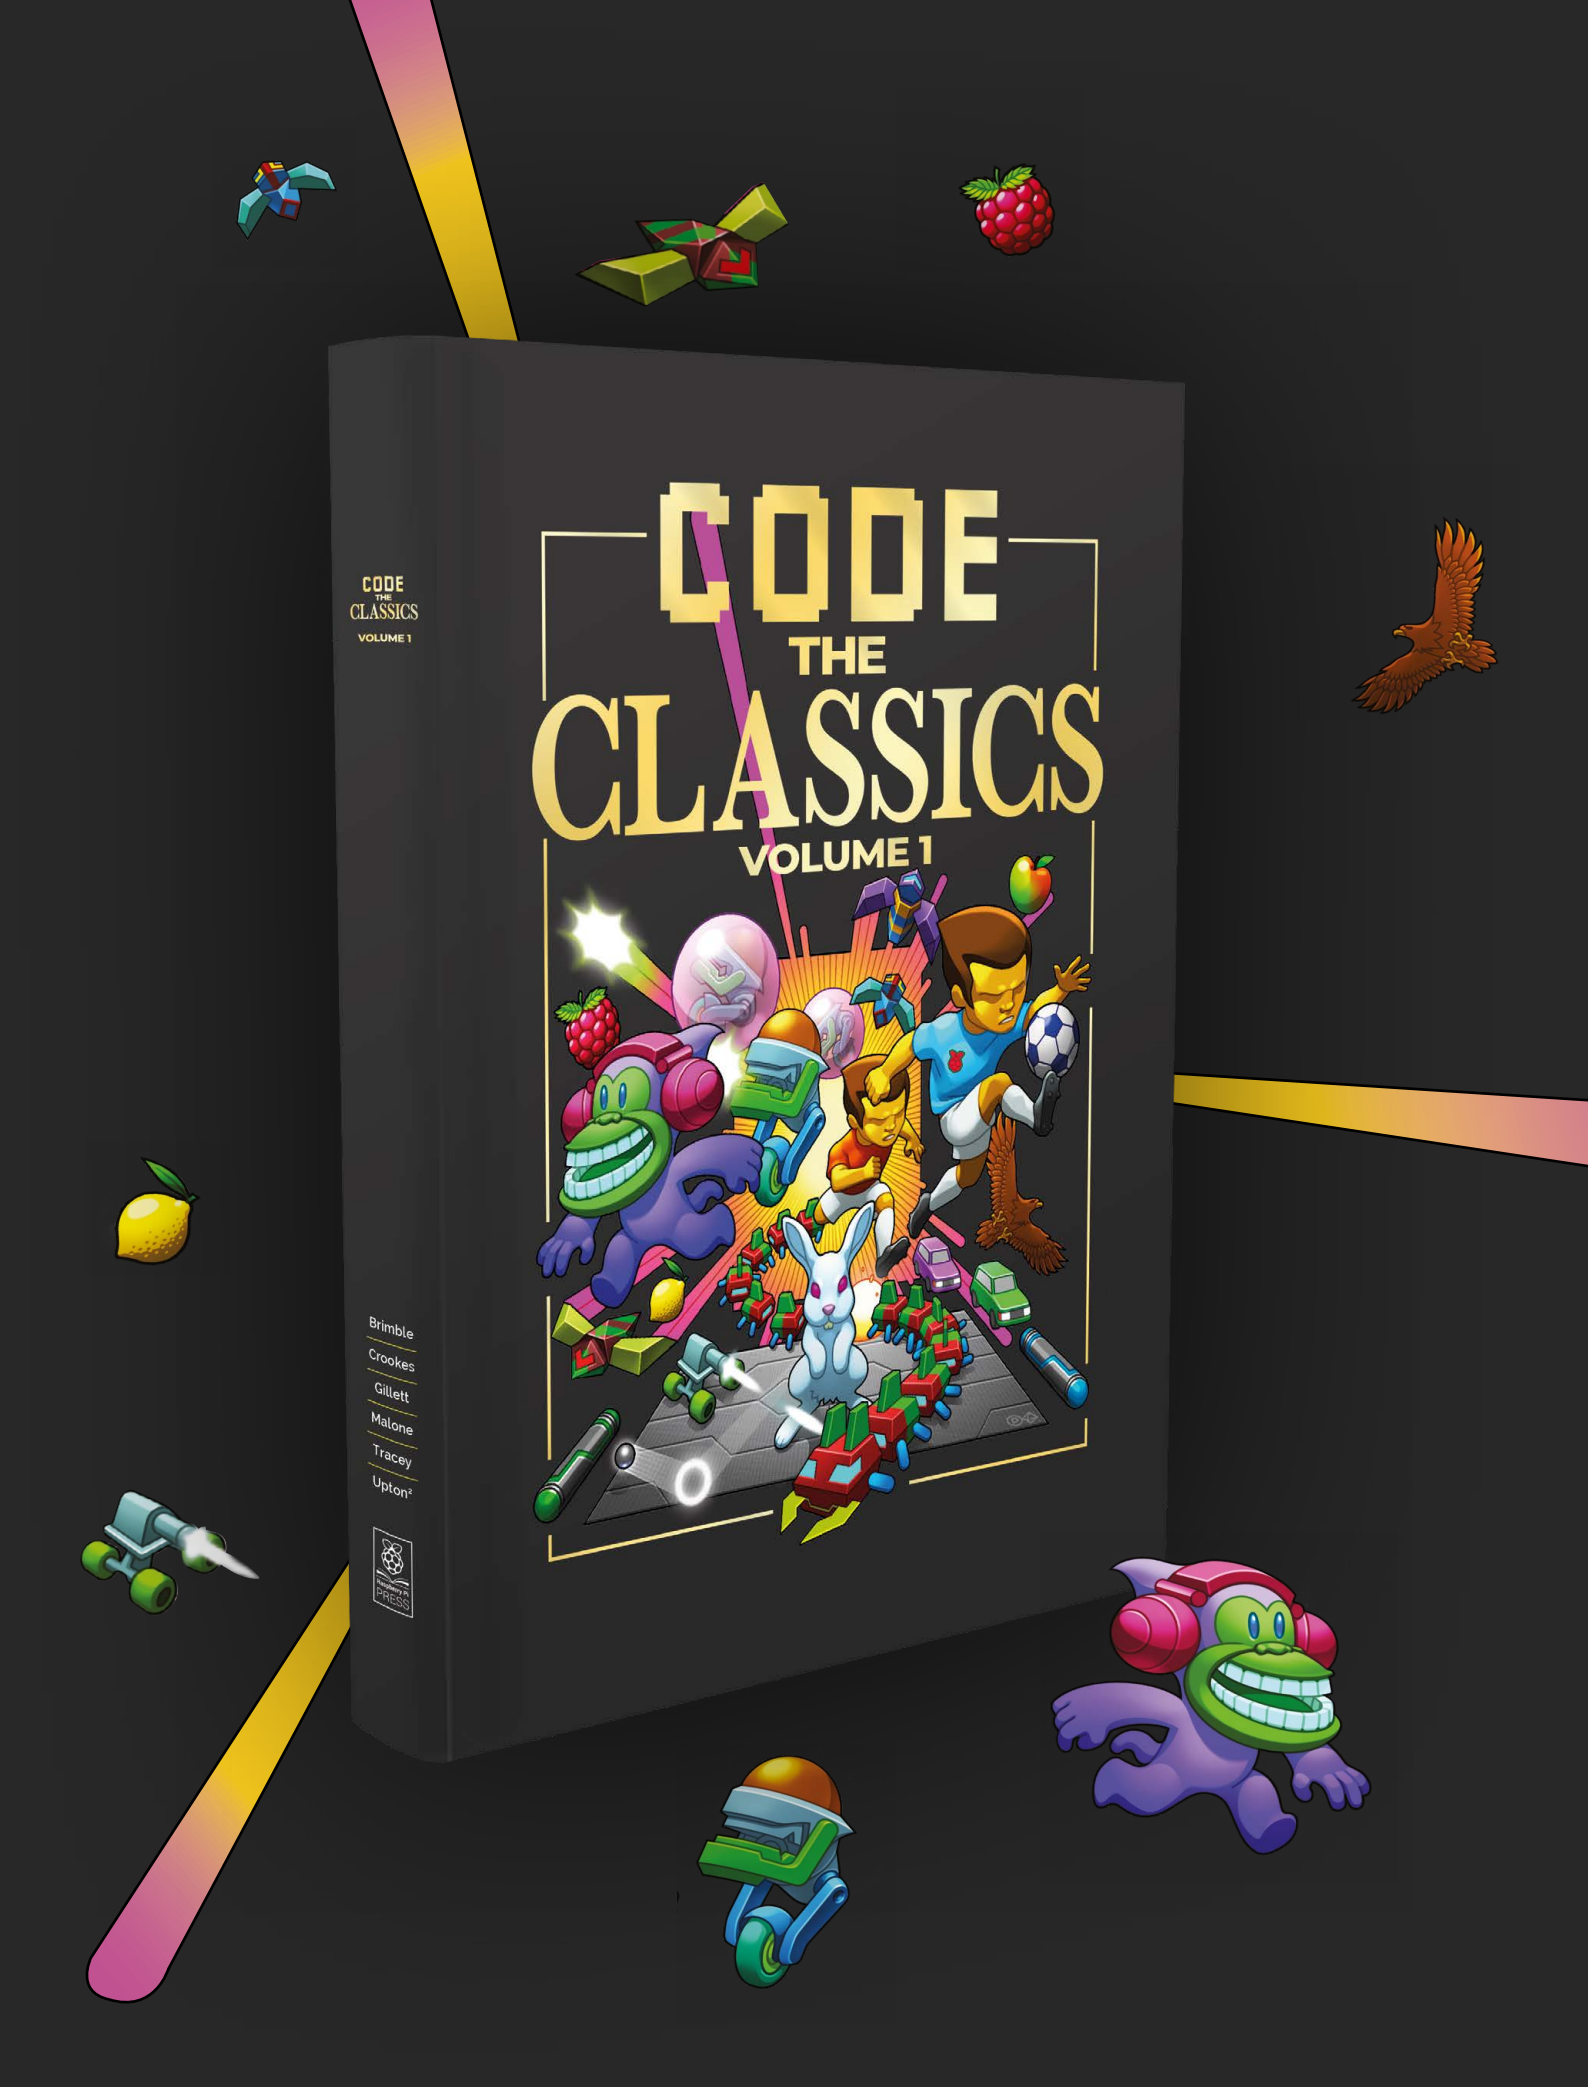
CODE
THE
CLASSICS
VOLUME 1
Brimble
Crookes
Gillett
Malone
Tracey
Upton²
Raspberry Pi
PRESS

CODE THE CLASSICS VOLUME 1

This stunning 224-page hardback book not only tells the stories of some of the seminal video games of the 1970s and 1980s, but shows you how to create your own games inspired by them using Python and Pygame Zero, following examples programmed by Raspberry Pi founder Eben Upton.

