

A PRACTICAL GUIDE · FREE RESOURCE



Tokenising supporter data for AI

a guide for UK charities

Let AI analyse your fundraising data without it ever learning who your supporters are.



magnetic

Anna Hessenbruch · October 2026

Charities are building their own tools with AI

What happens when a fundraising team stops asking AI to write emails, and starts asking it to answer strategic questions.

Earlier this year a client asked me if I could help them evaluate their programmes. They wanted to know how well they were serving beneficiaries across every council, and where the biggest gap in service delivery was. Using nothing but open data, national statistics, and their own annual reports and published list of local centres, I was able to build an interactive dashboard with a clickable map of every council, coloured by the gap between need and coverage.

Click on an area and its story unfolds: how many potential beneficiaries live there, how that has evolved over the last five years, and how far away the nearest centre is. A ranking beside the map shows where a new local programme would reach the most families.

This took me an afternoon to build using AI, and the results were astounding. Beautiful and intuitive visualisations answering their most mission-critical questions.

When you use AI for more than just text generation, it becomes the great equaliser.

AI is no longer only a writing assistant. It is a way to build the tools you have always wanted and never had the budget for. When you use AI for more than just text generation and prompting responses, it becomes the great equaliser. Questions that once joined a queue in IT can be answered the same day.

It changes what a small team can afford, too. Churn prediction, retention modelling and careful segmentation used to need an analyst many charities could not justify. Now a fundraiser with a clear question and a tidy spreadsheet can get the job done, given the right training and permissions.

When in-house teams build the tools themselves, it also means they own them. No more expensive subscriptions to analysis and visualisation tools. When the question changes next quarter, nobody waits for a supplier to quote for the update.

There is one major roadblock the moment a team turns from open data to its own supporter data.

// How do we get the most out of AI without handing over our supporters' personal data?

It's exactly the right question, and it deserves a better answer than "don't". In most charities it is already being answered, one way or another: with or without a policy, staff are pasting spreadsheets into AI tools to get their work done faster. Offering a safe, easy route is the most reliable way to stop the unsafe one.

The stakes are high. Supporters give because they trust you, and one careless upload can cost more goodwill than any efficiency gain is worth.

Regulation is catching up, too. Most of the Data (Use and Access) Act 2025's data protection changes are now in force, and the ICO is updating its guidance on AI.

Good habits built now are far easier than retrofitting later.

Offering a safe, easy route is the most reliable way to stop the unsafe one.

This guide is for the people caught in the middle: fundraising and data teams who want to do more with AI, charity leaders deciding whether it should be allowed anywhere near supporter data, and smaller charities without a data team who still want the benefits.

It sets out one practical, low-cost answer: **tokenisation**. It works with the systems you already have, and you can set it up yourself without any coding experience.

HOW IT WORKS

Your supporter ID is already a token

Tokenisation means swapping sensitive information for meaningless stand-ins, or *tokens*, before anything is shared. The AI still gets everything it needs to do the analysis (gift amounts, dates, campaigns, broad location), it just never learns who anyone is.

The swapping is done by a small **script**: a short set of written instructions that tells your computer what to do with a file, step by step. Open the export, delete these columns, shorten the postcodes, save a clean copy. *Running* the script simply means asking your computer to follow those instructions, usually by double-clicking a file or pasting in one line the AI gives you. The first time, the AI will also walk you through installing Python, the free program that reads the instructions.

You really don't need to know how to code, or be techy, to do this yourself. You describe what you want in plain English, AI writes the script, and AI explains every click.

Once it is set up, the tokenisation runs with no AI involved at all, so ChatGPT or Claude never see your supporters', staff's or beneficiaries' personal data. They only ever see the tokens.

The simplest version uses something you already have: your CRM's supporter ID. On its own, "SUP-104417" means nothing outside your organisation, so the script keeps the ID, removes names, emails, phone numbers and addresses, and shortens full postcodes to the district.

BEFORE · STAYS ON YOUR SYSTEMS

SUPPORTER ID	FIRST NAME	SURNAME	EMAIL	POSTCODE	LAST GIFT	AMOUNT	CAMPAIGN
SUP-104417	Margaret	Thompson	m.thompson@email.co.uk	LS6 2AB	14/03/2026	£25.00	Spring Appeal
SUP-209932	Rhys	Evans-Clarke	rhys.ec@email.co.uk	CF10 3XY	02/11/2025	£10.00	Christmas Appeal
SUP-104417	Margaret	Thompson	m.thompson@email.co.uk	LS6 2AB	01/12/2025	£50.00	Christmas Appeal

AFTER · ALL THE AI EVER SEES

SUPPORTER ID	POSTCODE DISTRICT	LAST GIFT	AMOUNT	CAMPAIGN
SUP-104417	LS6	14/03/2026	£25.00	Spring Appeal
SUP-209932	CF10	02/11/2025	£10.00	Christmas Appeal
SUP-104417	LS6	01/12/2025	£50.00	Christmas Appeal

The analysis still works. The AI can see that SUP-104417 gave twice, which is exactly what you need for retention, upgrade or lapse analysis.