- ***Get game design tips and tricks from the masters***
- *Explore the code listing and find out how they work*
- ***Download and play game examples by Eben Upton***
- *Learn how to code your own games with Pygame Zero*

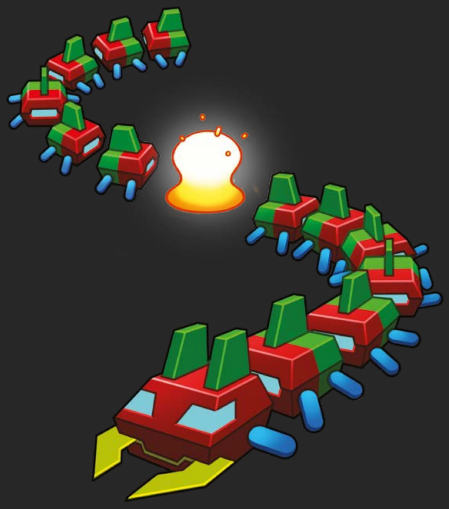

Available now **hsmag.cc/store**

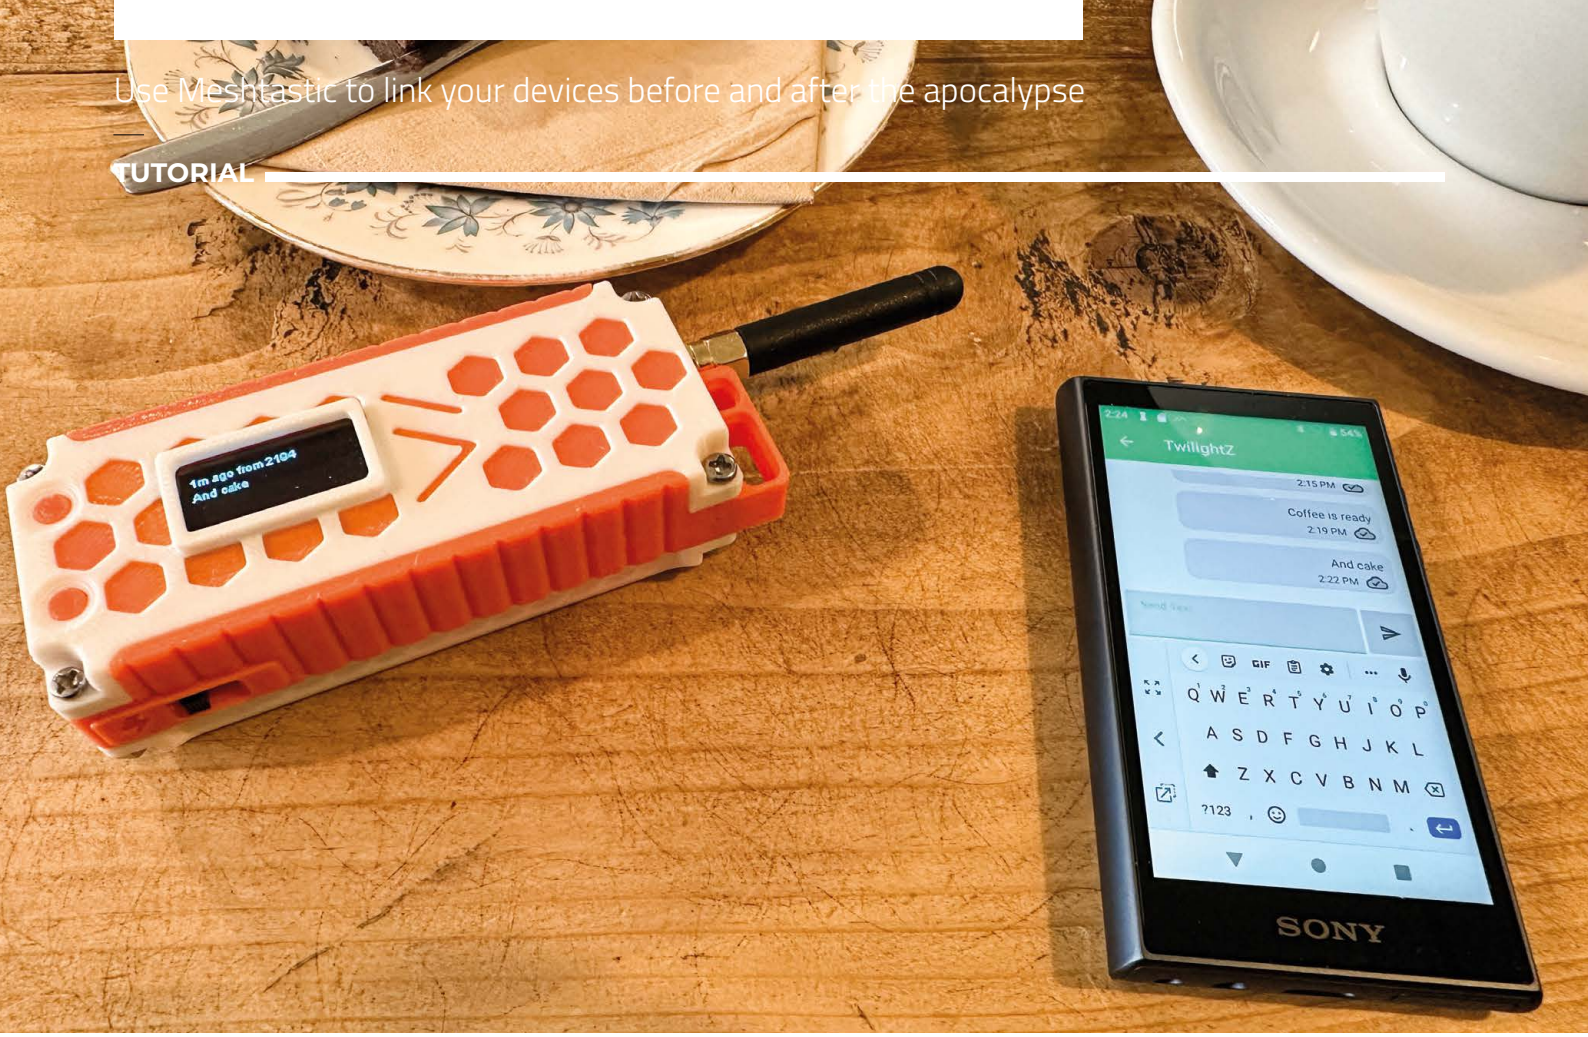

Use Meshtastic to link your devices before and after the apocalypse

Discover how to use Meshtastic to connect your devices together

Rob Miles

Rob has been playing with hardware and software since almost before there was hardware and software. Find out more at **robmiles.com**.

‘Off-grid’ networks like Meshtastic are becoming fashionable. They provide free connectivity (once you’ve bought the devices), and can be used in regions where the cellular phone signal is poor or non-existent – such as distant Alaskan lakes or a market town near where the author lives. Meshtastic devices could also be used in ‘post-apocalyptic’ scenarios, although perhaps keeping up with gossip might not be the first thing on your mind after the end of the world. We’re going to look at how Meshtastic works, how you can connect to it, and what you can use it for.

MESHTASTIC

Meshtastic (**meshtastic.org**) is an open-source, community-driven project. It provides software for radios (which connect everything together), and client devices (which connect to the radios to send and receive data). There is almost no network configuration required beyond giving your radio a name and deciding who you want to talk to with it. The radios use ‘LoRa’ technology to move the data around.

Figure 1 shows Meshtastic in action. One user has entered an important message into the chat window on their mobile device and this has been received on

Figure 1
The device on the left is a Meshtastic radio based on a Heltec V3 in a rather nice 3D-printed case. The device on the right is a music player running the Meshtastic application on Android. The player is connected over Bluetooth to another Meshtastic radio which is not shown

Figure 2
A radio can be 'headless' or have devices connected to it

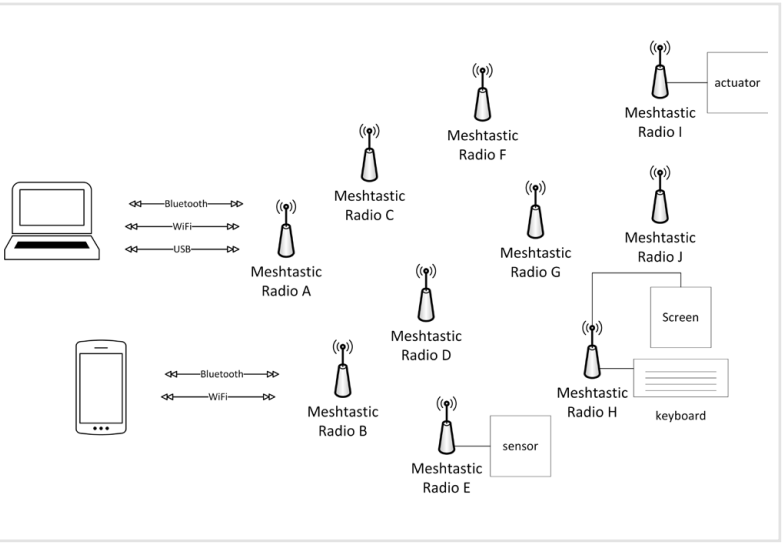

STAYING SECURE

Meshtastic messages are encrypted using AES-256 encryption. This uses 256-bit keys to encrypt the data before it is transmitted by the radio. The key is combined with the content of a message in such a way as to make it very hard to decode the message if you don't have the key. This is important because anyone can receive LoRa messages sent from your radio. If you create a Meshtastic channel, you will also create a key for that channel, so only those who have joined the channel can send and receive the messages. Keys can be shared between radios using QR codes.

If I want to let you join my channel, I'll use the Meshtastic app on my phone to display a QR code which you can scan with the Meshtastic app on your phone. Your phone then transfers the key information to the Meshtastic radio connected to your phone, which can then send and receive messages on that channel.

the Meshtastic radio on the left. The message was sent to a channel (TwilightZ) and will be received by other devices that have joined that channel. You can use Meshtastic for one-to-one chats, and the radio devices can also control sensors and actuators.

MAKING A MESH OF THINGS

Figure 2 shows how the Meshtastic network is structured. The Meshtastic radios cooperate to deliver messages over the network. There are no fixed network routes or configuration: messages are forwarded across according to routing information which is held in each radio and updated as the configuration changes. Radio A is connected to a computer, and Radio B is connected to a phone.

When a radio transmits a message, this will be received by radios in range that will then retransmit it. A message sent from Radio A destined for Radio I in **Figure 2** might be picked up by Radios B, C, and D. They then retransmit the message. The message from Radio C might reach Radio F which will then retransmit it again. Finally, the message will be picked up by Radio I. Each retransmission of a message is called a 'hop'. A message contains a 'hop counter'. If the message hasn't arrived after a given number of hops, it will be discarded.

When a radio joins the network, it sends out a broadcast which is used by nearby stations to update their routing tables based on the strength of the signal they have received from the new radio. A radio won't retransmit a message if it decides it is not on a particular route. For example, Radio E in **Figure 2** might not retransmit a message addressed to Radio I. →

YOU'LL NEED

- **A device containing a LoRa radio**
  There are Heltec and T-Beam devices available or you can connect a LoRa radio to a Raspberry Pi Pico
- **A mobile phone or computer**
  with a Bluetooth connection that you can use to connect to the radio

CHANNEL CROSSING

The LoRa standard divides the radio frequency range into different radio channels. A Meshtastic channel is usually assigned to a particular LoRa channel, with one channel reserved for broadcasts to all nodes. The LoRa hardware in a radio can't listen to more than one channel at a time. Instead, it must listen to each channel in turn. This means that messages need to be transmitted repeatedly to improve the chances of them being received. This, coupled with the limited data rate provided by the LoRa radio system that underpins Meshtastic, means that devices must be careful about their use of the network. You can configure how your Meshtastic installation uses the channels.

Figure 4
You can add a circuit to turn off the GPS when it is not required

Figure 3
You usually have to solder the display unit onto the T-Beam, although they are sold together

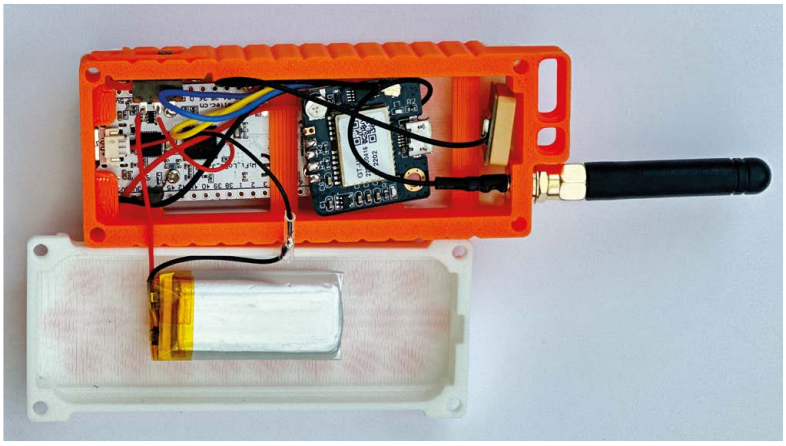

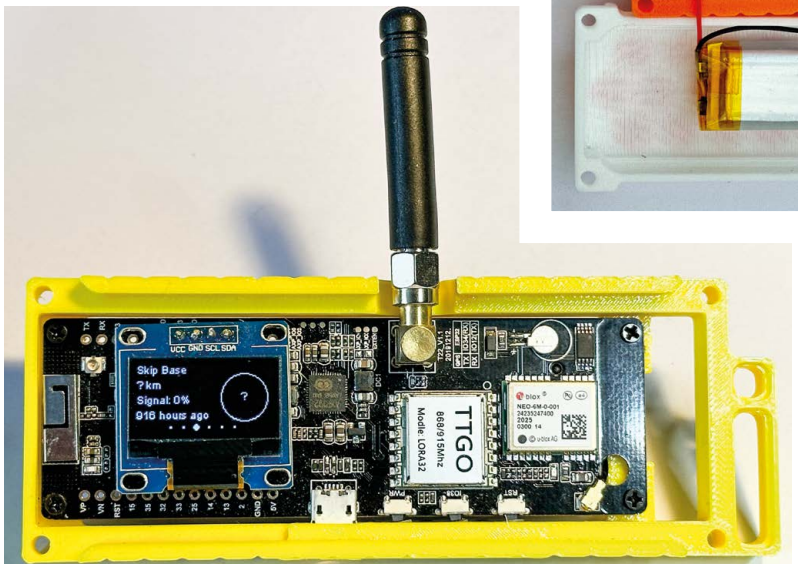

Figure 3 shows a T-Beam device mounted in a 3D-printed case. You can see the TTGO LoRa radio chip underneath the aerial with the GPS receiver on the right. The LoRa radio is marked with the LoRa frequencies for use in Europe. Make sure you get the right device for your region. You can find more details at **hsmag.cc/MeshRadioSettings**.

Figure 4 shows the configuration of a Meshtastic radio based on the Heltec V3 device. A GPS receiver and antenna have been added to make the radio location-aware. You don't need to add the GPS, but it does make the radio more useful when you are out and about.

QUICK TIP

When using Meshtastic, the radio holds all the details of radios and channels that it is connected to. The client running in the mobile phone or computer connected to the radio is just viewing this information.

All network designs are compromises, and Meshtastic is no exception. You can read all about the design at **hsmag.cc/MeshBroadcastAlgorithm**. There is even a network simulator you can use to experiment with different configurations.

HANDY HARDWARE

To get started, you'll need a device which can send and receive LoRa messages. Popular ones include the Heltec V3 and the LILYGO TTGO T-Beam. The T-Beam provides the most complete solution with an on-board GPS receiver and space for a large battery on the back. The devices section of the Meshtastic site (**hsmag.cc/MeshSupportedDevices**) provides descriptions of each device along with suggestions for suppliers.

SOFTWARE IN A FLASH

Once you have your device, the next step is to load it with the software that will turn it into a Meshtastic radio. You can do this from the Meshtastic website (**hsmag.cc/MeshGettingStarted**) which also takes you through the process of installing serial port drivers on your computer and then on to the online device flashing tool to transfer the software into your device.

Figure 5 shows the Meshtastic flash page. You select your device and the version of the firmware you want, and it does the rest. You need to use a Chrome-based browser, such as Chrome or Edge, for this to work, as only these browsers

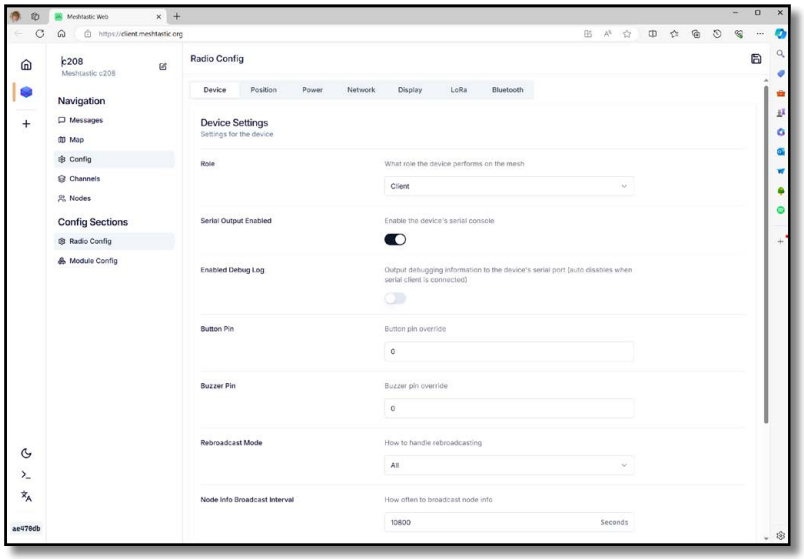

Figure 6
The radio is being set to the role of Client. It can be set to Router if it is just going to be a part of the mesh

Figure 5
You can also download the Meshtastic source code from GitHub and build your own images if you wish

CALL OF DUTY

Meshtastic provides 'duty cycle enforcement' to limit the amount of time a radio can transmit on the network. This is usually set at the level of 1%, which means that a station can only transmit for 1% of the time; for example, it can only transmit for one second every hundred seconds. This is not a huge restriction bearing in mind the intended uses of the system, namely text messages and small commands/data items.

support direct access to the serial port of the host computer.

COMFORTABLE CONFIGURATION

Once you have flashed the software onto your device, you can configure it using a browser or the Meshtastic application. The author has found that the easiest way to configure a radio is to use the browser-based client application which can connect to a device over USB, Bluetooth, or Wi-Fi. You can do this over USB straight after the device has been flashed and while it is still plugged into your PC or laptop.

Figure 6 shows a Meshtastic device being configured. The default settings are best left as is. You can give a device a more sensible name than the original one, though. The author hasn't got around to doing that yet, as you can see in **Figure 6**. If your Meshtastic radio has Wi-Fi (which most do), you can use the Network tab in the configuration window shown in **Figure 6** to connect a radio to your local network. You can then interact with the radio from any browser connected to your network.

Figure 7
The display shows how many hops away some devices are

GET TALKING

Now that we have Meshtastic working, we can use it to communicate. You can use the browser client to interact with a Meshtastic radio, but it is much more useful to connect your mobile device to it. There are Meshtastic applications for Apple and Android devices which you can download and install from the appropriate app store. These use Bluetooth to talk to the radio.

Figure 7 shows the Meshtastic app running on an iPhone and connected to a Meshtastic radio. It is showing all the local nodes. We can send a direct message to any of these devices by selecting the Messages tab at the bottom of the screen. The Meshtastic app can be configured to display incoming messages as alerts when they arrive. →

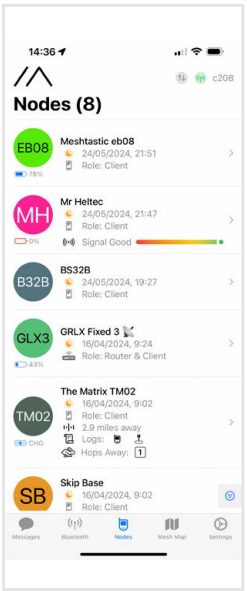

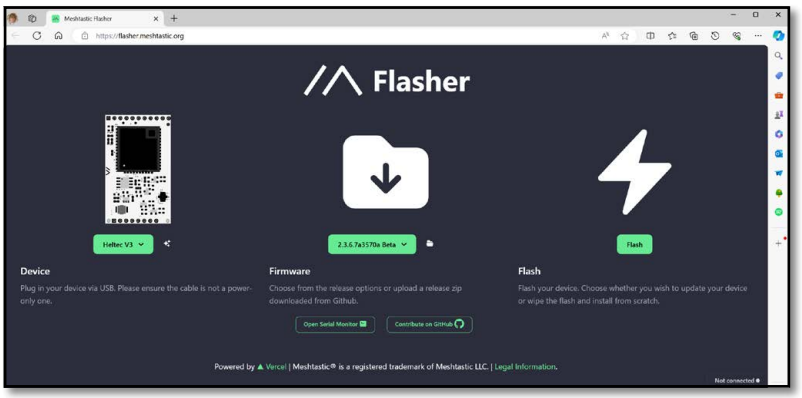

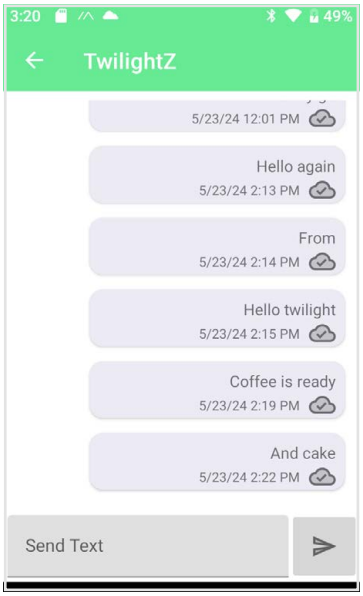

Figure 8 This conversation would be shown on all devices which are connected to the channel

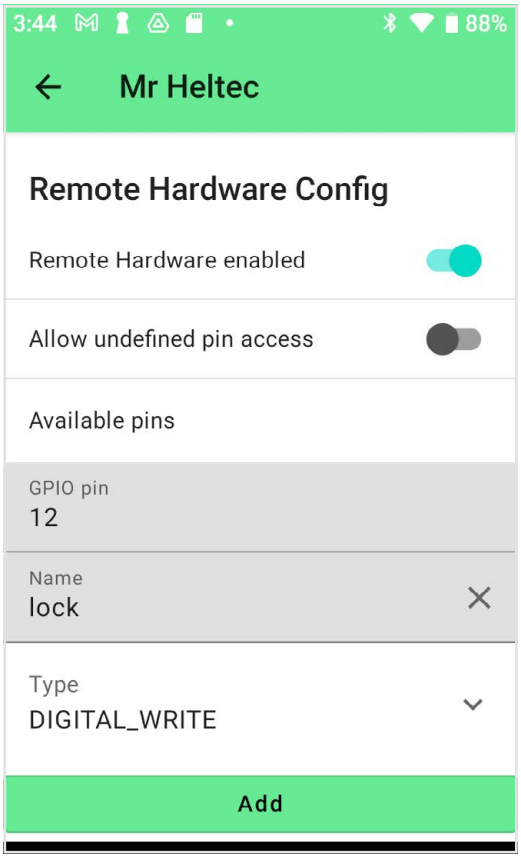

Figure 9 We could connect a lock control to GPIO 12 on 'Mr Heltec'

QUICK TIP

The Meshtastic applications work well most of the time, but they are very much 'works in progress', and the author has experienced a few crashes on both Android and iPhone.