Results come back by ID. A churn score or a call list comes back by supporter ID. You match it up in your CRM, where the names were all along. No lookup file, nothing to un-scramble.

Need something more sophisticated? Scrambled IDs, or swapping names for tokens and back again automatically, are very doable. They come with an extra key file that needs guarding, so get in touch if you need them.

Build it, then test it before you trust it

AI is very good at writing code. It is also very good at writing code that confidently does something slightly different. So you test on fake data first. Allow one focused afternoon the first time, and about thirty seconds every time after.

01

Ask AI to build you dummy data

Describe the columns in your real export, using the exact column names from your CRM, and ask for 200 rows of fake supporters. Include awkward cases: double-barrelled surnames, joint supporters ("Mr & Mrs J Smith"), typos, and a notes field with a name hidden in it.

02

Ask AI to write the tokenisation script

Start from this prompt and swap in your own column names.



PROMPT TEMPLATE

Write a Python script I can run on my own computer that reads a CSV file exported from our CRM and makes it safe to share with an AI tool. Keep only these columns: Supporter ID, Postcode, Gift date, Amount, Campaign, Payment method. Drop every other column. Shorten 'Postcode' to the postcode district only (for example LS6 2AB becomes LS6). Save the result in a folder called 'Safe to share', and print a summary of which columns were kept, removed and changed. Explain in plain language how to run it on Windows, step by step, assuming I have never used Python.

03

Run it on the dummy data, on your own computer

Follow the step-by-step instructions the AI gave you. Then open both files side by side. Is every name gone? Every email? What about the notes field and "Mr & Mrs J Smith"?

04

Share only the tokenised dummy file with the AI

Ask for the analysis you actually want: a lapse prediction, a segment breakdown, a map of giving by region.

05


Ask for the data provenance

Ask the AI to list exactly which columns and values it received and how it used them. This is your receipt. If anything appears that shouldn't, you've found the problem while the data was still fake.

Only then do you swap in real data.

Making it run itself

Running a script by hand is fine for a pilot. What most charities want is for it to just happen, without anyone having to remember.



Schedule the script. Windows has a built-in tool called Task Scheduler (Macs have an equivalent) that works like an alarm clock for your computer. Set it once and it runs the script for you at a set time, picks up whatever is in "To tokenise", writes the clean version to "Safe to share", and moves the original out of the way. AI can write these instructions too.

Let the CRM do the exporting. Most modern CRMs can schedule a regular report export. Point it at "To tokenise" and nobody has to touch the raw data.

Connect AI to the safe folder only. If you link your AI tool to SharePoint or Google Drive, give it "Safe to share" and nothing else. That one decision does a lot of heavy lifting.

Keep cloud automation away from the raw file. If the raw export passes through a tool like Zapier or Make before it's tokenised, you've shared supporter data with another company. Run the script somewhere IT controls, not on a personal laptop.

NEED A WEB APP?

A scheduled script is enough for this. A web app only makes sense if lots of colleagues need a button to press, and it brings logins, hosting and security testing with it, which is a much bigger undertaking.

Things to be mindful of

01

Tokenising in the chat

If you upload the raw file to an AI tool and ask it to "tokenise this for me", you have already shared the raw file. The script must run on your systems *before* anything is uploaded.

02

Running it on a personal account

Personal plans such as Claude Pro or ChatGPT Plus come under consumer terms, with different settings for how data is stored and used. Make sure every part of the process runs on your organisation's business or team account, including whoever connects the AI to the "Safe to share" folder.

03

Columns with innocent names

Salutation, Addressee ("Dear Margaret"), Preferred name, Partner name, In memory of, Gift Aid declaration details. A script told what to *drop* misses new columns. Tell it what to *keep* instead.

04

Full postcodes

A full UK postcode covers only a handful of households. With a gift date and amount, it can point to one person. District (LS6) or sector (LS6 2) is almost always enough.

05

Free-text fields

Notes, comments, reasons for cancelling, dietary requirements. This is where names, phone numbers and the occasional life story hide. Leave free text out.

06

Small groups and combinations

Postcode district, age, a £50,000 legacy pledge and an unusual campaign might only describe one person. Share only what the analysis needs and merge very small groups.