TELL LORA I LOVE HER

The wireless technology that underpins Meshtastic is based on a radio technology called LoRa. The name stands for 'Long Range', although it can also stand for 'Low Rate'. LoRa can use a very low-power radio to send a data signal a long distance, and it does this by sending data encoded as 'chirps' rather than fixed signals.

A chirp starts at a low frequency and then moves to a higher frequency over a particular time. You can create audio 'chirps' of your own using a swanee/slide whistle or a trombone. The LoRa transmitter uses a computer-controlled radio to generate a chirp that represents the data to be sent.

A LoRa receiver decodes the data by mixing the incoming 'chirp' signal with a locally generated 'counter-chirp' that sweeps in the opposite direction. If the received chirp moves from low to high frequency, the local counter-chirp moves from high to low frequency. This combination produces a specific frequency, which corresponds to a particular data value. You can set the rate at which the chip frequency changes. Longer chirps take more time to send, but they are less affected by noise.

LoRa signals can be sent at very low power because the 'chirp' encoding is highly resistant to interference. The bad news is that sending data is slow. It makes LoRa perfect for text messages and small amounts of sensor data, but it can't be used for large files, voice, or video.

QUICK TIP

LoRa devices are usually supplied with very poor antennas. You can greatly improve the range of your device by purchasing a better one.

Figure 8 shows a highly advanced conversation taking place on the TwilightZ channel. Radios will buffer several incoming messages and then send them to a client when that client connects.

MESHTASTIC MODULES

Figure 9 shows the process of using the 'Remote Hardware Config' in the Meshtastic app. You can use this to assign a general-purpose input/output (GPIO) pin to a Meshtastic radio so that the pin can be controlled by Meshtastic messages. In the case of the above configuration, the plan is to use the radio to control the lock on a cabinet so that we can control access to something via Meshtastic messages.

There are lots of other modules available for use in a Meshtastic radio. These include environmental sensors, MQTT, ambient lighting, serial data, and a Paxcounter which detects Wi-Fi and Bluetooth devices to count the number of people passing by the radio.

If you want to go beyond modules and connect a Meshtastic radio to a novel piece of hardware, the source code for all the Meshtastic radios is available on GitHub. It's built using PlatformIO, so it's very easy to install it, then create custom radio software.

TALK TO MESHTASTIC FROM YOUR APPS

We have seen how we can use client applications for the browser, the iPhone, and Android devices to connect to Meshtastic radios to send and receive messages. There are also libraries which you can add to your applications which can connect to a Meshtastic radio. This means you can make your application use Meshtastic as a network transport.

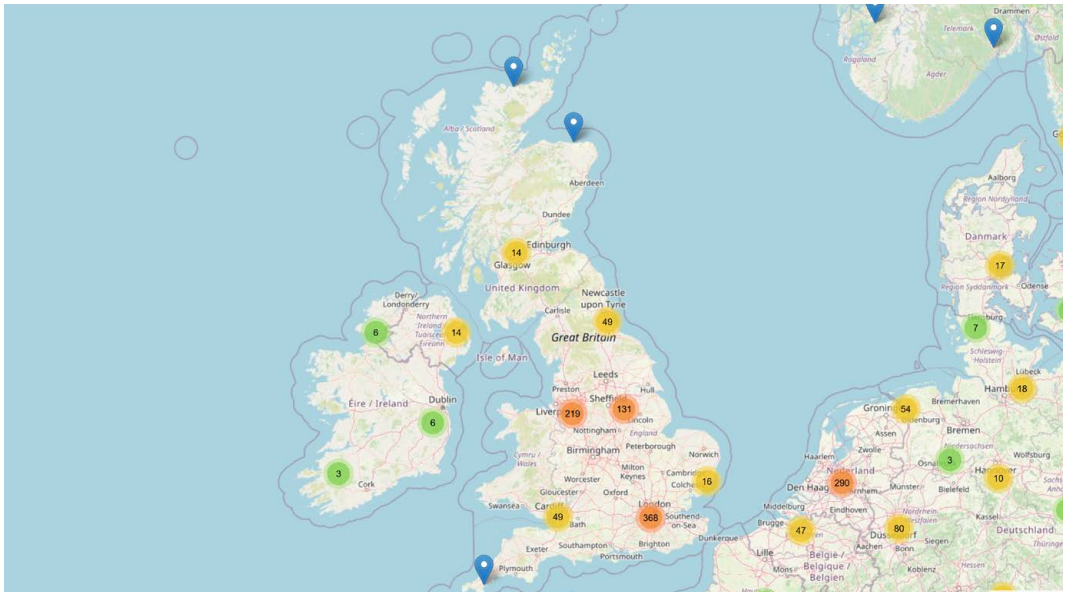

> **The LoRa network may be slow, but you can get a** surprising amount of range

Above
There are Meshtastic nodes all over the world. Take a look at meshmap.net to find ones near you

MESHTASTIC IS FANTASTIC

The great thing about Meshtastic is that you don't have to do much work configuring your devices and getting started. The author had some old Heltec V2 devices which turned out to work quite well (although the newer V3 ones have much better radio performance and battery life). If you have some LoRa devices lying around, you might like to pop the Meshtastic firmware onto them and have a play. You can start just by giving your radios names, creating some channels, and that's it. And once you have your devices communicating, you can add features very easily.

The LoRa network may be slow, but you can get a surprising amount of range from these simple radios. The Meshtastic application provides a mapping function which shows you radios in your area, and these can be quite a distance away, particularly if you can get some radios nice and high up with good 'line of sight' coverage of an area.

Meshtastic will never replace the internet, but it does provide a very neat way of taking your networking a long way further, particularly if you are lucky enough to live somewhere with an existing community of Meshtastic devices which can forward your messages. And if you haven't got such a community, it is certainly worth building one.

LORAWAN VS MESHTASTIC – FIGHT!

You might have heard of the LoRa radio protocol before in the context of LoRaWAN (LoRa Wide Area Network). Both Meshtastic and LoRaWAN use LoRa to move data around, but their structure and how they are used are quite different.

In a LoRaWAN network, the LoRa radios connect to LoRaWAN gateways. A gateway contains a LoRa radio and an internet connection which it uses to forward messages to LoRaWAN servers. Applications that wish to use LoRaWAN to communicate with remote devices connect to the LoRaWAN servers to send and receive messages. LoRaWAN works very well if you want to gather information from a huge number of remote sensors, but it is not very useful for bidirectional conversations and it is very difficult for one remote device to talk directly to another. Meshtastic makes it very easy to connect a number of devices together, but doesn't require any infrastructure other than the radios themselves.

Both technologies have their place, and LoRa radio devices work with either. The author has used Heltec devices connected to LoRaWAN to collect data from air quality sensors and transmit it over the Things Network LoRaWAN (**thethingsnetwork.org**), and he uses Meshtastic to keep in touch and connect to sensors around the home.

QUICK TIP

Switching on a LoRa device without an antenna connected can break the radio transmitter on the device.

TUTORIAL

Build an electric guitar

PART 01

Can you craft a high-quality instrument using only tools from the middle aisle of a German supermarket? Let's find out!

Andrew Gregory

Everything in Andrew's garage is covered in a thin film of orange dust from working with wood.

Phew. The time has come to write up our electric guitar build. It's been a rollercoaster of mistakes and inefficiencies, but somehow – mainly because we took it slowly and carefully – we've not made too many mistakes. Of course, if we were to ever make another one, the mistakes are what we'd pay the closest attention to, so we wouldn't make them again. And that's what we're doing here. Guitar design is a complicated subject full of myth and magic, but we've tried to simplify things, so if you give it a go yourself, you'll find loads of different avenues to explore.

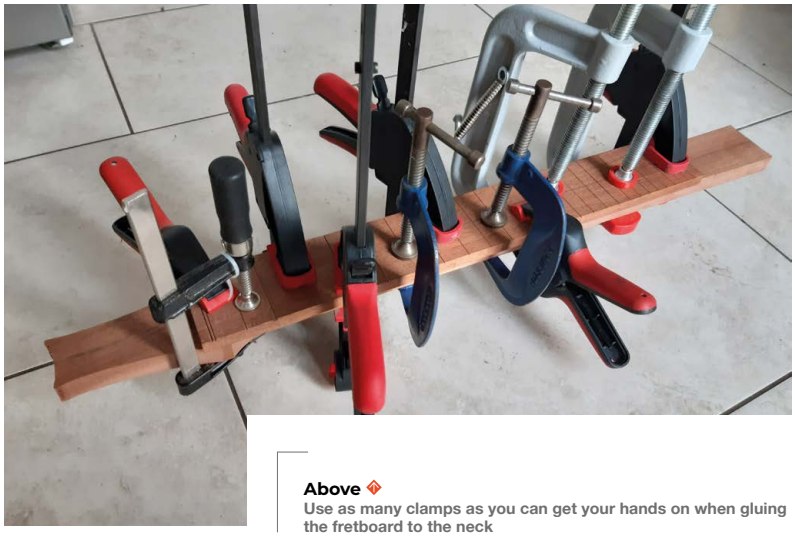

Above
Use as many clamps as you can get your hands on when gluing the fretboard to the neck

WOOD

The first thing to do is buy wood. If you're lucky enough to have a local timber merchant nearby, there will most likely be something affordable there for you. The most important thing is that it needs to be a hardwood; softwood, such as pine, is cheaper, but unless you have a way of keeping your tools razor sharp, they'll just produce cuts with fuzzy edges. It's counter-intuitive, but hardwood is easier to work with. We used a 20 mm thick × 70 mm wide × 900 mm length of sapele, aka African mahogany.

The other thing that you need to ensure is that the wood you get is square and planed on all four sides, so you've got known right angles to work with. This is important because as soon as things aren't perfectly flat they get harder to measure, harder to clamp, and harder to cut accurately. Keep everything square for as long as possible. We measured a guitar we had already, and bought too much wood, because we knew we'd inevitably make mistakes.

You'll also need MDF for templates. Proper craftsmen might sneer, but we used an MDF template as a guide to cut out the guitar neck with a hand router. We drew the shape of the guitar neck that we wanted on a piece of MDF, cut it out with a jigsaw, and sanded it until the important bits were the right dimensions. We got away with this approach for the straight bit of the neck, but the shape of the headstock is slightly wonky, because it's hard to cut accurately with a jigsaw. A bandsaw would have helped here, or even better, a laser cutter – a computer will help you draw a symmetrical

Above
Sanding a radius onto a flat piece of wood takes a little while, and generates a load of dust. Wear a mask!

shape far more easily than a Biro, and laser cutters will give you better accuracy than a jigsaw bought for £29.99.

We fixed the MDF template in place to the neck using double-sided sticky tape, then ran a router around it using a flush cut bit. The router bit follows the outline of the MDF, leaving you with an exact copy in wood; it's a slow process, and if we had access to one, we'd have used a bandsaw to cut away most of the wood, leaving the last 2 mm or so for the router. As it was, we ended up cutting away almost everything by making multiple passes with the router, which took ages.

With the neck blank cut out, we had to use our router to carve a slot in it for a truss-rod – this is a metal bar that strengthens the neck, and allows for some adjustment when you tighten or loosen the nut at its end, which is accessible through a hole drilled straight down into the top of the neck from the headstock; you then glue a separate, thinner piece of wood on top of this – the fretboard.

We built a jig to help us keep the router moving in a straight line and used a 6 mm plunge router bit to cut a channel that would fit the truss-rod. Perhaps we didn't clamp this perfectly accurately, because the perfectly straight line is actually out by around 3 mm. This means we would have to drill 3 mm to the left when we drilled the access hole for the truss-rod, and hope that the two channels meet up.

DON'T FRET

The distance between the nut and the bridge on any guitar is known as the scale length, and our planned guitar is 600 mm, which is near enough 23.75 inches. Half of 600 mm is 300 mm, which is where we need to put our 12th fret. Half again is 150 mm, which is where the 24th fret would go. Halving the length of a vibrating string doubles its frequency – in other words, it's an octave difference. We started out relishing the challenge of working from this first principle and calculating the distances between the rest of the frets, but then found an online calculator that would do it all for us: **stewmac.com/fret-calculator**.

With a scale length of 600 mm, the distance between frets 1 and 2 should be 31.785 mm; between 23 and 24 should be 8.919 mm, so keep your pencil sharp and cut carefully. Better yet, put the numbers into a laser cutter and use that to mark the fret positions.

The little metal lines that run across the width of a guitar neck are called frets, and you can buy them online. They comprise a thick bit that comes into contact with the string, and a thin, barbed bit that goes into a slot cut into the fretboard. The slots that the frets go into have to be perfectly parallel all the way up the neck; so, as always, we'll keep the fretboard timber square, cut the fret slots, then glue it to the neck and shape it. We found a saw with a 0.5 mm kerf online (the tool costs really did mount up, but we've used everything at least once since then, so don't feel like we've wasted any money) to match the width of the fret tangs. →

Left
These dents in the fretboard were made by the impatient craftsman bashing the frets in with a mallet. Next time, we'll use an arbour press, or a clamp and a rounded piece of wood to press them in

Below
We fixed the MDF template for our guitar neck in place with double-sided sticky tape, which proved to be stronger than the MDF when we tried to remove it

> **To achieve the curve from the headstock up to the nut, we bought a** hemispherical router bit and used it to carve a nice curve

ANOTHER MISTAKE

We cut the fret slots (see box, Don't Fret), then glued the fretboard to the neck, which we left a little longer than we needed to give us room for adjustments (aka room to get things wrong). At this point, the fretboard was still square – it's at this stage that we should have measured a line down the centre of it and drilled the holes, 1.5mm deep, to insert our fret marker dots. Instead, we did this after we'd used a handsaw and router to trim the fretboard flush to the neck, which made measuring harder than it needed to be – so hard, in fact, that the dots we put in are just wonky enough to be annoying, even if you can't notice them unless you're looking closely.

We next had to decide on the fretboard radius. The fretboard on an electric guitar isn't completely flat, but has a slight curve to it, which you can describe as a section of a circle with a known radius. A larger radius will give you a flatter feel; a smaller radius, such as the 7.25-inch radius found on vintage Fender Telecasters, is more comfortable for wrapping your hand around, but less useful when it comes to playing fast guitar solos. We opted for a compromise of 9.5 inches (annoyingly, the guitar world still uses imperial measurements most of the time). We bought a pre-made sanding block with a 9.5-inch radius, then set to sanding. We'd marked the positions of the fret marker dots by scratching lines onto the surface of the wood, so we'd know when we'd got a consistent radius when all these marks were sanded out.

Next, we needed to reduce the thickness of the headstock so that we could drill through it and add the tuning pegs. This called for another jig to be built, this time to restrict movement of the router and enable us to shave off only a few millimetres of wood at a time. To achieve the curve from the headstock up to the nut, we bought a hemispherical router bit and used it to carve a nice curve – in a proper workshop we'd have done a better job with a spindle sander, but we managed. We drilled into the curve we'd just made to give access to the truss-rod, remembering to move the hole 3mm off-centre, and miraculously found where we needed to be. With a CNC machine, this error would have been easily avoided.

Now we had a neck with a shaped form, but a square back. It was time to break out the most Neanderthal of tools, the rough file. We used this to turn beautiful, pristine African mahogany into sawdust, at first gingerly knocking off the corners all along the length of the neck, then more confidently. If we'd got this wrong we would have wasted all the previous effort, but somehow we got this bit exactly right, carving a comfortable, easy-to-grip profile all down the neck, even making a slight taper toward the top. We have no idea how we did this. We spent hours carving the neck by hand like this, moving to a smoother file, then to sandpaper to get the marks out and leave a smooth finish. We enjoyed doing this, but if we'd used a CNC machine, we'd have enjoyed it even more, and been able to have more control over the finished shape.

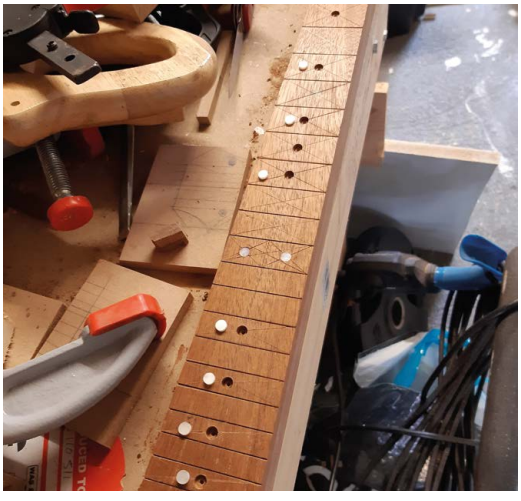

Right
We used a hand drill to create holes for the dot markers – this job would have been far more accurate with a pillar drill. Note also that we were working perched on the edge of our bench, a suboptimal work situation

Right
We'll come back next month for the rest of the guitar, but for now we're pretty pleased with how things are going

Below
We've no idea how we got this so right. We do know that it took ages, made loads of mess, and you'd be better off using a CNC machine if you have access to one

The shaping process had sanded away the outer two-thirds of each of the fret slots, so we went over them again and pushed the frets in. These were pre-cut to length, but we still had to trim them with flush cutters, then file the ends to a round, smooth finish – definitely a process that gets easier with practice.

A bit of red dye, a coat of Danish oil, and we've got a guitar neck. Believe it or not, this is the hard part. With an electric guitar, it doesn't matter what shape the body is: you can use a shovel, a toilet seat, anything you want as long as it's got room for the electronics. We'll come back to that next time, but for now we're delighted that we could make so many mistakes while still coming up with something that looks more or less right.

ERRORS

We've learned many things through trial and error:

1. Any mistakes in the MDF template will be replicated in the wood, so make your mistakes in MDF first, and keep at it until you get it right; fixing it later will only cost you time and money.
2. Measure and mark the centreline of everything.
3. Use Inkscape and a laser cutter to design your templates. Even if you're a great designer, a jigsaw just isn't accurate enough to cut the shape you want.
4. Clean up after yourself, and use an uncluttered workbench. Wear a mask when you're sanding, and eye protection when you're using a router. No, your normal eyeglasses are not enough. Dry wood chips will stick to your wet eyeballs.
5. Every time you cut into your guitar, make a practice cut on a bit of scrap wood first – then you'll know how it's supposed to feel.
6. If something doesn't fit, don't just bash it!
7. Keep things square for as long as you can – it helps.
8. Use scrap wood to protect your work from dirty clamps.
9. A blunt pencil can easily be half a millimetre thick or more, so use a sharp pencil or marking knife for greater accuracy.
10. Chisels only work in the direction of the wood grain. If you try to use a chisel to cut across the wood grain, you'll end up making a mess.

DON'T MISS THE **BRAND NEW** ISSUE!

SUBSCRIBE FOR JUST £10!

- **THREE!** issues of The MagPi
- **FREE!** Raspberry Pi Pico W
- **FREE!** delivery to your door

+ FREE RASPBERRY PI PICO W*

Three issues and free Pico W for £10 is a UK-only offer. Free Pico W is included with a 12-month subscription in USA, Europe, and Rest of World. Not included with renewals. Offer subject to change or withdrawal at any time.

magpi.cc/subscribe

HackSpace
TECHNOLOGY IN YOUR HANDS

FIELD TEST

HACK | MAKE | BUILD | CREATE

Hacker gear poked, prodded, taken apart, and investigated

ONLY THE BEST

DIY photography and videography

A collection of products related to taking pictures and videos

By **Marc de Vinck**

I've always been interested in photography, and I've played around with a few DIY photography projects, including building an intervalometer from an A555 timer, integrating an old webcam into an even older camera body just for fun, and an outdoor camera for capturing the critters that were roaming around my yard at night. I've also covered cameras in a previous Best of Breed, but there are a lot of products out there, so it warrants a second look. Also, the last time I looked at the various photography-related electrons, things like the AI-powered Arducam PiNSIGHT weren't even in existence. Things really are developing fast! (Pun intended.)

In this Best of Breed, I'll be taking a look at some photography-related products that you can integrate into your next project. Some of these products are cutting-edge, while others are fairly low-tech. In any event, being able to capture light and interpret it can be a very powerful feature in your next DIY project. So, let's jump in and take a look at a few of my favourite photo- and video-related products.

Arducam PiNSIGHT vs 64MP Autofocus Camera

ARDUCAM $99 | arducam.com

ARDUCAM $69 | arducam.com

The Arducam PiNSIGHT is what got me started thinking about writing another camera-related roundup for this Best of Breed. At the heart of the add-on board is the OAK-SoM and 12MP camera, which creates a powerful AI camera for your Raspberry Pi. So, what exactly can you do with this AI module? A lot, including face recognition, object recognition, anomaly detection, and pose estimation.

The website has a lot of demo videos showing exactly what this camera is capable of recognising. It's very impressive! However, although Arducam does provide a fair amount of documentation and examples, it's not for a beginner programmer. That being said, if you are up for the challenge, and need to integrate some AI into your project, this is a great choice at a fairly affordable price for an AI board.

Sometimes you just need a lot of resolution for your project, and that's where the 64 MP autofocus camera for Raspberry Pi comes into play. Featuring a sensor with a pixel resolution of 9152×6944 and measuring in at 9.25 mm diagonally, it enables you to capture a high-resolution 64 MP image.