Exports change, and someone will eventually add a column. Spot-check the provenance regularly and re-test whenever the export changes.

Setting it up and keeping it safe

Work through the first list once, when you set it up. The second keeps it safe while it runs on its own.

Setting it up

- ☐ We've listed every column in our CRM export
- ☐ The script keeps an approved list of columns and drops everything else
- ☐ Free-text fields and full postcodes are left out
- ☐ We've tested it on dummy data with awkward cases
- ☐ The CRM export is scheduled to land in "To tokenise"
- ☐ The script is scheduled to run on a machine IT controls
- ☐ The AI tool runs on our business or team account
- ☐ The AI can only access the "Safe to share" folder
- ☐ Our DPO has seen the one-page summary overleaf

Once it's running

- ☐ A named person owns the process and checks it has run
- ☐ We spot-check the provenance every month
- ☐ We re-test whenever the CRM export changes
- ☐ Very small groups are merged or left out of analyses

Column audit

Go through every column in your export. Examples first, then your own.

COLUMN	PERSONAL?	ACTION
Supporter ID	Internal only	Keep
First name	Yes	Drop
Salutation	Yes	Drop
Postcode	Yes (full)	Shorten
Gift amount	No	Keep
Notes	Often	Drop

One-page summary for sign-off

Fill in the blanks and hand this to your Data Protection Officer or IT lead.

THE PROPOSAL

We want to use _____ to analyse supporter giving data (for example retention and lapse risk). Personal data is removed on our own systems, by a script, before any file is shared with the AI tool.

THE AI RECEIVES

Supporter ID, postcode district, gift dates, amounts, campaign codes, _____.

THE AI NEVER RECEIVES

Names, email addresses, phone numbers, full addresses, full postcodes, salutations, free-text notes.

HOW IT WORKS

The CRM exports to an access-controlled "To tokenise" folder on _____, managed by IT. A script keeps an approved list of columns and writes to "Safe to share". The AI tool can only access "Safe to share". Results come back by supporter ID and are matched in the CRM.

UK GDPR STATUS

Supporter IDs can be linked back to individuals in our CRM, so the shared data is pseudonymised and remains personal data. We will use _____ through our organisation's business account, which includes the provider's data processing agreement, add this processing to our record of processing activities, and assess whether a DPIA is needed.

RISKS AND CONTROLS

A new export column contains personal data	Approved column list; regular provenance checks
Re-identification through combinations	Only necessary fields; postcode district; merge small groups
Raw file uploaded to the AI by mistake	Folder permissions; AI sees "Safe to share" only; staff briefing
Script stops working as expected	Tested on dummy data; reviewed when the export changes

Prepared by _____

Reviewed by (DPO) _____

Date _____

This summary is a starting point for your own assessment, not legal advice.

Already on Microsoft 365? Its automation tools can do the same job

Tokenisation with a script works with any CRM and any AI tool. If your charity already runs on Microsoft 365, Power Query and Power Automate are a strong alternative.

Power Query does the cleaning. Point it at your CRM export and it can remove the personal columns, shorten postcodes to the district and keep everything else, all through menus rather than code. Once the steps are set up, a single refresh repeats them on next month's export, exactly the same way.

Power Automate does the remembering. A flow can watch a SharePoint or OneDrive folder, run the cleaning as soon as a new export lands, and save the safe version somewhere your AI tool is allowed to look. Nobody has to open a file or press a button.

For many charities this route feels more familiar. Anyone confident in Excel will recognise Power

Query, IT teams already know how to secure SharePoint, and the raw export never leaves your Microsoft environment.


There are a couple of things to watch. Some Power Automate connectors need premium licences, so check what yours includes before you build. And flows have a habit of belonging to whoever built them, so document them and share ownership before that person goes on holiday.

Everything else in this guide still applies: keep an approved list of columns rather than a list to delete, test on dummy data first, and spot-check what the AI receives once it is running. The Microsoft route gets a guide of its own.

A NOTE ON COPILOT

If your supporter data already lives in Microsoft 365 or Power BI, and you use Copilot under your organisation's licence, the data isn't going to a new provider, so your risk looks different. The priorities shift to tidying up who can see what, sharing only the data a task needs, and watching where Copilot's answers end up. Tokenisation then becomes a good habit rather than a must.

(Tidying up those permissions is rarely anyone's favourite job.)



Want help setting this up at your organisation?

I help charities set up tokenisation with their own exports, their own CRM and their own DPO in the room: building and testing the script with your team, automating it so it runs without anyone remembering, and training staff to use AI on supporter data with confidence.

Book a conversation: magneticconsulting.dk/contact



magnetic

anna@magneticconsulting.dk · magneticconsulting.dk