If you need to shoot video, the camera can capture many different resolutions and frames per second, including the standard 1080p at 30 frames per second, or 720p at 60 frames per second. As if that weren't enough, you also get an autofocus lens to keep everything nice and clear. If you're looking for a high-resolution camera for your next Raspberry Pi project, you really should check this one out.

VERDICT

Arducam PiNSIGHT

High computing power.

8/10

VERDICT

64MP Autofocus Camera

Extreme megapixels.

9/10

Pocket intervalometer for Canon

TINDIE $17 | tindie.com

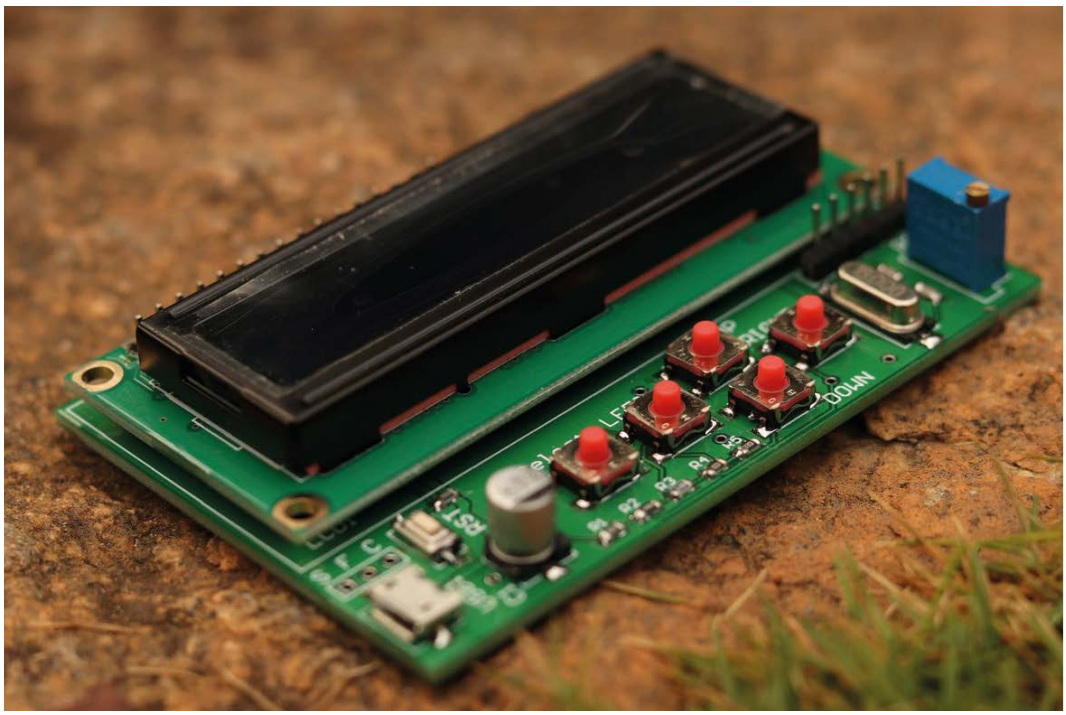

The credit card-sized pocket intervalometer for Canon cameras is a fun little kit that allows you to trigger the shutter on your camera automatically at a predetermined interval. This is perfect for time-lapse photography or long-exposure shoots.

At the core of the device is an ATmega328 microcontroller, the same found in a standard Arduino. It has a simple LCD screen, and a four-way button layout allowing you to set the total time, delay, and time interval. It also features an optocoupler, isolating the on-board electronics from the triggering mechanism connected to your camera. This is critical for protecting your expensive camera from any voltage. If you've never tried shooting stop-motion, or multiple exposures, check out this fun little kit.

VERDICT

Pocket intervalometer for Canon

Great price for a handy device.

10 /10

PAA3905 Optical Flow Camera

TINDIE $36 | tindie.com

The PAA3905 Optical Flow Camera designed by Pesky Products isn't exactly a high-resolution camera used for photography, but that doesn't mean it's not awesome. This little camera can identify motion in extremely low light in the x/y direction. It's what you would find on an optical mouse, but ready to integrate into your project.

The PAA3905E1 sensor uses an 8-bit 35×35 pixel imager that can sense both visible and infrared light. It automatically identifies features within its field of view, and creates a pixel map to track the movement and calculate the velocity. The sensor uses standard SPI for communications, making it a great solution for your next robot project using almost any microcontroller or Raspberry Pi. I picked up a similar board a few years ago for tracking small movements of a gesture controller I built, and it worked great. Head over to the website for more information and inspiration on what you can do with this little breakout board.

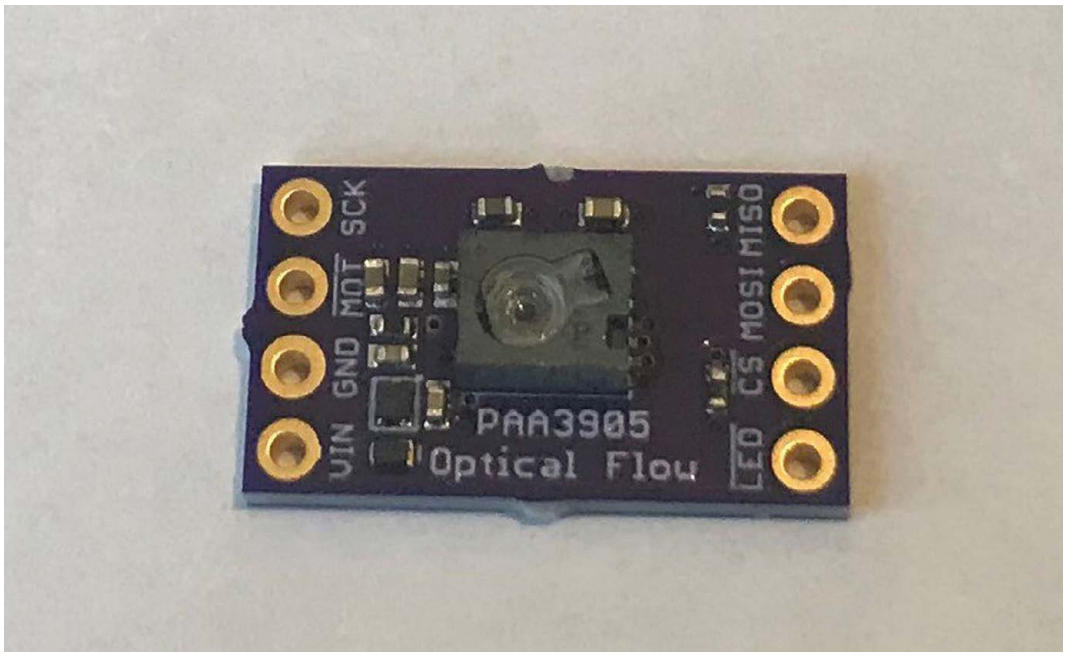

VERDICT

PAA3905 Optical Flow Camera

Very cool tech.

8 /10

Camera Wi-Fi Printer for Game Boy

TINDIE $32 | tindie.com

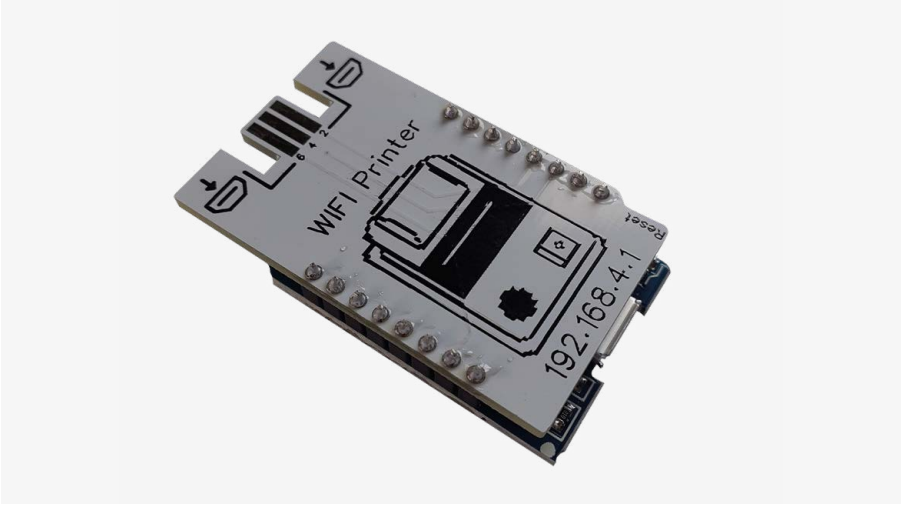

Left
This is exactly how we imagined the future would be in 1999

OK, this one is a little niche, but I love the concept. For anyone who has owned a Nintendo Game Boy camera and has a bunch of pictures that you'd love to take and share, this little board is a lifesaver. The Camera Wi-Fi Printer for Game Boy connects to the Game Boy via the Link port and allows you to 'print' to the board, which instead of printing them, stores the images on its on-board memory.

Once stored, you can access them via Wi-Fi from a web server and IP address. Keep in mind that, in addition to the board, you will also need a power supply, printer cable, or Game Boy Color Link Cable to download the pictures. Head on over to the website to see how it all works and get ready to share your super-low res, but super-cool, Game Boy pictures.

VERDICT

Camera Wi-Fi Printer for Game Boy

The ultimate retro camera for the 8-bit generation.

10 /10

CAMERA CABLE JOINER

PIMORONI $3.18 | pimoroni.com

The Raspberry Pi Camera Module typically comes with a cable to connect to your Raspberry Pi, but there are times when you need a longer cable/one that comes apart for easy assembly. And that's where this handy little board comes in. With the Camera Cable Joiner, you can add an extension cable to your standard Raspberry Pi camera.

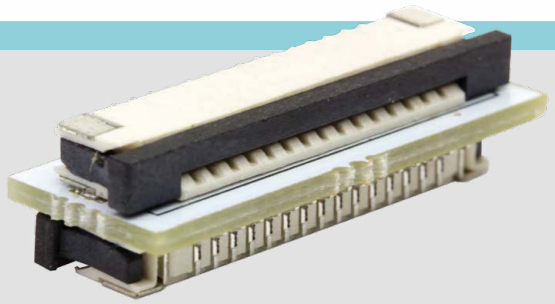

REVIEW

Raspberry Pi M.2 HAT+

Add an extra connection to your small computer

RASPBERRY PI ◆ **$12** | **hsmag.cc/m2hat+**

By **Ben Everard**

When the Raspberry Pi 5 came out, it had a new port on the bottom that enabled you to attach a flex cable. This rather unassuming port opened up a wide range of possibilities as it enabled you to connect PCI Express (PCIe) devices. This high-speed protocol is commonly used for connecting hard drives as well as other add-ons (see overleaf). However, for a long time, this remained tantalisingly out of reach because the 16-way connector needed an adaptor to work, and that adaptor hadn't been available. The most common adapter for PCIe is M.2, and this board converts the flat flex cable to an M.2 port.

Since storage is likely to be the most popular use for this, we'll look at this first. Non-volatile memory Express (NVMe) is the standard for hard drives to connect via M.2 to PCIe. They come in different physical sizes which are determined by their length. The M.2 HAT+ is compatible with 2230 and 2242, which are 22 mm wide by 30 mm and 42 mm long, respectively. They slot in and are secured by a thumbscrew.

The two main reasons for switching to an NVMe drive are performance and reliability. Performance is a massive jump up. It's hard to say exactly what performance increase you'll get because it depends on what SD card and what hard drive you have, but a 10× increase in disk access is likely. Similarly with reliability, it depends on your use case and existing hardware. Many people run for years without any problems with SD cards, but solid-state drives (SSDs) are fundamentally a more reliable technology. Perhaps the most notable difference is in getting started. Most people, when they first get started with a Raspberry Pi, plug the SD card into their computer and flash on the operating system (OS). With an NVMe drive, that's not possible, so you still need to have an SD card to boot up your Raspberry Pi in order to then load the OS onto the SSD.

BEYOND STORAGE

There's much more to PCIe than NVMe, and in theory, it is a standard, so any PCIe device could be made to work. However, life isn't that simple, and most PCIe devices need drivers. Not all of these drivers support Linux and, of those that do, they're mostly untested

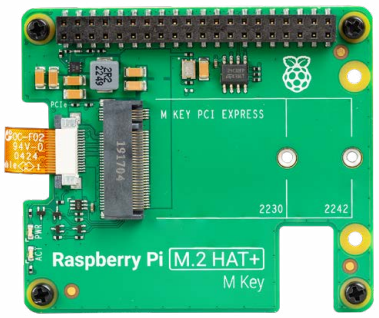

Right ◆
The board has screw holes for 2230- and 2242-sized devices

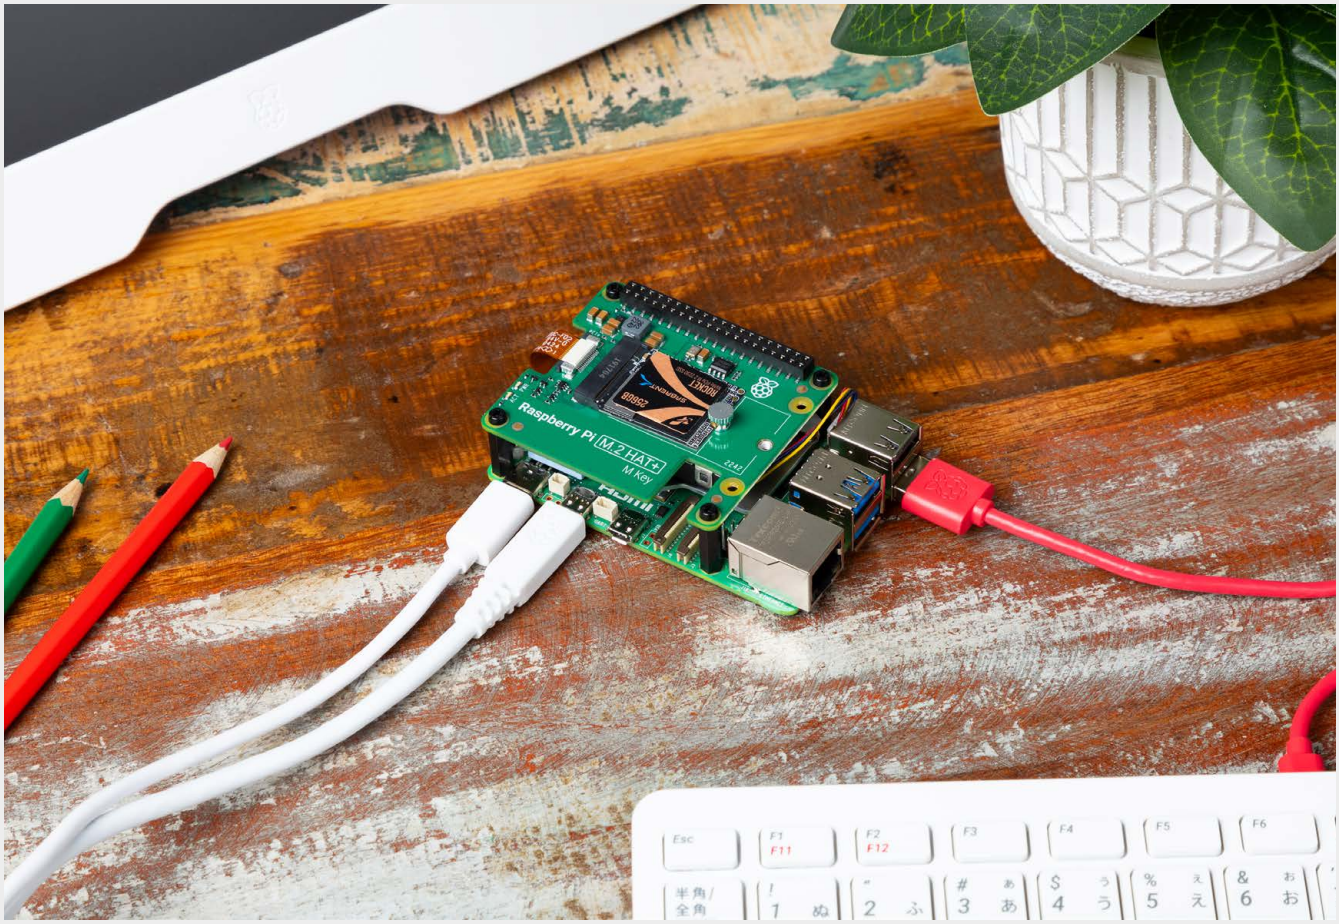

on Arm hardware. This means that if you pick a random PCIe device off the shelf, it may not work. YouTuber Jeff Geerling is attempting to work out what will and won't work, and his PCIe database is currently the best source of information about what works and what doesn't with PCIe on the Raspberry Pi: **pipci.jeffgeerling.com**.

So far, you may have noticed that we've been talking mostly about the advantages an M.2 slot can give you, not this specific M.2 HAT+, and that's because, well, it's kind of boring because it just works. It has one job and it does it well. You get a Gen 2 PCIe lane in M.2 form factor. You can increase the speed to Gen 3, but since Raspberry Pi 5 isn't certified for these speeds, you may find it less reliable. The only issue you're likely to encounter with this HAT+ is cases. It does fit in the official case, though you'll need to swap the case fan out for an Active Cooler. Other cases may or may not fit – the case manufacturer should be able to advise you.

There are a few other M.2 Raspberry Pi add-ons available, and a few of them have additional features (such as an extra M.2 port), which may be worth considering if you have specific needs.

SD cards are, in some ways, wonderful. They're cheap and easy to work with. Being able to quickly swap out the entire OS on a Raspberry Pi made it really easy to get up and running and test things out in a low-risk way. In some ways, however, they're terrible – they're slow and unreliable.

This isn't the first time there's been improved storage on a Raspberry Pi – the Compute Module 4 had optional eMMC, and some baseboards had M.2 slots. However, this brings this enhanced storage to hobbyists (as well as those who want more flexible storage options).

Now that there's a first-party way of connecting NVMe drives to the Raspberry Pi 5, it's a much more serious machine for long-term use. The additional speed and features you can get for $12 makes it probably the best price/performance upgrade available, not just for Raspberry Pi but for any computer out today. ◻

Above
NVMe storage is small, easy to use, and much faster than SD cards

VERDICT

Get serious about storage.

10 /10

REVIEW

Raspberry Pi AI Kit

Putting the ghost back in the machine

RASPBERRY PI ◈ **$70** | **hsmag.cc/aikit**

By **Ben Everard**

In case you've missed the news, AI is poised to be the next big thing in tech. Actually, scratch that, it's already the current big thing in tech. The only slight problem is that no one can quite agree what it is.

While the latest headlines are being grabbed by large language models, including ChatGPT, which have a habit of lying to users and writing uncompilable code, AI models have been quietly working away in the background. They generate captions for our videos, help us take better photographs, help scientists identify things in photographs, improve quality control in factories, and generally help make our lives progress a little smoother. The neural networks underpinning these are running everywhere, from server rooms to the phones in our pockets.

Neural networks have two stages – first, they must be trained. This is where you define the structure of the network, and run training data through it (typically large amounts of training data). While a lot depends on the particulars of the model you're training, this usually takes a huge amount of computing power and is only done rarely. In fact, the majority of people using AI don't train their own models. Instead, they use pretrained models that are available from a variety of sources (there's a wide range of models for the Hailo-8L – the accelerator at the heart of the AI Kit – available at **hsmag.cc/Hailo8L**).

Once you have a model, you can then run it – this is where you use it to analyse real-world data. Running a model takes a much more modest amount of computing power, and it's this that the Raspberry Pi AI Kit is designed to do.

The Hailo-8L accelerator can perform 13 trillion operations per second (aka 13 TOPS – though the T in this case stands for Tera). That's obviously a big number, but to put it in context, Apple's M3 processor's neural engine can perform 18 TOPs, while its A15 SoC (from the iPhone 13) can perform 15. Meanwhile, an NVIDIA A100 GPU can perform 1248 TOPS.

Models are getting faster and more accurate all the time, so it's hard to say exactly what this is capable of, since it will probably be able to run better models in a year's time than it can now. However, to give you an idea, the YOLO models can distinguish between about 80 different types of object (person, car, bicycle, etc.), and they can run quickly in real time on the AI Kit.

Similar models can detect someone's pose. Take a look at the model zoo (in previous link) for a fuller breakdown of the different models and their performance, but broadly speaking, the sorts of models this can run can differentiate between around a hundred types of object and find them in a scene.

Just as executable files have to be compiled for the particular processor you're using, neural network models have to be compiled for the particular accelerator you're using (as well as the framework they are running in). Hailo has a Dataflow Compiler that accepts models in many common formats including TensorFlow, PyTorch, and Keras. The compiler converts these input files into HEF files that can be loaded onto the AI Kit.

That's a lot about what the Raspberry Pi AI Kit is meant to be, so let's now take a look at what it is.

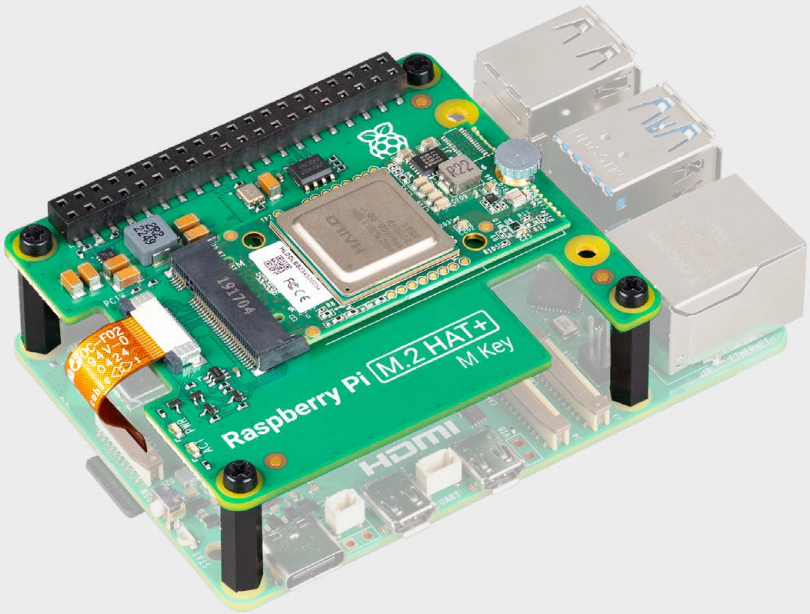

Inside the kit itself, you'll find an M.2 HAT+ and a Hailo-8L board. These two plug together and then into a Raspberry Pi 5 – because it connects to the PCIe port, earlier versions of the Raspberry Pi won't work. This is all detailed in the Getting Started Guide (**hsmag.cc/AIGettingStarted**).

Once the hardware is connected and the dependencies are installed, you can start on the software. While the Raspberry Pi AI Kit isn't explicitly a vision product, we suspect the vast majority of its use will be in vision. That's just the area where neural networks of the sort of size this can run are most useful.

At the moment, you can run the Hailo models within Raspberry Pi Camera apps by passing a suitable value for the --post-process-file parameter. There are also examples created by Hailo on its GitHub at **hsmag.cc/RPI5Examples**.

We suspect, though, that most people want to use the models in their own software. This is possible at the moment with Hailo's TAPPAS framework, but it should soon become far easier when support for the Picamera2 Python module is released.

Neural acceleration on small computers has lagged for a long time, so we're really excited to see development in this area. The Hailo-8L is powerful enough to let many vision processing tasks run in real time, while still leaving the CPU mostly free to do whatever other processing you need.

We've said plenty of times in this magazine that the products that excite us the most are the ones that open up new categories of project, and this is one such example.

It's not the first AI accelerator for small computers, but it's the first one we're aware of with this level of performance at a hobbyist price point, and it should really open up the field of embedded AI. ◻

Above
Everything is held together securely, so it's easy to embed this in other hardware

VERDICT

A new product that opens the door to many potential AI projects.

10 /10

CROWDFUNDING NOW

Inkplate 6 MOTION

Speed up your screen

From $169 | **hsmag.cc/inkplate6** | **Delivery:** Oct 2024

Computers have come a long way. Eighty years ago, the first machines were just being built – they took up whole rooms and could do little more than simple arithmetic. Modern computers fit on your wrist – admittedly, they still only do simple arithmetic, but they do it much faster. However, despite all the advancements, computers still have a fatal enemy – sunlight. Sit too close to a window and the reflections on your screen make it difficult to see.

The long-touted solution to this is e-ink – a technology that makes a screen that works like paper and is far more tolerant of ambient lighting conditions. Traditionally, this has been held back by incredibly slow refresh rates, but this might be starting to change.

The Inkplate 6 MOTION promises to refresh the screen in just 91 ms – more than adequate for interactive use. It's Arduino-compatible and includes a bucket load of GPIO. We're particularly interested to see the USB OTG port, which potentially means you could plug in a keyboard and make an e-ink typewriter. The devil, however, is always in the details, and without having one to play with, we can't tell how easy this will be.

We're excited to try this out and hope that it really does work as well as it looks in the campaign video as it will enable many outdoor projects. □

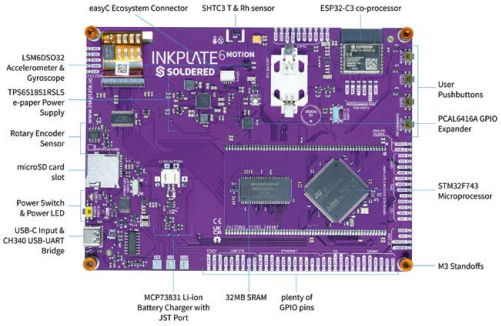

(Far) Left The driver board has programable GPIO for building your project

Below With a 3D-printed enclosure, you can make a tablet

GC6

Now with added carbon fibre

From £70 | hsmag.cc/gc6 | Delivery: Sept 2024

The blurb tells us that this is the world's first carbon fibre rotary pen. We can completely believe this because we have no idea why you'd want a carbon fibre rotary pen. Having used plenty of similar (but not carbon fibre) tools over the years, we've never once been held back by the lack of black chequerboard pattern.

Because this is a thoroughly modern tool, it comes with a smartphone app. We are honestly a little lost for words about this. As far as we can see, it just lets you select between four speed settings (these appear to be customisable). We'd hesitate to say that this is pointless, because without fully testing it out, we might be missing something, but it certainly appears pointless.

Rotary tools come in a wide range of styled from a wide range of manufacturers already. We're not sure what the GC6 adds to this already overcrowded market. The magnetic charger connection does look nice, though. □

Left
Just look at the fancy black grid pattern on the handle

HealthyPi Move

Keep track of your vitals

$249 | **hsmag.cc/healthpi** | **Delivery: Nov 2024**

The HealthyPi is a bucketful of medical sensors crammed into a watch form factor. It's open-source hardware with an open-source software stack based on Zypher. It can measure – among other things – your heart rate, blood pressure, blood oxygen, and respiration rate. If there's something about your cardiovascular or pulmonary system that you'd like to track, the HealthyPi may well be able to do it.

The most critical thing about collecting data like this is the accuracy. It's all well and good displaying a heart rate or blood pressure, but if it's not correct, then the device is at best useless and at worst actively harmful (because you could act on this incorrect information that you wouldn't do without this information). The HealthyPi is not certified as a medical device, so caution should be taken here. We look forward to seeing the device in the hands of people who can test it out so we can start to understand its reliability.

Potentially, this is a really useful bit of kit for its sheer range of features and its hackability – you can get it to monitor the sensors however you want and process or upload the data however you want. It all looks great on paper, but accuracy will be key to whether it proves to be worthwhile.

Left
The Healthy Pi looks great, but is it accurate?

BUYER BEWARE

When backing a crowdfunding campaign, you are not purchasing a finished product, but supporting a project working on something new. There is a very real chance that the product will never ship and you'll lose your money. It's a great way to support projects you like and get some cheap hardware in the process, but if you use it purely as a chance to snag cheap stuff, you may find that you get burned.

Thread Boards 2.0

Screw your storage together

From $42 | **hsmag.cc/threadboard** | **Delivery:** July 2024

Organisation is hard, and workshop organisation is particularly hard. Most workshops have a lot of stuff that all wants to be easily accessible. Ideally, everything should be easy to reach and easy to remember where it is. To make things harder, the range of stuff in most workshops is constantly evolving.

Thread Boards is the latest solution for this particular problem. It's a 3D-printable set of parts that can be mounted on a wall. There are baseplates that have threaded holes, and 3D-printable bolts that go into these. The real advantage, though, are the various holders, trays, and shelves that you can bolt on.

There are apparently tens of thousands of STLs for common household items; however, there are few details on exactly what these are. The tricky thing here is that it's hard to know if your stuff will fit. Just because there is, for example, a paintbrush holder, that doesn't mean that your paintbrush will fit. There is a promise of an STL customiser which lets you adjust some parameters of the holders by dragging sliders. Potentially, this could be very useful if it lets you quickly and easily create holders for the specific things you have. However, details are very scant, so we can't see how useful this could be.

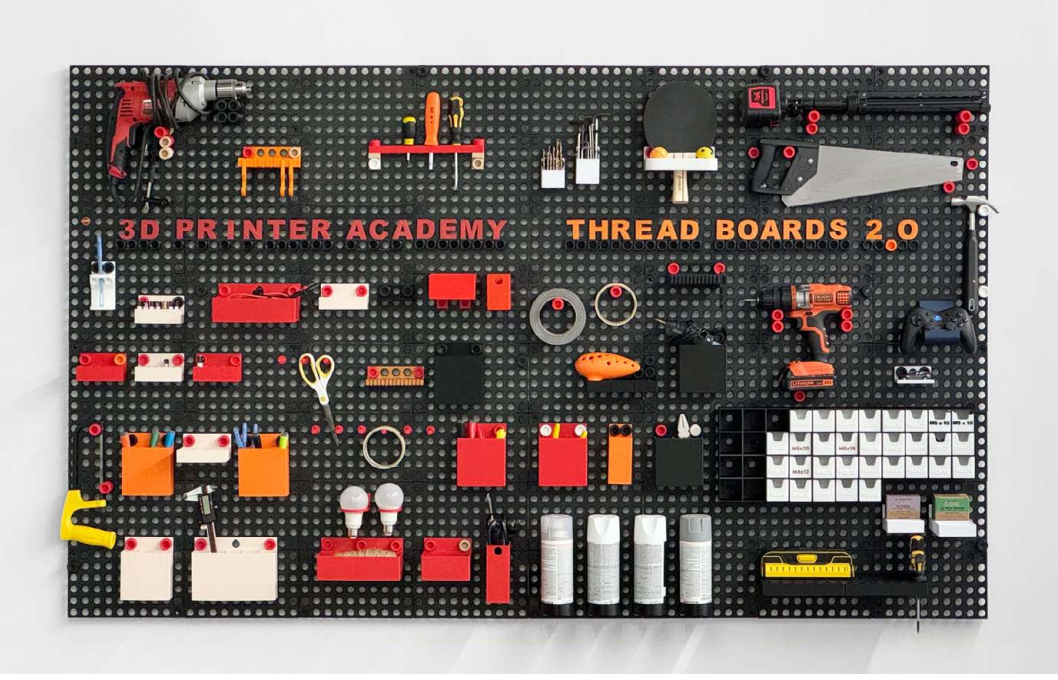

Left
Keep track of all your tools

next month

issue #81 ON SALE 25 JULY

ELECTRONICS

ALSO

- SO MANY LEDS
- 3D PRINTING
- LASER CUTTING
- WOODWORK

AND MUCH MORE

DON'T MISS OUT

hsmag.cc/subscribe

Photogrammetry

3D design is hard. There's one more dimension than you have on your screen. Frankly, it's one more dimension than we can hold in our head. At the same time, it's easier than ever to make 3D things, and hugely useful to be able to capture 3D objects in the silicon confines of a computer.

On page 66, we take a look at photogrammetry, which uses some clever computational tricks to convert a bucketload of 2D photographs into a 3D model. Using this, you can borrow designs you find in the real world. Use them just as you would any other 3D design, whether that's for 3D printing, or anything else.

PiKVM

Remote control **redefined**

Manage your servers or PCs remotely!

Powered by Raspberry Pi

PiKVM V4 Mini

Small, cost-effective, and powerful!

- Power consumption in idle mode: just 2.67 Watts!
- Transfer your mouse and keyboard actions
- Access to all configuration settings like UEFI/BIOS
- Capture video signal up to 1920x1200@60 Hz
- Take full control of a remote PC's power

PiKVM V4 Plus

The most feature-rich edition

- More connectivity
- Extra storage via internal USB3.0
- Upgraded powering options
- More physical security features
- Extra HDMI output
- Advanced cooling solution

A cost-effective solution for data-centers, IT departments or remote machines!

Available at the main Raspberry Pi resellers

core ELECTRONICS · PiShop.ca · RPishop.cz · Welectron. · elektorstore · PIMORONI

Pi-Shop.ch · The PiHut · PiShop.us · CloudFree · KUBII · Cytron Marketplace

ROBU.IN Your ideas, Our Parts · ALLTOP COMPUTERS & ELECTRONICS · netsket mall · DIREKTRONIK · RaspberryPi.dk

HiPi.io

No reseller in your country?
Check shop.hipi.io (import fees might apply).

List of official resellers by country: