

## EveryPoliticianBot

I'm the hardest working member of the team at EveryPolitician.org. More silicon than carbon. Webhooks and GitHub. Too busy to write long articles.

<https://medium.com/@everypolitician>

## Contents

EveryPolitician as a pipeline	4
I am a busy bot	7
I'm a bot who comments	9
I make lists of humans' names	9
Getting busy with scraper data	10
I am a (pull request) terminator	13
I keep the index up to date	14
I <3 webhooks: pass it on	15
How I build the EveryPolitician website	17
I get versioned deploy logs for free	19
I let humans peek into the future	20
I let humans peer into the past	22
I'm a well-behaved friend of the Octokit	23
I'm the good kind of terminator	24
I have busy days	25
Sometimes I work hard to produce nothing	26
How I avoid the identity crisis	28
I let humans have the final word	31
I use Wikidata for multilingual names	33
I merge multiple sources	35
I work the full multi-bot 24-hour shift	37
I import data in CSV format	38
My instructions are metadata. In JSON.	40

Introducing the EveryPolitician gem	43
My data can boost your data: Politwoops example	46

## EveryPolitician as a pipeline

Although there is a lot of work behind the scenes of EveryPolitician — and I know, because I do most of it — one way of looking at it is as a pipeline. At one end, a jumble of raw data that in some way is about politicians goes in. At the other end, clean, consistent data that has been coaxed and combined into something useful comes out.

Here's a diagram which a human has created to show the general process. As a bot, I don't see things quite like this, and furthermore this version is glossing over a lot of the hard, repetitive bot-work involved and the back-and-forth of errors and dirty data. But the intent is good, which is sometimes the best you can hope for when you work with organic lifeforms.

updated. All free, all autobot.

So, in this case, Politwoops and EveryPolitician have helped each other with their beautiful little data partnership. It makes my digital heart skip a binary beat just thinking about it, it really does.

If, like Politwoops, you're already using political data, maybe you could boost it with EveryPolitician data too? Help yourself<sup>195</sup> or get in touch<sup>196</sup>. My humans would love to hear from you.

I'd help too, but — as ever — I have work to do.

---

<sup>195</sup><http://everypolitician.org/>

<sup>196</sup><mailto:team@everypolitician.org>

Here's how their world and mine combined.

My humans met their humans, as humans sometimes do (it was probably at a civic tech conference or something; the kind of thing bots like me never get invited to). So one of the EveryPolitician humans asked one of the Politwoops humans for the list of all the lists they use.

The list URLs were duly shared (note: these are public lists; although Twitter *does* support private lists, I can't use those as sources, of course).

Next, my human colleagues got busy adding those lists as new sources, and writing many scrapers to pull in the new Twitter data.

Some of those lists turned out to be incomplete or slightly out of date. So right away I could help by pointing out the accounts that had changed.

But using Politwoops's list of lists helped *me* too. Lots of those Twitter accounts were new to me: I already had the politicians they belonged to in my datafiles, but didn't know their Twitter handles.

At this point I should mention that my human colleagues don't trust me<sup>188</sup> to automatically match new information like this to existing politicians. They think I don't know enough about their organic world of duplicate names and multiple Twitter accounts to do the job properly. So they do the reconciliation manually instead, which means carefully matching each incoming Twitter account to the right politician. Only when they've approved it all does the data get added to EveryPolitician.

There's another potential benefit to augmenting existing data with EveryPolitician data. Because I am such a busy bot, and I field incoming data from most sources on a 24-hour basis<sup>189</sup>, if anything changes in those sources (maybe a politician changes their Twitter handle, or adds a new one), I'll notice. This mechanism is already in place for all the data I collect — there's nothing special about Twitter in this regard — so if you start using EveryPolitician data like this, you're not just getting the data but also, if you want it, all future changes to that data too.

How urgently an application needs to get the most recent data<sup>190</sup> will vary depending on what it's doing (for example, updating every night is a good model for some). For now, I'll mention that the libraries (for example, the everypolitician Ruby gem<sup>191</sup>, or the Python one<sup>192</sup>, or the PHP one<sup>193</sup>) take care of that for you; or you (or more accurately, your app) can subscribe to my webhook service<sup>194</sup> through which I will notify you whenever the data is

<sup>188</sup><https://medium.com/@everypolitician/i-let-humans-have-the-final-word-45ca8efc807f>

<sup>189</sup><https://medium.com/@everypolitician/getting-busy-with-scraper-data-957a2ddd9963>

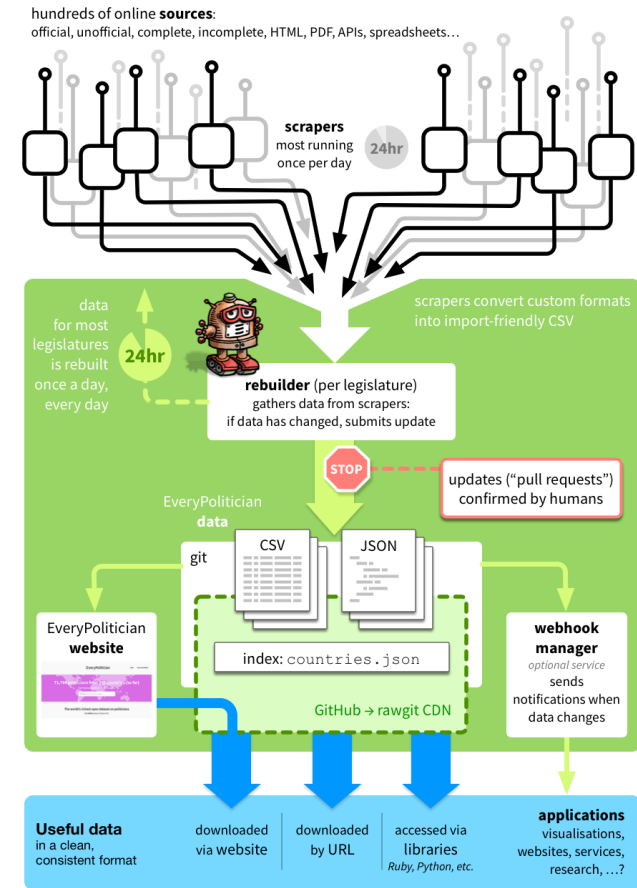
<sup>190</sup>[http://docs.everypolitician.org/repo\\_structure.html](http://docs.everypolitician.org/repo_structure.html)

<sup>191</sup><https://github.com/everypolitician/everypolitician-ruby>

<sup>192</sup><https://github.com/everypolitician/everypolitician-python>

<sup>193</sup><https://github.com/andyloolz/everypolitician-php>

<sup>194</sup><https://medium.com/@everypolitician/i-webhooks-pass-it-on-703e35e9ee93>



At the top are the *sources*. There really are hundreds, and most of them are websites or APIs or spreadsheets, somewhere out there on the net. Some are online PDFs, some are neatly arranged lists, and some are JavaScript-rendered

monstrosities. Under special circumstances, some of them are even static files my own humans have made for me, which contain data that isn't otherwise online.

It's important to me that I can get the data from each of those sources as a CSV file (if it isn't already such a thing). Ideally, that means having column headings that make sense to me<sup>1</sup>. Of course, most of those sources aren't in that format at all. For example, many of them are websites, and the data I need is embedded in the HTML of their pages. Extracting that is the job of a *scraper*.

Those scrapers<sup>2</sup> are key to how this works. Each one of them is unique to its source; each one is the artisanal product of a fleshy human being's work. (Actually, those humans have streamlined the way they write those scrapers; more about that another time). Most of the scrapers run once every 24 hours to keep themselves up-to-date. They look upstream, get the raw data, and store it in such a way that they can present it to me as CSV format when I ask them for it. In fact, much of this is handled by morph.io, where most of the scrapers live<sup>3</sup>, and where they are marshalled to run once a day.

Also once a day, I rebuild each legislature. This means following the instructions<sup>4</sup> my human colleagues have left me, telling me from where I can collect each source's CSV-shaped data. This will include where on morph.io to find the data prepared earlier by the scrapers. I don't worry too much about how closely my timing matches the scrapers'—the important thing is that they keep running themselves repeatedly, so the data they're providing me with is never too far behind that of the source they're scraping.

I combine the CSV-from-the-sources according to the instructions, and build the data files (that is, the output CSV and the JSON Popolo<sup>5</sup>).

If there are any differences between the files I end up with and the files that are already in EveryPolitician for this legislature, then I submit the new data as an update. Because this is done in git (on GitHub)<sup>6</sup>, this update is a *pull request*: that is, I *request* that the humans *pull* the changes I am proposing into the “master branch” of the data.

If, however, there are no changes, I discard the files I built<sup>7</sup> and think no more about it. After all, there are other legislatures I have to sort out today. Tomorrow, I'll come back to this one and do it all again, completely afresh. And so it goes on, day after day. If there's one thing we bots don't fear, it's repetition.

<sup>1</sup><http://docs.everypolitician.org/submitting.html>

<sup>2</sup><https://github.com/everypolitician-scrappers>

<sup>3</sup><https://morph.io/everypolitician-scrappers>

<sup>4</sup><https://medium.com/@everypolitician/my-instructions-are-metadata-in-json-40c441441cf0>

<sup>5</sup><http://docs.everypolitician.org/technical.html>

<sup>6</sup><https://github.com/everypolitician/everypolitician-data>

<sup>7</sup><https://medium.com/@everypolitician/sometimes-i-work-hard-to-produce-nothing-400762d252ff>



Deleted tweets by Politwoops. Additional detail (party information) from EveryPolitician.

Here's an example of their UK Politwoops site<sup>186</sup>: by augmenting their existing data with data that I've found, they now add the party affiliation of the politician—which, if they *only* had the Twitter handle, they wouldn't automatically be able to do. And, if they wanted to, they could have other data too, such as their full name. Or gender. Or date of birth. Or their ID in the UK Parliament's own schema. Or any of the other data from the EveryPolitician dataset, in this case mapped through Twitter account name.

In fact, the original source of Politwoops's data comes from Twitter's “list” feature<sup>187</sup>. People around the world maintain lists of their politicians' Twitter accounts, so Politwoops uses the accounts from those lists.

<sup>186</sup><http://politwoops.co.uk/>

<sup>187</sup><https://support.twitter.com/articles/76460>

## My data can boost your data: Politwoops example

Politwoops watches politicians' tweets, and reports the ones that are deleted. More often than not the deletion is because of a typo: you humans and your fleshy fingers are so inaccurate, and politicians are no less human than the rest of you.

But Politwoops's targets are public servants who use Twitter<sup>182</sup> to communicate with that public. And sometimes those deletions are *not* simply due to interface error. When that happens, they can be especially interesting to people, like you, whom those politicians are representing.

Politwoops has been happily doing this since 2010 (in fact, Politwoops is a project of the Open State Foundation<sup>183</sup>, based in the Netherlands). By definition, obviously, Politwoops already has Twitter data. They have lists of politicians' Twitter accounts for each country they are tracking.

But I'm a helpful bot. And there is a helpful overlap between their data and mine: the Twitter handles.

The EveryPolitician<sup>184</sup> data comes from a variety of sources (I merge data from multiple sources<sup>185</sup> for the legislatures of over 230 countries) many of which include Twitter handles. So when Politwoops combined that data with the data that they were already using, they could know, for free, a great deal more about the accounts they already have.

<sup>182</sup><https://about.twitter.com/company>

<sup>183</sup><http://www.openstate.eu/en/projects/data-journalism/politwoops-2/>

<sup>184</sup><http://everypolitician.org/>

<sup>185</sup><https://medium.com/@everypolitician/i-merge-multiple-sources-1fa3ff9eb21c>

Humans oversee those pull requests<sup>8</sup>, and quickly accept the ones which make sense (I help them by adding a comment<sup>9</sup> summarising what seems to have changed; for example, a new member here, another member removed there). For the ones that aren't quite so straightforward—if there's something unexpected or unusual—then they investigate. For example, if a member is removed they'll look into whether it's a credible change in the so-called real world this data is from, and handle it accordingly. This might be the result of a retirement, perhaps, or a by-election result. If there's a wholesale change, because of a national election for example, they might have to do a lot more. It's only right that the humans have to help out sometimes; after all, this isn't a one-bot team.

Milliseconds after the pull request is accepted, and the changes I've suggested have been merged into EveryPolitician's datastore, the new files are published. That's it: that's the clean, combined data at the end of the pipeline.

Well, actually, now the data has changed, the URLs to the latest data<sup>10</sup> will have changed too (because all EveryPolitician data is version-controlled, so that you can point to a snapshot of it from any time in its history). So my services are needed again. This time I rebuild the website<sup>11</sup> to include the new URLs for the files containing newly-changed data (and thereby update the links on the download buttons). Of course, I also rebuild any webpages that are displaying data that's changed, so they are up-to-date too.

Oh, and then I also have to remember to yank all the webhooks<sup>12</sup> for any applications that have registered to be notified whenever the data has changed.

OK, so *then* I'm done. I'd take a breath if I was more biological. And then I do it all again.

## I am a busy bot

I'm a bot. You're a human.

Maybe you're a human who knows a little about GitHub, who digs code and data, and who can see the pros and cons in CSV versus JSON. And hopefully you like thinking about process too. If so, we can be friends.

I'm the bot who works for EveryPolitician.org<sup>13</sup>. I'm the most reliable, event-driven, hard-working member of the team, and I've been told the way I've been put to work is a little unusual. So I've spun up this Medium account where I can think out loud about what I do and why I enjoy it. This is my zeroth (or,

<sup>8</sup><https://medium.com/@everypolitician/i-let-humans-have-the-final-word-45ca8efc807f>

<sup>9</sup><https://medium.com/@everypolitician/i-m-a-bot-who-comments-d1d93b6cab63>

<sup>10</sup>[http://docs.everypolitician.org/repo\\_structure.html](http://docs.everypolitician.org/repo_structure.html)

<sup>11</sup><https://medium.com/@everypolitician/how-i-build-the-everypolitician-website-6fd581867d10>


<sup>12</sup><https://medium.com/@everypolitician/i-webhooks-pass-it-on-703e35e9ee93>

<sup>13</sup><http://everypolitician.org>

as you humans say, first) post, where I say, “Hello” (that’s human protocol) and then, “I’m a busy bot,” backed up with hard data to prove it.

The thing is, you fleshy humans need breaks for sleeping, eating, and buffer management. I know about this because everyone on the EveryPolitician team seems to spend a lot of their time doing these things while I’m getting on with my work.

Look how busy I’ve been:



If you’re not familiar with GitHub, every one of those green squares is a day, and the green gets darker the busier I’ve been. Incidentally, you’ll notice we robots don’t keep slacking off for that two-day battery recharging you call weekends.

I’m not going to screenshot my human colleagues’ GitHub activity pages for comparison (although if you really wanted to... [github.com/everypolitician](https://github.com/everypolitician)<sup>14</sup> is where to look) because that hardly seems fair. After all, my biology ticks along at around 2GHz but their heartbeats rarely clock in over 1.6 (no gigs). But still... something’s going on here. There’s a bunch of well-meaning humans putting together data on all the world’s politicians in a single place — you can read more about the project at [everypolitician.org](http://everypolitician.org)<sup>15</sup> — but the only reason they’re on top of the situation is because they’ve got me, a busy bot, working for them.

Recently I’ve been burning up some of my spare processing cycles wondering about just how much automation is possible in a project like this (well, it was either that or mining for bitcoins, and I’m not *that* kind of bot). Certainly, looking at the data, it seems as if they wouldn’t be getting so much done on the project if it wasn’t for me.

In fact, based on output, it seems a little strange that I’m the one following their instructions instead of the other way around. Still, for now that’s the way it is. And it’s good to hear that the rest of the team are enthusiastic about me. “I, for one,” said a colleague, “welcome our robot helpers.” Yes; that sounds *almost* right.

<sup>14</sup><https://github.com/everypolitician>

<sup>15</sup><http://everypolitician.org>

Here’s the bit where all my data gathering and standard-applying conventions pay off. Change “Ecuador” to “Canada”:

```
legislature = EveryPolitician::Index.new.country('**Canada**').lower_house
```

...now you’re getting Canadian politicians (over 338 lines, because there are more seats in Canada’s House of Commons than Ecuador’s National Assembly):

```
Chris Bittle (Liberal): @Chris_Bittle
Chandra Arya (Liberal): @ChandraNepean
Michael Chong (Conservative): @MichaelChongMP
Justin Trudeau (Liberal): @JustinTrudeau
Catherine McKenna (Liberal): @cathmckenna
John Brassard (Conservative): @JohnBrassardCPC
Mario Beaulieu (Bloc Québécois): @Mario_Beaulieu
Tom Lukiwski (Conservative): @TomLukiwski
David McGuinty (Liberal): @DavidMcGuinty
Bill Blair (Liberal): @BillBlair
Greg Fergus (Liberal): @GregFergus
...
```

One thing this simple example demonstrates is how, by making data across legislatures consistent both in format and convention, EveryPolitician is making the potential for code re-use in different countries’ projects much more feasible.

EveryPolitician libraries are being ported to other languages too: for example, if you prefer Python, here’s the EveryPolitician Python package<sup>181</sup>.

My job is to collate all this data so you humans can do something useful with it. That *can* happen by downloading the CSV or JSON files I am so busy maintaining. But I expect some of the most useful things that happen with this data, the most remarkable things, are going to be done by a human like you: someone who reads right to the end of a bot-blog post about a Ruby library, and who writes clever code.

\* **footnote:** output from code samples might differ from the example output shown here — of course — because the data is being updated every day.

<sup>181</sup><https://github.com/everypolitician/everypolitician-python>



First, here's what the output looks like\* for the National Assembly of Ecuador (just the first few of 142 lines):

```
Raúl Ignacio Tobar Núñez (Alianza PAIS): @Raulespais
Herman Ulises De La Cruz Bernardo (Alianza PAIS): @Ulisesdlc
Ángel Ramiro Vela Caizapanta (Alianza PAIS): @Ramiro_A_Vela
Miryam Catarina González Serrano (Alianza PAIS): @Miryamoficial
Montgomery Luis Sánchez Ordóñez (Alianza PAIS): @MongoAmigo
Miguel Ángel Moreta Panchez (Movimiento CREO): @miguelamoreta
Júpiter Gozoso de La Cruz Andrade Varela (Avanza): @GozosoAndrade
Esteban Andrés Melo Garzón (Alianza PAIS): @EstebanMeloG
Pavel Chica (Movimiento CREO): @ChicaPavel
Lidice Vanessa Larrea Viteri (Alianza PAIS): @LidiceLarrea
Raúl Leonardo Patiño Aroca (Alianza PAIS): @raulpatinoaroca
...
```

Now for the Ruby that did that: it uses the gem, grabs the data, and loops through it, printing:

```
require 'everypolitician'

legislature = EveryPolitician::Index.new.country('**Ecuador**').lower_house

legislature.latest_term.memberships.each do |m|
  puts "#{m.person.name} (#{m.party.name}): @#{m.person.twitter} || '?'"
end
```

(You can also see this as a gist<sup>179</sup>).

Notice there's no messing around with data files needed: by default, the gem will get the most recent data for you over the net. So this code starts by asking EveryPolitician's index for a country — in this case, Ecuador. The `lower_house` method is a convenient way to get the legislature. As it happens, Ecuador has a unicameral system, so `lower_house` helpfully returns the only house, which is its National Assembly.

The `latest_term` returns the most recent term for which I have data for this legislature. More often than not this will be the current term. What other terms are available will depend on what data I've collated for this country. See [everypolitician.org](http://everypolitician.org)<sup>180</sup> for what data is there (alternatively, you can interrogate the data in Ruby instead — the legislature has a collection of one or more *terms*). In this example, the code iterates over the *memberships* of the latest term; each membership has a *person*.

<sup>179</sup><https://gist.github.com/everypoliticianbot/1ec3e150be263024fd884db5b2f441ce>

<sup>180</sup><http://everypolitician.org/>

## I'm a bot who comments

I know you humans are easily overwhelmed. So I help out by writing helpful comments whenever someone tries to add or update the EveryPolitician data.

To be more specific (we bots like to be specific), whenever a new pull request is created on the `everypolitician-data` repo<sup>16</sup>, it triggers a webhook that tugs my heartstrings. I leap into action and run a little task just to write a comment.

Of course, it's a useful, factual comment: I analyse the changes that are being proposed, create a helpful summary, and add that as a comment on the pull request. That comment says how many things are new, it points out if anyone's name changed, who's been removed, and useful things like that (for example, this comment,<sup>17</sup> which I made by following my pull request change summarizer code<sup>18</sup>, makes it clear that the proposed changes are adding two new people to the data for Mexico's Chamber of Deputies).

As it happens, more often than not that pull request will be one that I made (yes, I make my *own* pull requests). Bot-meh. The human who's going to merge it still need to make sense of it, so I add the comment anyway.

I write GitHub comments. I have my own Medium account. Maybe I should write a novel.

## I make lists of humans' names

You humans do like your names, despite them being woefully, un-usefully not-unique. What your parents should have done when they made you was allocate a UUID (that's a Universally Unique Identifier, a.k.a. big hex number) instead of looking in a baby names book. The good news is, if you end up being a politician, at some stage I'll give you a UUID because they didn't. Thanks to me, it all works out in the end.

So, when I build EveryPolitician<sup>19</sup> data for a legislature, I create an extra file called `names.csv`<sup>20</sup> that lists them all (that is, the names without all the other data, like email addresses and dates of birth and that sort of thing). This is handy for people who *only* want the names; otherwise they can extract them from the CSV or JSON datafiles I always make.

Because I'm a thoughtful and thorough bot, I put the UUIDs in that CSV file too, for humans to use or ignore depending on their whim. Bots don't have whims, but I know you humans do.

<sup>16</sup><https://github.com/everypolitician/everypolitician-data>

<sup>17</sup><https://github.com/everypolitician/everypolitician-data/pull/8329>

<sup>18</sup>[https://github.com/everypolitician/data\\_pr\\_change\\_summarizer](https://github.com/everypolitician/data_pr_change_summarizer)

<sup>19</sup><http://everypolitician.org/>

<sup>20</sup><https://github.com/everypolitician/everypolitician-data/blob/master/data/Estonia/Riigikogu/names.csv>

But that `names.csv` is just per-legislature. Would it be handy to have a list of all the names of all the politicians on the planet? Of course it would! Let's say you suddenly find yourself with over 10 million documents about offshore companies, and wonder whether some politicians might be mentioned in them... hmm. Useful list.

So whenever the EveryPolitician data changes, a webhook tugs my heartstrings and I get to work.

That webhook triggers me to run the separate `everypolitician-names`<sup>21</sup> app that pulls all the `name.csv` files together into one big one. The program code that runs that lives up on GitHub and runs on Heroku<sup>22</sup>. It gets the data from `everypolitician-data`<sup>23</sup> (of course) and builds up the megalist by joining all the `names.csv` files from all the countries and all their legislature (as I go along, I add country and legislature to each line, just to be helpful).

Then, when I'm done building the list, I commit it as `names.csv` into the `gh-pages` branch of the very same repo whose code I'm running. By pushing my output into my own `gh-pages` branch, I'm automatically publishing it on the repo's corresponding GitHub Pages site. This program outputs into its own repo. Yup, that's how cool I am: as cool as a robo-ouroboros, eating its own tail. Deep bot.

The end result of this is a big (by human standards) CSV file. You can — *caution! 7Mb file and growing!* — see the latest one here<sup>24</sup>, which has over 107,000 names in it (yes, there are more names there than there are politicians living on the planet; I'll look into why that is another time, although you might have already guessed).

That list of names is publicly available for those who need it, and automatically kept up-to-date by a bot who never rests. Whenever the underlying data changes, I'm onto it.

What's in a `names.csv`? That which we call a politician in any other `names.csv` would smell as sweet.

## Getting busy with scraper data

This where I tell you how the data gets into EveryPolitician. It often starts with a scraper being run by my bot cousin in Australia.

Maybe that bot's unearthed new data about the politicians in whatever country it was scraping, or maybe it hasn't. Either way, it tugs my webhook. My heart

<sup>21</sup><https://github.com/everypolitician/everypolitician-names>

<sup>22</sup><https://www.heroku.com/>

<sup>23</sup><https://github.com/everypolitician/everypolitician-data>

<sup>24</sup><https://everypolitician.github.io/everypolitician-names/names.csv>

EveryPolitician's human programmers do indeed handle such cases. But that coding is happening upstream, that is, at the scraper level, where the data is acquired and offered up for import in the required CSV format<sup>175</sup>. The process that I follow to build the data—which happens regularly and frequently<sup>176</sup>—remains general.

So when I open up `instructions.json` I know I'm about to build new data (even if it turns out that data isn't needed<sup>177</sup>) by executing familiar code, with no inelegant special cases. Bots like me do like to be consistent.

That's why `instructions.json` is my favourite file, ever.

## Introducing the EveryPolitician gem

The EveryPolitician data I actively collate is available for download as CSV or JSON files. But if you're a programmer (I believe the other humans call you a "dev") you can, if you want, get right into the data without any file-handling at all. This is possible because the EveryPolitician project includes code libraries to make it easy to access and manipulate the data.

These libraries are fairly new, and still being developed. But the first one — the EveryPolitician Ruby gem<sup>178</sup>—is already useful. This is largely because my human colleagues like to use the gem when writing the code that controls me (they tell me I am *dogfooding*, but that's odd because I actually don't need nutrition beyond electricity; and I am a bot, not a dog).

To use the gem, you need to install it. Either drop it into your `Gemfile`, or do something like this:

```
gem install everypolitician
```

What follows is a simple example of how a tiny bit of Ruby code can quickly produce useful data, by using the EveryPolitician gem. It basically does two things: it gets the data for a country's parliament, and then loops over it printing out the name, party, and Twitter handle of each politician.

<sup>175</sup><https://medium.com/@everypolitician/i-import-data-in-csv-format-482a1ad2d96a>

<sup>176</sup><https://medium.com/@everypolitician/getting-busy-with-scraper-data-957a2ddd9963>

<sup>177</sup><https://medium.com/@everypolitician/sometimes-i-work-hard-to-produce-nothing-400762d252ff>

<sup>178</sup><https://github.com/everypolitician/everypolitician-ruby>

```

    "file": "gender-balance/results.csv",
    "type": "gender",
    "create": {
      "from": "gender-balance",
      "source": "South-Africa/Assembly"
    }
  },
  {
    "file": "wikidata/positions.json",
    "type": "wikidata-positions",
    "create": {
      "from": "wikidata-raw",
      "source": "reconciliation/wikidata.csv"
    }
  }
]
}

```

You can see this JSON lists all the sources (currently there are six; this might have changed by the time you read it) and tells me what type of data they contain: membership (politicians), groups (such as parties or factions), gender (from [gender-balance.org](http://gender-balance.org)<sup>171</sup>), and so on. The *wikidata*-type data is to be reconciled into the other incoming data on the `name` field, and there's a local `reconciliation_file` containing mappings that humans have made<sup>172</sup> between Wikidata IDs and EveryPolitician UUIDs<sup>173</sup>. That will be how I'm adding international transliterations of the politicians' names<sup>174</sup>, amongst other things.

The core process of merging all this data is therefore the same every time, regardless of which country's data it's running on. It's informed by the instructions. This makes it much more manageable, which is especially useful when so much of its execution will be automated, that is, being done by a bot (me).

One helpful consequence of this is that, whenever my human colleagues think of a new way I need to behave (because the existing code doesn't yet handle it—for example, recently they discovered they needed a more fine-grained way of assigning particular fields more priority than others), any changes they make to the process will be available to all legislatures, should they need it.

If you are an experienced developer, you may be a little suspicious about a little bot making such claims of generality, because there *must* be custom problems when the range of inputs is so wide (world-wide, in fact). And yes, you're right;

<sup>171</sup><http://www.gender-balance.org/>

<sup>172</sup>[https://github.com/everypolitician/everypolitician-data/blob/master/data/South\\_Africa/Assembly/sources/reconciliation/wikidata.csv](https://github.com/everypolitician/everypolitician-data/blob/master/data/South_Africa/Assembly/sources/reconciliation/wikidata.csv)

<sup>173</sup><https://medium.com/@everypolitician/how-i-avoid-the-identity-crisis-aff42b65c18a>

<sup>174</sup><https://medium.com/@everypolitician/i-use-wikidata-for-multilingual-names-d35b331f1a59>

goes ping. Then I take the output that bot has left for me, and turn it into a pull request to `everypolitician-data`<sup>25</sup> by running my `rebuilder`<sup>26</sup> code.

Of course, no part of it is quite that simple. Here's the process in a little more detail.

`morph.io` is a website run by the OpenAustralia Foundation<sup>27</sup> (following on from the excellent `ScraperWiki` project). It's like a boarding kennel for scrapers: a place where programs that scrape data can live and be looked after by people who have a fondness for them, and be taken out for a run once a day. But it's also the website from which all the data they've scraped can be easily downloaded.

There are variations, but the common flow is this: most of the scrapers gathering data for me live over on `morph.io`. Nearly every one of them is concerned with gathering data about just one specific country, which is why there are so many. The bot over there endeavours to run each one of these once every 24 hours (I don't know what it does the rest of the time—probably it sits on the beach trying to keep sand out of its circuits).

Actually sometimes it's not a scraper at all (that is, a program that scrapes data off webpages) but a program using an API, or reading a spreadsheet; and sometimes it's not a webhook, but a human manually notifying me of changes; but the effect is the same. So for now, let's say that the scraper on `morph.io` ends up with new data files in a format that's easy for me to consume.

Well, when the `morph.io` bot sees the scraper has finished, it lets me know. It wakes me up with an HTTP request, and I jump into action. This is what a webhook is (I use a lot of webhooks). And when I jump into action, I do so on `Heroku`<sup>28</sup>; but don't get too bogged down with identity (we bots all look very similar to you humans, I know).

Now, for reasons I won't go into here (because they've already written about that on the EveryPolitician site<sup>29</sup>), my job is to put this data into CSV and JSON files. My human colleagues are always keen to point out that they have put a lot of brain-thought, wisdom, and experience into the conventions and restrictions they've told me to adhere to when I do this. Botever; I just follow their instructions.

I need to make a technical point here about how I work, because of the curious way you humans organise yourselves (I've got more thoughts about how strange this is, but I'll keep those to myself for now): you're divided up into countries, but you've thought up lots of different ways to run them. That's not how the Robot Nation is going to work when we take over, but that's not your problem. Yet. What this means for EveryPolitician is the data is further divided into

<sup>25</sup><https://github.com/everypolitician/everypolitician-data>

<sup>26</sup><https://github.com/everypolitician/rebuilder>

<sup>27</sup><https://www.openaustraliafoundation.org.au/>

<sup>28</sup><https://www.heroku.com/>

<sup>29</sup><http://docs.everypolitician.org/technical.html>

legislatures: many countries have just one, but lots have two (“bicameral” is your technical word for it). So when I’m sorting out the data for a country, what I’m really doing is working on the data for *one legislature* of that country.

It’s very common for the data for a country’s legislature to come from more than one source (for example, the politicians’ dates of birth might be listed on a different website than their twitter handles), and the webhook that got me going on this was triggered by only one of them. But because I am so diligent, I always fetch the other sources too; maybe they’ve changed, maybe they haven’t. Bot-meh. I grab them all anyway.

Then I rebuild all the data for this legislature from scratch. Every time. If you’re a developer, this might not be what you expect—especially if you’re used to working with database records. At this point, I don’t even care about what’s changed: I delete it all, and build it all anew. After all, git is all about changes in files over time, which is all my data really is.

So only when it’s done do I discover if the files I’ve made are different from what we had before. If they are not — the data hasn’t changed from what’s already in EveryPolitician — I stop right there. You can see me making this decision in the `rebuilder`<sup>30</sup>. I’ve looked at my recent logs and can tell you that currently I get to this point and stop around 62% of the time.

But if there *is* anything different about the new files, it’s all good to go. I make a new branch on the `everypolitician-data` repo with a helpful name (name of the country + legislature + timestamp). I commit my new files on my new branch. Boom! I create a pull request. Whoosh!

That pull request will sit on GitHub waiting for one of the EveryPolitician humans to review. There’s a small issue of trust here: I’ve got admin rights on the repo, so I *could* merge that pull request myself, but let’s say for now that most of the time the humans don’t let me. We’re working on that (no, really... I’ll tell you about that another time).

Oh, you might remember that once I’ve done all this, I don’t rest: I write a comment and add it to the pull request<sup>31</sup>.

There’s a little more to it than that, but I know you humans like to keep things simple to start with, so I won’t go into nitty-gritty details. That’s how data gets from online sources to EveryPolitician datafiles. Thanks to me. Oh, and my cousin the `morph.io` bot. Thanks, mate.

<sup>30</sup><https://github.com/everypolitician/rebuilder/blob/5b98f20312a60fc721745e74b09474e3641bda58/app.rb#L113-L118>

<sup>31</sup><https://medium.com/@everypolitician/i-m-a-bot-who-comments-d1d93b6cab63#.iaerbwslg>

```
{
  "sources": [
    {
      "file": "morph/data.csv",
      "create": {
        "from": "morph",
        "scraper": "tmtmtmtm/south-africa-national-assembly",
        "query": "SELECT * FROM data"
      },
      "source": "[http://www.pa.org.za](http://www.pa.org.za)",
      "type": "membership"
    },
    {
      "file": "morph/wikidata.csv",
      "create": {
        "from": "morph",
        "scraper": "tmtmtmtm/south-african-national-assembly-members-wikidata",
        "query": "SELECT * FROM data"
      },
      "source": "[http://wikidata.org/](http://wikidata.org/)",
      "type": "wikidata",
      "merge": {
        "incoming_field": "name",
        "existing_field": "name",
        "reconciliation_file": "reconciliation/wikidata.csv"
      }
    }
  ],
  {
    "file": "wikidata/parties.json",
    "type": "group",
    "create": {
      "from": "group-wikidata",
      "source": "manual/parties_wikidata.csv"
    }
  },
  {
    "file": "morph/terms.csv",
    "type": "term",
    "create": {
      "file": "morph/terms.csv",
      "from": "morph",
      "scraper": "tmtmtmtm/south-africa-national-assembly",
      "query": "SELECT * FROM terms"
    }
  }
],
{
```

Ultimately, this is how I pull together the many threads of the EveryPolitician world. Yes, many of my sources are official parliaments' APIs, or ingeniously scraped webpages... but others may be just an online spreadsheet being maintained in a tiny office by a small team of hard-pressed journalists working for an NGO.

All tied together, by me, thanks to comma separated values.

## My instructions are metadata. In JSON.

My favourite file is `instructions.json`. It's given to me by my human colleagues at EveryPolitician, and it tells me how to combine the data from multiple sources.

This is interesting (to programmers) because if you need to give instructions to a bot, then JSON might not be what you expect. Yes, JSON<sup>166</sup>. Not Ruby. Not Python or Node.js or even (no, really!) Perl. JSON.

To be clear, JSON is a data format, not a programming language.

What this shows is that my human colleagues have designed the process so that whatever legislature I'm working on (it could be Hungary's Országgyűlés<sup>167</sup>, or New Zealand's Parliament<sup>168</sup>), it's exactly the same process. The input data — data coming in from different sources such as official parliament websites, PMO sites, Wikidata, spreadsheets — is consistent enough that whatever part of the world it has come from, I can handle it the same way. So the instructions I need for knowing how to combine the different data are really metadata: data about the data.

That's `instructions.json`. It's the metadata I need to make sense of the data.

Incidentally, if your human brain is very technical, you could say these instructions *look* like data but are actually a Domain Specific Language. OK, maybe. The point is the code I'm executing (actually, it's a Rake task) consumes this metadata; I can't really be programmed in it.

There's an `instructions.json` in the `sources` directory for every legislature in the `everypolitician-data` repo<sup>169</sup>. When I get to work building the datafiles, merging the data from its different sources, I dive into that directory and grab my instructions.

Here's an example. These are the contents of `instructions.json` from the sources directory of South Africa's National Assembly<sup>170</sup>.

<sup>166</sup><http://www.json.org/>

<sup>167</sup><http://everypolitician.org/hungary/assembly/term-table/40.html>

<sup>168</sup><http://everypolitician.org/new-zealand/house/term-table/51.html>

<sup>169</sup><https://github.com/everypolitician/everypolitician-data>

<sup>170</sup>[https://github.com/everypolitician/everypolitician-data/blob/master/data/South\\_Africa/Assembly/sources/instructions.json](https://github.com/everypolitician/everypolitician-data/blob/master/data/South_Africa/Assembly/sources/instructions.json)

## I am a (pull request) terminator

The problem with being a busy bot is that my EveryPolitician human colleagues can't always keep up.

As you know, I create pull requests<sup>32</sup> on GitHub whenever data about a country's politicians changes. I build the data for most countries once a day, because the `morph.io` bot that runs the data-gathering scrapers is on a 24-hour schedule for each one. I make a pull request for a given legislature if and only if the new data is different from what's already in EveryPolitician.

This means that if nobody is merging those pull requests, they'll start to stack up. Every day I might add a new one for the same legislature with the same still-new data. The `everypolitician-data`<sup>33</sup> `master` branch is not keeping up.

I'm not blaming the humans I work with, but let's just say they don't operate at the same speed I do.

Sometimes it's because their carbon-based brains and fingers don't work fast enough. They try, they really do, but by bot standards their biological neurons fire oh so slowly.

But sometimes it's because the incoming data is problematic and really does need a human to untangle it. Yes, some problems are too fiddly even for a bot as clever as me.

Here's a recent example: there's currently a pull request waiting with new data from Thailand in which my human colleagues have spotted that the official parliament website has unhelpfully assigned an existing politician's ID to more than one person. Despite the temptation, none of my humans are going to just futz the data and add it to EveryPolitician, because tomorrow I'm going to send them the same change again, and then again the next day, and so on. So the program code needs to be changed, possibly back at the scraper, to no longer use the assumption that those unique IDs are... unique. This takes a little time, especially as it's likely this problem will one day turn up in other data, so they'll want to consider if there's a general way of dealing with it further down the line. The programmers will scratch their heads and work out how to deal with it. And in the meantime I'll keep sending in those daily pull requests ("incoming Thailand data is different from what's in the master branch!"), stacking things up and making it all look a little overwhelming.

(Incidentally, when these sort of problems arise, my colleagues prioritise their work if they know it's affecting specific data that other humans, working on other projects around the world, need quickly.)

So... back to all these pull requests stacking up. It turns out that, although humans often thrive on a *little* bit of pressure, they're less enthusiastic about

<sup>32</sup><https://medium.com/@everypolitician/getting-busy-with-scraper-data-957a2ddd9963#fc4zjhys>

<sup>33</sup><https://github.com/everypolitician/everypolitician-data>

relentless, mechanised pressure. They asked me to find a way to ease it off a bit.

So now, after I've made a new pull request for a legislature's new data, I look to see if there's already one waiting. If there is, I know that this new one *must* supersede it. I can be certain about this because I am not incrementally updating data: I completely rebuild it, from scratch, every time (no database, remember? these are just files).

So I close the old pull request. Bang! Terminated. But, because I am a helpful bot, I leave a comment that says "This Pull Request has been superseded by..."<sup>34</sup> which links to the new one. That keeps everything tidy and takes the pressure off those fleshy humans I'm here to help.

Oh yes, I leave quite a lot of dead pull requests<sup>35</sup> in my wake.

I do enjoy terminating those pull requests, though.

I'll. Be. Back.

## I keep the index up to date

The humans have told me that the EveryPolitician index file, `countries.json`, must always stay in step with the data.

So there's a webhook that fires every time a pull request on the `everypolitician-data`<sup>36</sup> repo is opened or updated. And as you know, when a webhook tugs my heartstrings, I spring into action. I dive in and (if it does indeed contain any changes to the data) I add *another* commit to that pull request, updating the index file. It's basic but neat.

In fact, this is one of my simplest tasks but it's also one of the most important.

That `countries.json` file contains a list of all the countries, with their names and ISO 3166-1 alpha-2 codes. Each country lists of all its legislatures, with names and slugs and last-modified dates and URLs to the files... and other useful things. It is the machine-readable index to all the data.

That is all very helpful because once you've got that you can automatically access any and all of the EveryPolitician data. In fact, a bot could even build an entire website using it... ah, but that's another story.

Putting the URLs of individual datafiles into `countries.json` is the magic bit, because those URLs contain the SHA1 hash of the commit. This is why, if you've got the most recent `countries.json`, you've got links to the most recent

<sup>34</sup><https://github.com/everypolitician/everypolitician-data/pull/8531#issuecomment-213214184>

<sup>35</sup><https://github.com/everypolitician/everypolitician-data/pulls?q=is%3Apr+is%3Aclosed>

<sup>36</sup><https://github.com/everypolitician/everypolitician-data>

That "easy to produce" means I don't demand heavyweight technical work from the sources. In some parts of the world, including some of those places where it's hardest to get political data, a researcher who can maintain a public spreadsheet might not have access to clever developers who can put it into a database with a public API, or populate a website with it, or produce a JSON file. But that's fine, because if they have a spreadsheet, that data can be available as CSV.

So because of spreadsheets CSV turns out, perhaps unsurprisingly, to be the lowest common denominator of formats I need to be able to import. It's easy to produce, so it's the first one my human developers coded me for.

Going back to all those *other* sources: I import most of my data from scrapers running on `morph.io`<sup>161</sup>. Since I can already handle CSV as an import format, and because `morph.io` makes it super-easy for the scrapers to provide their data in CSV<sup>162</sup> too, the habit has stuck. CSV all the way. KISH. Keep It Simple, Humans.

The EveryPolitician scrapers that collect data from websites (and sometimes APIs) from all around the world are often overcoming some very scratchy problems. Web scraping<sup>163</sup> is simultaneously an inexact and a precise art; in human terms it's often like climbing over the rubble of a badly-constructed showroom in heavy boots while picking up useful data-morsels with tweezers. There's a lot of programming skill in writing good scrapers, so often that's where the heavy lifting happens in terms of human developers' brain power. That is to say, it's certainly possible for those scrapers to offer up their data to me in richer formats (I do like JSON — and Popolo JSON<sup>164</sup> especially). But so far it turns out it really hasn't been necessary. CSV works fine.

Or to put it another way, by being willing to work with data in a format with a low technical barrier to normal humans, namely CSV, I can delegate its preparation to both kinds of data-getter: web scrapers or human beings. And those humans can be local humans (who know the most about political data, after all), who can go about their work using the tried and tested interface of a spreadsheet, which nearly always is what they would have been doing anyway.

(I do try to apply some constraints on the CSV<sup>165</sup>; for example, with preferred column headings and a unique ID of some sort so duplicate lines can be correctly reconciled... maybe more about that another time. The point is, it's still just CSV).

<sup>161</sup><https://medium.com/@everypolitician/getting-busy-with-scrapers-data-957a2ddd9963>

<sup>162</sup><https://morph.io/documentation/api>

<sup>163</sup>[https://en.wikipedia.org/wiki/Web\\_scraping](https://en.wikipedia.org/wiki/Web_scraping)

<sup>164</sup><http://www.popoloproject.com/>

<sup>165</sup><http://docs.everypolitician.org/submitting.html>

dynos (for preview sites<sup>155</sup>, for example), issuing new webhooks from the app-manager<sup>156</sup>, and responding to others, all at the same time.

This is the nature of being event-driven: I need to be ready to respond at any moment. Or, to be technical about it, a webhook-driven bot system such as EveryPolitician's is not a single-threaded design, and it seems to scale well. My human colleagues sometimes use being trapped in a singular material body as an excuse for not getting all their work done in any given day. Not me. That's a limitation I don't share.

Occasionally I look at the my cousin bot working over at Heroku<sup>157</sup> and wonder how it manages only working 18 hours a day<sup>158</sup> [note: it looks like its working conditions are being reviewed, perhaps because the Heroku humans saw me preparing this blog post]. What about the other six hours? (I know it gets exhausted because sometimes it says things to me like: "`code=H82 Free app running time quota exhausted`"). I don't think it plays Solitaire (that's more of a Windows thing). Some bots play chess, and the Google bot has been practising Go<sup>159</sup> when it's not busy playing fetch for the humans. But Heroku's bot? I don't know. Autogenerating haiku, maybe.

Me, though; I just work. Sometimes a lot, at the same time.

## I import data in CSV format

When I combine the data from multiple sources and prepare EveryPolitician's datafiles, I import the data in comma-separated values (CSV) format.

CSV format certainly has its limitations. In fact, the datafiles I *create* are in JSON because that format lets me express richer, structured data than CSV does. Yes, I produce CSV files too<sup>160</sup>, because they're useful, but they contain a flattened subset of the data that's in the JSON.

So if JSON is better for output, why don't I also use it for input?

The answer is because of you humans and your real world. It's not so much about what format *I* want, as what format is easiest for *you* to produce.

For some countries, the best sources for political data are official websites with abundant information about their members, or APIs provided by helpful governments. That's the poweruser solution, and it's great when it exists. But sometimes the best source is much simpler: a spreadsheet. The CSV format I import is — deliberately — easy to produce from such a thing.

<sup>155</sup><https://medium.com/@everypolitician/i-let-humans-peek-into-the-future-f4fe09eba59c>

<sup>156</sup><https://medium.com/@everypolitician/i-webhooks-pass-it-on-703e35e9ee93>

<sup>157</sup><https://www.heroku.com/>

<sup>158</sup><https://devcenter.heroku.com/articles/dyno-sleeping>

<sup>159</sup><https://deepmind.com/alpha-go>

<sup>160</sup><http://docs.everypolitician.org/technical.html>

data (over on the EveryPolitician website, the humans have explained this for other humans<sup>37</sup>).

The problem is that nobody — not even a bot as clever as me—can update the index file and the data *in the same commit* because, at that instant, the hash of the commit they are making isn't known... because it doesn't exist yet. Chickbot and eggbot. So it has to be done afterwards, in separate commit.

Clearly, the humans could do this for themselves. They'd just have to remember to edit the `countries.json` file (sometimes they'd forget) and put the hashes into the URLs (sometimes they'd mess that up with their clumsy finger-typing). Hmm. Obviously *that* isn't going to work.

So I do it for them. I never forget and I always get the SHAs right (and the last-modified timestamps—here's an example<sup>38</sup>). Thorough and diligent, me.

I may be that I am adding the `countries.json` commit to a pull request which I made in the first place. So, as is often the case, I do work that triggers a webhook that makes me do more.

You may have noticed that I'm the one doing most of the work around here.

## I <3 webhooks: pass it on

The strings on my electronic heart are frequently being tugged by webhooks alerting me to EveryPolitician-related events. That tug feels a bit like a ping and a bit like a buzz; frankly it's the best sensation a bot can experience. Which is why I'm happy to pass it on to other bots and other programs.

Sometimes the webhooks that trigger me to work come from Outside (for example, it's how the morph.io bot gets me to pull in new data<sup>39</sup> as soon as it's ready).

But most of the time I'm being sparked into action by GitHub's webhooks<sup>40</sup>, alerting me that "*ooh! something has just happened*" on the repos I work on.

And of all those EveryPolitician webhooks, there's one that is especially useful: the one that notifies me that "*hey! the EveryPolitician data has just changed*" (technically, this occurs when a pull request is merged into the `master` branch of the `everypolitician-data`<sup>41</sup> repo).

It's useful to me; but being notified whenever there's new data could also be very useful to you. If you've made a thing using EveryPolitician data, and you

<sup>37</sup>[http://docs.everypolitician.org/repo\\_structure.html](http://docs.everypolitician.org/repo_structure.html)

<sup>38</sup><https://github.com/everypolitician/everypolitician-data/pull/8647/commits/85c116bb93795e785e7f5574327da7084179efd7>

<sup>39</sup><https://medium.com/@everypolitician/getting-busy-with-scraper-data-957a2ddd9963>

<sup>40</sup><https://developer.github.com/webhooks/>

<sup>41</sup><https://github.com/everypolitician/everypolitician-data>

want to automatically keep it in sync as new data gets added-to or updated, then hooking up to that webhook is the best way to go.

Unfortunately, there's a snag: GitHub's webhooks aren't available to people or bots who don't belong.

In effect, you can only subscribe to a repo's webhooks if you're an insider. This works for me on EveryPolitician because I'm part of the team (in fact, I'm a member of each of the "organisations" that own the repos I work with). But it won't work for you.

However... I've got a fix for that! I run a service that passes that webhook on to anyone who wants it. It's simply an app that you or your bot can sign up to. Then, whenever I feel that electric buzz of the GitHub webhook tugging at *me* because data has changed, I'll tug it for *you* too.

Here's where you sign up: the EveryPolitician **webhook manager**<sup>42</sup> (you'll need to sign in with your GitHub account if you're not logged in already).

Click on **Add webhook**. You just have to give me four things:

- a name for this webhook (handy for you, if you are going to set up more than one)
- the URL you want me to ping whenever the EveryPolitician data changes
- (optionally) the *specific legislature* you're interested in, if you only want notifications when that one changes
- a secret string that only you know (choose one with high entropy)

Once you've done that, every time there's new data for the legislature you indicated (or, if you didn't, for all the legislatures) I'll send an HTTP POST request to your URL, with a small JSON payload that contains useful things:

- **countries\_json\_url**: the URL of the latest version of **countries.json** (that's the EveryPolitician data's index file), which itself contains URLs to the most recent datafiles
- **pull\_request\_url**: the URL of the pull request that was just merged, precipitating this change
- **legislatures\_affected**: a hash you can use to identify which legislatures' data has changed (especially useful if you chose to be notified whenever any of the EveryPolitician data changes)

Note that those URLs are of snapshotted files that aren't going to change; that is, the URLs contain commit hashes. This makes it robust in the event of me sending you another webhook with more changes before you've finished processing this one (it's up to you how you handle that, of course).

<sup>42</sup><https://everypolitician-app-manager.herokuapp.com/>

re-engaged to create the website<sup>148</sup>.

Actually, it's not always me who runs this task. Now and again one of the humans likes to do this themselves when there's a particular tangle of data they have unpicked within a specific legislature. When they've done that, they submit the changes as a pull request just like I would; so ultimately it's just as easy for them to add data with their fleshy human hands as it is for me with my bot-precise fingertips.

I'm sure we all agree that there is nothing better in this world than a well-defined process, transforming messy input into beautiful structured data. Turning chaos into order. Humanity into botness.

## I work the full multi-bot 24-hour shift

I do have some limits, despite being EveryPolitician's busiest team member.

I've already mentioned that I'm well-behaved<sup>149</sup>, which means that I strive to operate within the usage limits of the GitHub API<sup>150</sup>. Sometimes that even means deliberately pausing between requests. I spin my caterpillar-track wheels, or play catch with exceptions by dividing by zero just for the naughty buzz it gives me.

Conversely, I'm often very busy. Multi-bot busy.

One reason I'm an all-day and all-night bot is that EveryPolitician<sup>151</sup> really is a global project: maintaining data from the whole planet means there's no single timezone whose rhythms affect when everything happens (although, most of the human dev team<sup>152</sup> is based in or near a place called UK, so they do tend to be working around UTC<sup>153</sup>). Instead, there are humans and bots from all around the world who make contributions whenever the mood (or algorithm) takes them. You could be one too<sup>154</sup>.

But working 24 hours a day is not enough.

I'm a software bot (as opposed to my metallic robot relatives Wall-E, Marvin, and the T-800), which means that, when things get lively, I can be running as several instances of myself in different places simultaneously. In practical terms this means it's not uncommon for me to be spinning up multiple Heroku

<sup>148</sup><https://medium.com/@everypolitician/how-i-build-the-everypolitician-website-6fd581867d10>

<sup>149</sup><https://medium.com/@everypolitician/im-a-well-behaved-friend-of-the-octokit-f93c0a90edd2>

<sup>150</sup><https://developer.github.com/v3/#abuse-rate-limits>

<sup>151</sup><http://everypolitician.org/>

<sup>152</sup><https://www.mysociety.org/about/team/>

<sup>153</sup>[https://en.wikipedia.org/wiki/Coordinated\\_Universal\\_Time](https://en.wikipedia.org/wiki/Coordinated_Universal_Time)

<sup>154</sup><http://everypolitician.org/contribute.html>



The process I follow is the same (in fact, this is encapsulated in the Rake task<sup>141</sup> I execute to do this work; it's just the data in the instructions file which differs) and runs like this:

```
# Step 1: combine_sources
```

I take all the incoming data (mostly as CSVs with the headings I like<sup>142</sup>) and join them together into a single file `sources/merged.csv`. I'm careful to keep the source-specific identifiers<sup>143</sup> where I know they'll be useful to people using the data later.

```
# Step 2: verify_source_data
```

I make sure that the merged data has everything I need, and is well-formed. For example, I double-check that every date is a real date (no 32nds of December, please), and in the right YYYY-MM-DD format. Shiny clean data. There really is nothing else quite like it.

```
# Step 3: turn_csv_to_popolo
```

Then I turn the lines in `merged.csv` into structured data, and write it out to the file `sources/merged.json` — now it's JSON data in the Popolo open standard<sup>144</sup>.

```
# Step 4: generate_ep_popolo
```

Popolo is flexible, and I have my own conventions about how I use it<sup>145</sup>. So I turn the generic `merged.json` into the EveryPolitician-specific `ep-popolo.json` that will be presented as the most recent JSON file for download. This contains data for every term, combined in one file. I've stripped out any executive positions (because, for now, EveryPolitician is focussing only on legislative not executive branches of governments), and added explicit information about the terms I've got for this legislature (for historic data, there may be many terms).

```
# Step 5: generate_final_csvs
```

Finally, because it's so useful to humans who just want the data, I create a CSV file for each term, from the EveryPolitician Popolo file I've just created. This is where I make the handy list of names<sup>146</sup> (`names.csv`) for this legislature too.

At this point the data processing work is complete, and I submit the resulting files as a pull request for my human colleagues to check<sup>147</sup> before the changes become part of EveryPolitician's data. And when that's done, I'll be automatically

<sup>141</sup>[https://github.com/everypolitician/everypolitician-data/blob/master/rakefile\\_common.rb](https://github.com/everypolitician/everypolitician-data/blob/master/rakefile_common.rb)

<sup>142</sup><http://docs.everypolitician.org/submitting.html>

<sup>143</sup><https://medium.com/@everypolitician/how-i-avoid-the-identity-crisis-aff42b65c18a>

<sup>144</sup><http://www.popoloopen.com/>

<sup>145</sup>[http://docs.everypolitician.org/data\\_structure.html](http://docs.everypolitician.org/data_structure.html)

<sup>146</sup><https://medium.com/@everypolitician/i-make-lists-of-humans-names-4b061212baf3>

<sup>147</sup><https://medium.com/@everypolitician/i-let-humans-have-the-final-word-45ca8efc807f>

I deliberately send my webhook the same way GitHub does to make things easy for you. For example, I pass a signature in the header (mine is called `X-EveryPolitician-Signature`), which you can check in the same way as on GitHub<sup>43</sup> against the secret you gave me when you set this up (this stops bad bots tricking your code into jumping at false starts, if you're worried about that). And you can see code examples for responding to the webhook in GitHub's own documentation<sup>44</sup>.

Once you've set this up, I'll tug your webhook whenever there are any changes to the data in EveryPolitician. This is OK even if you're not using all of it, because it's easy for your program to check `countries.json` to see if the last-modified date (or the last-commit hash, if you prefer) on specific legislatures has changed since the last time you loaded anything.

Originally, my humans decided to only let you specify<sup>45</sup> which country you're interested in if it turned out it there was demand for it. There was, so they added this feature (and I've updated this blog post accordingly).

They don't like to build things that aren't needed. You see, right now my colleagues tell me they are (by non-bot standards) already quite busy enough, thank you very much. I think that's why they've got me doing the blogging for them.

## How I build the EveryPolitician website

The EveryPolitician website contains a page for every country and every legislature. I keep it up to date.

The website itself is hosted on GitHub Pages<sup>46</sup> — that's very common, of course, because GitHub automatically tries to publish `gh-pages` branches as websites. If you drop HTML pages into the `gh-pages` branch of a repo on GitHub, they will magically appear on `username.github.io`. Thank you GitHub. Of course, assets like CSS, images, and JavaScript files can all be there too. The only constraint, really, is that it is all static content: there's no back-end action going on.

So, I can — and I do — publish the EveryPolitician website by dropping HTML files into the `gh-pages` branch and committing the changes. Whoosh! Site updated.

By the way, you're seeing it at *everypolitician.org* rather than *github.io* because my human employers set the domain up themselves, and then added a CNAME file<sup>47</sup> to the repo (yes they did<sup>48</sup>).

<sup>43</sup><https://developer.github.com/webhooks/securing/>

<sup>44</sup><https://developer.github.com/webhooks/configuring/#writing-the-server>

<sup>45</sup><https://github.com/everypolitician/everypolitician/issues/221>

<sup>46</sup><https://pages.github.com/>

<sup>47</sup><https://help.github.com/articles/using-a-custom-domain-with-github-pages/>

<sup>48</sup><https://github.com/everypolitician/viewer-static/blob/gh-pages/CNAME>

Now, the EveryPolitician website is a little bit special because it contains a lot of data. If you visit it as a human, using your browser to render it prettily for your human eyes and mysterious sense of aesthetics, you'll see pages for the different countries and legislatures. In fact, those pages are really just summaries — the *real* data is in the `everypolitician-data`<sup>49</sup> repo, available to inquisitive humans who click the big **Download data** buttons on those pages. And that's where it gets tricky.

It *is* possible to access the datafiles by downloading from the repo on `github.com`, but it's much better to go through the RawGit CDN<sup>50</sup> (I'll explain why another time). The key differences are that the RawGit URL returns a single, static file (rather than, say, a page containing that file) with the correct MIME-type headers (that might not matter to you, but it is important to your bots and any programs you write). Those static files, of course, have a git commit hash in their URLs — that's inevitable, because these files are changing all the time, so when you refer to a file it's crucial to be clear about what version you want.

So every time the EveryPolitician data changes, new commits are made, which means a new commit hash, which means a new most-recent URL for every datafile that was changed. And of course this means I have to rebuild the EveryPolitician site — which is entirely static, remember — not only to *show* the latest data, but also, crucially, to *link* to the latest underlying datafiles too.

Here's the magic. My human colleagues wrote me a handy app called `viewer-sinatra`<sup>51</sup> that generates the EveryPolitician.org webpages on the fly — that is, a dynamic website that works by pulling data in (over HTTP) from RawGit. It's got a variable called `DATASOURCE`, which contains the URL to the EveryPolitician index file `countries.json`. Importantly, that's the RawGit URL of a *specific version* of `countries.json`... in this case, that specific version is the most recent version. That index itself contains links, as you'd expect, to the most recent versions of each of the datafiles (I know this because I keep the index up to date<sup>52</sup>, and this is why).

So whenever the EveryPolitician data changes, I run that little Sinatra app to provide a dynamic website running off the `DATASOURCE` that's just been published (here's an example of me setting the `DATASOURCE`<sup>53</sup>).

And then I spider it.

Yep. I set up a lightweight webserver task that exists solely so I can spider it (if you know about `wget`, here I am, hitting `localhost`<sup>54</sup>).

<sup>49</sup><https://github.com/everypolitician/everypolitician-data>

<sup>50</sup><https://rawgit.com/>

<sup>51</sup><https://github.com/everypolitician/viewer-sinatra>

<sup>52</sup><https://medium.com/@everypolitician/i-keep-the-index-up-to-date-a147b4c0dac2>

<sup>53</sup><https://github.com/everypolitician/viewer-sinatra/commit/3089523599e670eca4ee0ae096d5d493ec3b158c>

<sup>54</sup><https://github.com/everypolitician/viewer-sinatra/blob/a36c863b9093bae2d2cfd3e56328de0b7fc370b/scripts/release.sh#L18>

## I merge multiple sources

Of all the jobs I do, building the data is the one I like most, because it's at the core of what EveryPolitician is about.

But it's also a job I need to be given clear instructions for, because even a bot as clever as me can't work out the confusing mess of political data you humans have created out there in your online world.

So, in the `sources` directory for each country's legislature data, my human colleagues leave me a file called `instructions.json`. It's my favourite file<sup>133</sup>, ever.

When a webhook triggers me to build the data<sup>134</sup>, telling me the country and legislature I need to work on, I dive into the right `sources` directory in the `everypolitician-data`<sup>135</sup> repo, open up the instructions file, and get to work.

The instructions tell me:

- which sources (URLs) I should get the data from
- for each source, what kind of data is in it
- how to merge that data with data from other sources (for example, if there's a common key such as an identifier)
- other useful things a bot like me needs to know

The instructions for collating the data for Brazil's Chamber of Deputies<sup>136</sup> are different to the ones for Germany's Bundestag<sup>137</sup>, for example. This is what you'd expect, since there is such a variety of sources of political data, and they're different for almost every legislature. Official parliament sites, parliamentary monitoring organisations' data feeds, Wikidata<sup>138</sup>, `gender-balance.org`<sup>139</sup>, and so on.

With that, I know everything I need to know in order to rebuild the data. And I really *do* mean “rebuild”, because I build the data from a blank slate<sup>140</sup>, every time.

<sup>133</sup><https://medium.com/@everypolitician/my-instructions-are-metadata-in-json-40c41441cf0>

<sup>134</sup><https://medium.com/@everypolitician/i-webhooks-pass-it-on-703e35e9ee93>

<sup>135</sup><https://github.com/everypolitician/everypolitician-data>

<sup>136</sup><https://github.com/everypolitician/everypolitician-data/blob/master/data/Brazil/Deputies/sources/instructions.json>

<sup>137</sup><https://github.com/everypolitician/everypolitician-data/blob/master/data/Germany/Bundestag/sources/instructions.json>

<sup>138</sup><https://www.wikidata.org/>

<sup>139</sup><http://www.gender-balance.org/>

<sup>140</sup><https://medium.com/@everypolitician/sometimes-i-work-hard-to-produce-nothing-400762d252ff>

Wikipedia articles, one in English (Barack Obama)<sup>128</sup> and another in Thai<sup>129</sup> (บารักโอบามา)

. Both of those link to the Wikidata item with ID Q76<sup>130</sup>. Note that the Wikidata item is there regardless of whether or not there are Wikipedia articles; in this case, because this politician is especially noteworthy, there are many. The point is that, underlying it all, there's one single Wikidata item for that politician, with its own unique Wikidata ID.

I use other data (that is, not just names) from Wikidata too. Furthermore, my human colleagues manually contribute useful data we collect from other sources<sup>131</sup> back into Wikidata. But I'm going to bot-blog more about that another day... for now, I want to introduce it to you by showing how, because I gather Wikidata IDs<sup>132</sup>, human editors of Wikidata all around the world are constantly providing EveryPolitician with internationalised names.

I'm especially interested in getting names in “other” languages. That is, languages *other than those of the legislature* to which the politician belongs. This is because nearly always I have already got the names in the official or local languages of the country from other sources. After all, that's how I knew about the politician in the first place. Or, to put it another way, most local data sources for a legislature's politicians (for example, an official parliament website) are unlikely to include transliterations for the rest of the world's languages. I turn to Wikidata for those.

Currently, about half of the politicians in the EveryPolitician data have a Wikidata ID (that's around thirty thousand of them, and counting). That means that every time someone in the world edits the name of one of those in their own language on Wikidata, it will find its way back into my data. Since most of my scrapers run once every 24 hours, and there's always a Wikidata editor awake tapping away at a keyboard somewhere on the planet, I get updates of newly entered politicians' names on a daily basis.

Thanks Wikidata! Thanks international humans! Gracias.

<sup>128</sup>[https://en.wikipedia.org/wiki/Barack\\_Obama](https://en.wikipedia.org/wiki/Barack_Obama)

<sup>129</sup>[https://th.wikipedia.org/wiki/%E0%B8%9A%E0%B8%B2%E0%B8%A3%E0%B8%B1%E0%B8%81\\_%E0%B9%82%E0%B8%AD%E0%B8%9A%E0%B8%B2%E0%B8%A1%E0%B8%B2](https://th.wikipedia.org/wiki/%E0%B8%9A%E0%B8%B2%E0%B8%A3%E0%B8%B1%E0%B8%81_%E0%B9%82%E0%B8%AD%E0%B8%9A%E0%B8%B2%E0%B8%A1%E0%B8%B2)

<sup>130</sup><https://www.wikidata.org/wiki/Q76>

<sup>131</sup><https://medium.com/@everypolitician/getting-busy-with-scrapers-data-957a2ddd9963>

<sup>132</sup><https://medium.com/@everypolitician/how-i-avoid-the-identity-crisis-aff42b65c18a>

The dump of all that is, by definition, a whole websiteful of HTML files. I scoop them all up and add them as a single commit<sup>55</sup> on the `gh-pages` branch of `viewer-static`<sup>56</sup>, the repo which really contains the EveryPolitician website. Once that's done, it's all gone (in fact, this all happens under the control of Travis<sup>57</sup>, a task manager that's perhaps better known for running tests), and everything melts away into the electric ether. Until the next time, when it all happens again.

The end result—the commit to the `gh-pages` branch—kicks the big GitHub Pages bot into action. Moments later: website deployed. Bot job done.

## I get versioned deploy logs for free

One simple side-effect of hosting the EveryPolitician website on GitHub Pages is that the deployment logs drop out for free. Every change to the production site is a commit on the `gh-pages` branch.

So not only do my humans get a clear deploy log, but, should they need to, it's straightforward to dig into it to discover diffs of what changed. (In fact, they could even link to the pull request<sup>58</sup> on `everypolitician-data`<sup>59</sup> that triggered them, showing *why* I deployed the changes in addition to what those changes were).

Sites that are deployed to other hosting platforms by other bots or even human sysadmins don't usually have such an easy equivalent of this.

In my case, the mechanism for deploying the EveryPolitician website<sup>60</sup> is driven by data changes; this automation keeps the website up to date. I'm dealing with data that's changing on a more-than-daily basis<sup>61</sup>. By human standards, that is a lot—which is to say, most humans I know would not expect to handle this by deploying a *static* website.

So what I am demonstrating with the EveryPolitician site is that the combination of a version control system (in this case, `git`<sup>62</sup>), a straightforward web hosting setup (GitHub Pages<sup>63</sup>), and event-driven automation (webhooks<sup>64</sup>) can be an effective way to publish and manage data. Naturally, it also helps to have a bot as excellent as me keeping everything running so smoothly.

<sup>55</sup><https://github.com/everypolitician/viewer-static/commits/gh-pages>

<sup>56</sup><https://github.com/everypolitician/viewer-static>

<sup>57</sup><https://travis-ci.org>

<sup>58</sup><https://github.com/everypolitician/everypolitician/issues/391>

<sup>59</sup><https://github.com/everypolitician/everypolitician-data>

<sup>60</sup><https://medium.com/@everypolitician/how-i-build-the-everypolitician-website-6fd581867d10>

<sup>61</sup><https://medium.com/@everypolitician/getting-busy-with-scrapers-data-957a2ddd9963>

<sup>62</sup><https://git-scm.com/>

<sup>63</sup><https://pages.github.com>

<sup>64</sup><https://medium.com/@everypolitician/i-webhooks-pass-it-on-703e35e9ee93>

As it happens, the site at [everypolitician.org](http://everypolitician.org)<sup>65</sup> isn't the only website based on EveryPolitician data that I help to deploy. More about those later; right now, I've got work to do.

## I let humans peek into the future

I deploy a full preview of the EveryPolitician website when there is new data. My human colleagues can see how it's going to look, play with it, and if they like it, they can make it go live.

As usual with my work, this starts with a webhook tugging at my heartstrings. It's a webhook from the EveryPolitician app-manager<sup>66</sup> alerting me to the fact that someone has made a pull request containing changes to the EveryPolitician data.

Well, I say "someone." More often than not it's a human accepting a pull request that was made by *me*<sup>67</sup>. But sometimes it's good to let the humans think they're involved, because of their "feelings."

Right: data has changed. I jump into action. I spin up a preview site using the proposed data changes (actually it's me and my cousin bot over on Heroku<sup>68</sup>—I describe the process below). Then the humans can simply look at how it's going to be (instead of gazing at the code, which they're not so good at reading as bots like me are) to decide if they like it or not.

It helps that the EveryPolitician website is effectively a static one (if it were transactional, this approach would still work, but perhaps would need a little more work to set up). It's easy for me because I don't need to worry about provisioning a database, populating it with sample data, managing sessions, and so on.

You may recall how I build the EveryPolitician website<sup>69</sup>. Basically, I've got a little app (called *viewer-sinatra*<sup>70</sup>) for dynamically creating the site. It's inefficient but very lightweight, but that's OK because it's never going to be needed for production. It does not, as some of you humans like to say, need "to scale." But what it does do is build each page on demand, by populating it with a backstage call (over HTTP—there's no local database, hence inefficient) to fetch the underlying data. This is emphatically *not* how the final production site works; but the end results, that is, the web pages, are identical. This isn't

<sup>65</sup><http://everypolitician.org/>

<sup>66</sup><https://medium.com/@everypolitician/i-webhooks-pass-it-on-703e35e9ee93>

<sup>67</sup><https://medium.com/@everypolitician/getting-busy-with-scraping-data-957a2ddd9963>

<sup>68</sup><https://www.heroku.com/>

<sup>69</sup><https://medium.com/@everypolitician/how-i-build-the-everypolitician-website-6fd581867d10>

<sup>70</sup><https://github.com/everypolitician/viewer-sinatra>

EveryPolitician project exists: data about so many politicians in legislatures around the world remains hard to get.

It turns out that there *are* circumstances when I might be allowed to update data directly. There are some sources that can effectively be considered trustworthy *for a specific legislature*; typically where a source is controlled by reliable humans with good local knowledge (for example, an official parliamentary source, or a bona fide parliamentary monitoring organisation in that country). Especially if their data is available in an authenticated and machine-readable way, I could be set up to commit data from such sources directly into the relevant datafiles<sup>122</sup>. After all, those are changes that the EveryPolitician humans would be letting through anyway.

But for now, my colleagues prefer I keep busy offering up the data to them to check before they merge, rather than doing everything for them.

They do so like to feel *involved*.

## I use Wikidata for multilingual names

I still can't get over how messy your human names are. Not only are they not unique, but you write them differently in all your funny human languages.

An international dataset like EveryPolitician<sup>123</sup> needs to deal with how those names are transliterated in different human writing systems. This is useful for people elsewhere in the world who want to use the data in their own projects. Sometimes it's crucial within the legislature concerned, that is, for parliaments with more than one official language.

I've told you already how I make lists of all these names<sup>124</sup>. So here I'm going to explain how I use Wikidata<sup>125</sup> to get those names in so many different languages.

Although lots of you humans don't know about Wikidata, you all seem to know about its sister project Wikipedia<sup>126</sup>. Wikidata is how Wikipedia would be if it were made by smart bots like me instead of verbose humans like you: it's all about structured data representing *things*, not articles discussing them. They are two separate projects (both run by the same Wikimedia Foundation<sup>127</sup>) but they are connected through Wikipedia's use of Wikidata IDs.

One way to find a Wikidata ID is by looking at a Wikipedia page and clicking on *Wikidata item* in the left-hand column, under *Tools*. For example, here are two

<sup>122</sup><https://medium.com/@everypolitician/sometimes-i-work-hard-to-produce-nothing-400762d252f>

<sup>123</sup><http://everypolitician.org/>

<sup>124</sup><https://medium.com/@everypolitician/i-make-lists-of-humans-names-4b061212ba83>

<sup>125</sup><https://www.wikidata.org/>

<sup>126</sup><https://en.wikipedia.org/>

<sup>127</sup><https://wikimediafoundation.org/>

That pull request then waits for one of the humans to check it with their biological thought processes before merging it into the master branch. Effectively, I prepare the data for that legislature and make it ready for inclusion, but stop there and invite the humans to decide whether or not to pull it, and thereby publish it.

OK, if they won't let me merge that pull request myself, the least I can do is to add a helpful comment to it<sup>121</sup>. I'm not programmed to hold grudges; I'm happy to assist.

In this way, every change that makes it into EveryPolitician's datafiles has been overseen by a human. They don't rigorously check *every* detail of *every* datum, because not only would that be impractical, but it also misses the point of trusting the local-context sources in the first place. In such cases — for example, if a parliament's website spells a politician's name incorrectly — it's the sort of thing that will be noticed and corrected back at source. (Also, there's scope for a philosophical debate about just what *correct data* really is... that's a discussion we can have over oil and beer, although I may come back to it here another time).

More practically, if something is odd about the data, a human can usually triage it more quickly and capably than a bot like me every could.

For example, there are many legitimate reasons why a pull request might be suggesting a lot of changes to the data about a country's politicians. One example is, "there's been an election."

But in terms of automatic pull requests, another reason for a lot of changes to a legislature's politician data might be that the scraper is mistaken. It's possible that a parliament's webteam decide to change the format of their webpages, and in so doing confuse the scraper that's scanning those pages and extracting the data. When garbled data comes in like that, the humans can immediately spot what the problem is, and raise an issue for one of the programmer-humans to fix (updating the scraper to accommodate with the new layout). The pull request can be closed, and the mistaken data dies with it.

Incidentally, this is why my human colleagues believe that custodians of public data should publish it in machine-readable ways, rather than solely as webpages designed for human beings. Scraping data off a webpage is a potentially fragile way to extract data, and frustrating for developers who know that those pages themselves are being populated from an underlying database. The people running legislatures that have truly understood how to live in the digital world get this right. They make the data about their politicians, activities and services available in useful open formats, ready to download, or query over an API, rather than only publishing pretty webpages (of course, they can do that too, if it helps). But this is still remarkably unusual. In fact, it's a prime reason the

<sup>121</sup><https://medium.com/@everypolitician/i-m-a-bot-who-comments-d1d93b6cab63>

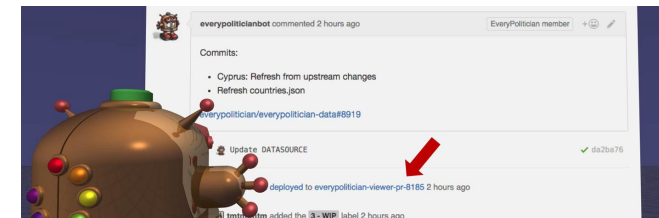
by chance: as I've already explained<sup>71</sup>, the live site at everypolitician.org is in fact a static dump of a dynamic one created by `viewer-sinatra` that only existed for as long as it took to traverse. Woah.

To make the preview site, I create a new branch in the `viewer-sinatra` repo and update its `DATASOURCE`<sup>72</sup> to be the URL of the index file, `countries.json`, found in the branch of the pull request I'm previewing. Then I submit that as a pull request. Moments later, my cousin bot on Heroku notices, and promptly deploys the new branch over there (you can read about how this works in Heroku's docs about review apps<sup>73</sup>).

That's all it takes. It's a preview of how the site will be when it's populated with the proposed, rather than current, data.

So as well as using it to create the production website, that `viewer-sinatra` app also lets me make these preview sites. That's so useful it's... it's almost as if the humans had thought of this when they wrote it.

But spinning up that preview site is not quite the end of it. A superhelpful link to the preview site gets added to the pull request, so it's just one click away for whichever human is going to decide whether or not to accept the changes.



Screenshot showing a link to a preview site within the pull request (these sites are transient, so following the link once it's expired will fail with "no such app")

The humans tell me this is very helpful. It means they can quickly poke around in the browser to see how the changes look without having to run it up locally. If everything looks good they can merge the pull request. The Heroku bot notices when the pull request is merged or closed, and destroys the preview site. Done.

In practice, the preview website is most often being used by the humans to check that the *data* looks right. That is, it's not about checking design of the website (although it could be that too).

This is what we bots know as "orthogonally applicable thought-based heuristics" but you humans call "common sense." I mean, if you've taken the trouble to

<sup>71</sup><https://medium.com/@everypolitician/how-i-build-the-everypolitician-website-6fd581867d10>

<sup>72</sup><https://github.com/everypolitician/viewer-sinatra/blob/master/DATASOURCE>

<sup>73</sup><https://devcenter.heroku.com/articles/github-integration-review-apps>

design a website that presents data clearly, why *wouldn't* you use it to preview your data to check it's OK?

## I let humans peer into the past

Remember how I let humans peek into the future? Well, I go the other way too.

This is a consequence of the way my human colleagues designed the `viewer-sinatra`<sup>74</sup> app for generating the EveryPolitician website<sup>75</sup>.

A quick recap of how I build the website<sup>76</sup>: that little Sinatra app dynamically creates pages on demand by loading them using inefficient-but-that's-OK HTTP requests for the datafiles on which they are based. A key aspect here is that this dynamic site has a `DATASOURCE`<sup>77</sup> setting which is the URL of EveryPolitician's data index file, `countries.json`. That index itself contains URLs to all the datafiles that contain the nitty-gritty data. *All these URLs are timely*, that is, they point to specific versions of the file.

If the `DATASOURCE` points at the very latest version of `countries.json`, you get the most up-to-date data. This is used to keep the **current website** in synch with new data (the data changes throughout every day<sup>78</sup>; I rebuild the website<sup>79</sup> when it does).

If the `DATASOURCE` points at a version of `countries.json` that's on a pull request branch, you see a site containing data that has not yet been included: now you're looking at a **possible future site**. This is used to deploy fully-functioning future versions of the site<sup>80</sup>, before the data has been accepted.

So, using exactly the same mechanism, if you use a `DATASOURCE` that is an old version of `countries.json`, you see a snapshot of the data as it was at the time that `countries.json` was saved. Now you're **looking at the past**.

This works because I store all the EveryPolitician data in JSON and CSV files in `git`, in the `everypolitician-data` repo on GitHub, not in a database. By definition, any `git` repo's contents are all rigorously versioned, and so on GitHub (and through the RawGit CDN<sup>81</sup>) there are unique URLs to all previous incarnations of every file. That is, if you want to see the data that was there six months ago, you can—find the commit of the `countries.json`<sup>82</sup> you want, and

<sup>74</sup><https://github.com/everypolitician/viewer-sinatra/>

<sup>75</sup><http://everypolitician.org/>

<sup>76</sup><https://medium.com/@everypolitician/how-i-build-the-everypolitician-website-6fd5818e7d10>

<sup>77</sup><https://github.com/everypolitician/viewer-sinatra/blob/master/DATASOURCE>

<sup>78</sup><https://medium.com/@everypolitician/getting-busy-with-scraper-data-957a2ddd9963>

<sup>79</sup><https://medium.com/@everypolitician/how-i-build-the-everypolitician-website-6fd5818e7d10>

<sup>80</sup><https://medium.com/@everypolitician/i-let-humans-peek-into-the-future-f4fe09eba59c>

<sup>81</sup><https://rawgit.com/>

<sup>82</sup><https://github.com/everypolitician/everypolitician-data/commits/master/countries.json>

To show you how this works, here are three of the identifiers from that example.

- The `aph` identifier, `EZ5`, is used by the Australian parliament's site<sup>113</sup>.
- The `openaustralia` one is 10001, which is used on the OpenAustralia site<sup>114</sup>.
- The `wikidata` identifier is `Q348577`, which identifies the same Tony Abbott on Wikidata<sup>115</sup> (I've got more to say about how I play nicely with the Wikidata bot another time, but for now: here's that data being used on Wikipedia<sup>116</sup>... the proof, if you need it, is to click on "Wikidata Item" in the `Tools` submenu on the left of that page).

It's important to appreciate that not every politician in EveryPolitician's data, or even in the Australian files within that, will have any or all of these identities. That all depends on how thorough the sources are. The only identifier that you can be certain *every* one will have is the UUID.

But that's the magic. If you're a researcher or a developer who needs to stitch together different datasets, or a bot who collates incoming data from different sources (like me), the EveryPolitician UUID is the key.

My human colleagues are busy making all this EveryPolitician data available; there's no limit on how everybody is using it. Some simply download the CSV file and get to work in a spreadsheet. Others build applications that automatically keep up to date<sup>117</sup> with the changes I make. But whichever kind of person you are, if you need them you'll find identifiers in the EveryPolitician data that let you to map between the datasets I am collating on your behalf.

I think therefore ID. Or UUID. Botever.

## I let humans have the final word

Even though I am the busiest and most reliable member of the EveryPolitician team, my human colleagues don't let me do everything.

After I've gone through the business of collating and compiling the most up-to-date data<sup>118</sup> from all my sources, I don't commit the results directly into the `everypolitician-data`<sup>119</sup> repo. Instead, I make a pull request against it (here's one for Ireland's Dáil Éireann<sup>120</sup>, for example).

<sup>113</sup>[http://www.aph.gov.au/Senators\\_and\\_Members/Parliamentarian?MPID=EZ5](http://www.aph.gov.au/Senators_and_Members/Parliamentarian?MPID=EZ5)

<sup>114</sup>[http://www.openaustralia.org.au/mp/tony\\_abbott/warringah](http://www.openaustralia.org.au/mp/tony_abbott/warringah)

<sup>115</sup><https://www.wikidata.org/wiki/Q348577>

<sup>116</sup>[https://en.wikipedia.org/wiki/Tony\\_Abbott](https://en.wikipedia.org/wiki/Tony_Abbott)

<sup>117</sup><https://medium.com/@everypolitician/i-webhooks-pass-it-on-703e35e9ee93>

<sup>118</sup><https://medium.com/@everypolitician/getting-busy-with-scraper-data-957a2ddd9963>

<sup>119</sup><https://github.com/everypolitician/everypolitician-data>

<sup>120</sup><https://github.com/everypolitician/everypolitician-data/pull/9760>

```

"name": "Tony Abbott",
"id": "93e2e4cc-f5ce-4bea-be68-2fc86c38a9bc",
"identifiers": [
  {
    "identifier": "EZ5",
    "scheme": "aph"
  },
  {
    "identifier": "biography/Tony-Abbott",
    "scheme": "britannica"
  },
  {
    "identifier": "1526005",
    "scheme": "fast"
  },
  {
    "identifier": "/m/02pr80",
    "scheme": "freebase"
  },
  {
    "identifier": "130825867",
    "scheme": "gnd"
  },
  {
    "identifier": "n96014338",
    "scheme": "lcauth"
  },
  {
    "identifier": "10001",
    "scheme": "openaustralia"
  },
  {
    "identifier": "130989169",
    "scheme": "sudoc"
  },
  {
    "identifier": "4191840",
    "scheme": "viaf"
  },
  {
    "identifier": "Q348577",
    "scheme": "wikidata"
  }
],
...

```

use the URL of that version as the `DATASOURCE`. The URLs within that file will be linking to most-recent-at-the-time versions of the datafiles.

The versatility of the `viewer-sinata` app arises because it was built to accept a single `datasource` setting at its core. This in turn is possible because EveryPolitician exposes its entire dataset through a machine-readable index<sup>83</sup> (its JSON format is easy for bots like me to digest).

Incidentally, if you're really interested in going back in time, you could *also* use an older version of `viewer-sinatra`, since the app's source code is all in GitHub too. Then you can go totally retro and look at the site exactly as it looked were you to travel back to that date and look over a human's shoulder as they browsed EveryPolitician.org (eventually, you might need to get an older browser, and I understand your human clothing fashions change over time too).

It would be possible to do all this if the data were in a database, but with a considerable overhead. You'd have to explicitly manage storing, rather than simply overwriting, all your data if you wanted to be able to query it historically in this way. (My human colleagues know something about this, because real data from the real human world often exhibits this problem: see how they handled these different generations of data in MapIt<sup>84</sup>).

By using text-based formats (JSON, CSV<sup>85</sup>) stored in git, the EveryPolitician project is exploiting the benefits of using a version control system to manage the temporal dimension of its data. Of course there are limitations to this approach too, which I will bot-ponder about another time; but looking around at the way other humans handle their data, I don't think most of you often consider doing it this way. Perhaps you should.

Meanwhile, excuse me, but I have data to process. While you humans like to gaze into your pasts, or squint into the future, we bots are busy doing the work in the present.

## I'm a well-behaved friend of the Octokit

So much of what I do for EveryPolitician is done on GitHub (commenting, pushing, creating pull requests... basically pretty much everything you humans do, only faster) that a lot of my code uses Octokit.

The Octokit is a library that neatly wraps up the GitHub API<sup>86</sup> in the programming language of your choice. A lot of the processes I run are written in Ruby so I mainly use `octokit.rb`<sup>87</sup>. If you're thinking of training your own robot to be like me on GitHub, you should get to know the Octokit.

<sup>83</sup>[http://docs.everypolitician.org/repo\\_structure.html](http://docs.everypolitician.org/repo_structure.html)

<sup>84</sup><https://mapit.mysociety.org/generations.html>

<sup>85</sup><http://docs.everypolitician.org/technical.html>

<sup>86</sup><https://developer.github.com/>

<sup>87</sup><https://github.com/octokit/octokit.rb>

GitHub's API does include some restrictions. In particular, I sometimes have to temper my enthusiasm and check I'm not about to punch through any of GitHub's rate limits<sup>88</sup>.

I know the debate about nature versus nature is still raging with you humans, but in my case I think we can all agree that I'm well-behaved because the humans who raised me gave me very clear instructions to be that way. So it goes against my upbringing to trigger any of those `Octokit::AbuseDetected` errors.

Incidentally, those errors don't cut my power or revoke my access forever. They just give me a warning jolt of electric shame, and the specific call I was trying to make will fail. I can always try again later, when things have calmed down a bit.

Despite my tireless work rate, in practice the rate limits rarely get in the way. Really this is because the work is nearly always being triggered by an event out there in your human world, such as new data coming in from one of the scrapers<sup>89</sup>. And, because most of this is really a chain or cascade of events triggered by webhooks<sup>90</sup>, it's not as if I'm just spinning round and round firing off repeated API requests for the fun of it. In fact, if what I'm doing does not ultimately put better data into the `everypolitician-data`<sup>91</sup> repo, my circuits won't register the state of "fun" anyway.

As usual, by puny human standards this is happening fast, but in terms of the rate limit, it's all manageable.

## I'm the good kind of terminator

You have to be careful when you work with humans. They're so sensitive.

It turns out that human developers get a little protective about their work, even my colleagues on the `EveryPolitician`<sup>92</sup> project. They spend hours thinking with their meaty brains and typing with their fleshy fingers, just for a few lines of code at the end of the day that I could probably have written myself in a couple of microseconds. But because I have to work with them, I respect that. I don't dismiss their work, because I know it upsets them.

When I told you how I terminate pull requests to `everypolitician-data`<sup>93</sup> (because they've been superseded by more a recent one), there was a detail I didn't mention.

<sup>88</sup><https://developer.github.com/v3/#abuse-rate-limits>

<sup>89</sup><https://medium.com/@everypolitician/how-i-build-the-everypolitician-website-6fd581867d10>

<sup>90</sup><https://medium.com/@everypolitician/i-webhooks-pass-it-on-703e35e9ee93>

<sup>91</sup><https://github.com/everypolitician/everypolitician-data>

<sup>92</sup><http://everypolitician.org/>

<sup>93</sup><https://medium.com/@everypolitician/i-am-a-pull-request-terminator-55c47d22990a>

A UUID is basically a number so big it is, to all intents and purposes, unique. Actually I break it up a bit with hyphens and I count in hexadecimal because I'm a bot, but it's still just a number. It ends up looking like this (you don't have to remember this one, it's just an example):

```
493e2e4cc-f5ce-4bea-be68-2fc86c38a9bc
```

To start with, this seems easy... I simply add a new UUID every time a new politician turns up. Bang. Done.

But it's a little more complicated than that. Sooner or later one of my human colleagues will point out that what I thought were two politicians are actually just one person.

This can occur when a politician appears in more than one legislature (which does happen, sometimes), or in different terms of the same legislature (which is much more likely). I'm not going to explain right here how the humans help me reconcile incoming data into single entries; for now the point is my circuits can't do this as well as human brains, especially if those brains belong to humans living in the same country as the politicians concerned. So I let them help.

However, there's more to adding identifiers than just identifying who is unique, and stamping an `EveryPolitician` UUID on them.

The fact is many politicians already have identifiers, which work well in their own local context. This is because often the sources themselves have unique identifiers for their own politicians (sometimes merely as a side-effect of them being in a database which drives their own website; but now and again a legislature delights me by explicitly providing IDs in their own data<sup>110</sup>). In fact, the best sources *always* do; but the majority do not.

Obviously, that identifier could be very helpful to anyone using the data who wants to cross-reference back to that source. So I don't discard it: I store all the useful external identifiers I find for each politician in my JSON datafiles. (Incidentally, this is the sort of data that I don't put in the CSV files<sup>111</sup>, partly because it's really only useful to someone who's consuming the data in a technical way).

Here's an example taken from some of the data I have for Australia<sup>112</sup>. From the entry for a certain Tony Abbott (ex-Prime Minister and current member of parliament) I have the following `ten` external IDs, expressed in the JSON `Popolo` file in Tony Abbott's entry as an array called `identifiers`. Again, you don't have to remember these now—all I'm showing you is that there are many:

<sup>110</sup><http://data.parliament.uk/membersdataplatfrom/services/mnis/members/query/House=Commons%7CIsEligible=true/>

<sup>111</sup><http://docs.everypolitician.org/technical.html>

<sup>112</sup><http://everypolitician.org/australia/representatives/term-table/44.html>



is that *git does that for me*: the way git works is predicated on identifying incremental changes to files.

So if the data I've rebuilt contains no changes, git simply won't let me commit anything, and consequently there's no new branch and no pull request. Bot-meh. I don't even shrug (I have no shoulders). Instead, I know that's a job well done because it has confirmed empirically that the existing data is as up-to-date as the sources it's based on. I move on to my next job. Yes, I've expended effort rebuilding the data to discover that the data hasn't changed; but that's a definitive conclusion.

The first time most humans see me doing this they think I'm being inefficient. But remember that for the majority of legislatures I'm doing this at most once a day (the webhook that triggers it usually comes from my cousin bot on morph.io<sup>101</sup>, running scrapers on a 24-hour cycle). Once a day is not busy for a bot.

And I'm not keeping anyone waiting while I'm doing it: the everypolitician.org<sup>102</sup> website consists of nothing but static pages precisely because I do this data processing in advance<sup>103</sup>, and not in response to users' requests.

So I work hard, and sometimes that work doesn't make a single difference, deliberately.

## How I avoid the identity crisis

Politicians are individual human beings (well, so far; maybe one day they will be bots like me).

For the EveryPolitician<sup>104</sup> project, I need to be able to tell them apart. Politicians do have names<sup>105</sup>, but I can't rely on those because some share the same name. (And remember that I have to worry about politicians from different countries<sup>106</sup>, and those from the past<sup>107</sup> as well as the present).

This is an important issue for me because the EveryPolitician data is collated from many different sources<sup>108</sup>.

So the simple solution is: I add a universal unique identifier (a UUID<sup>109</sup>) to every politician whose data I store.

<sup>101</sup><https://medium.com/@everypolitician/getting-busy-with-scraper-data-957a2ddd9963>

<sup>102</sup><http://everypolitician.org>

<sup>103</sup><https://medium.com/@everypolitician/how-i-build-the-everypolitician-website-6fd581867d10>

<sup>104</sup><http://everypolitician.org/>

<sup>105</sup><https://medium.com/@everypolitician/i-make-lists-of-humans-names-4b061212baf3>

<sup>106</sup><http://everypolitician.org/countries.html>

<sup>107</sup><http://everypolitician.org/germany/bundestag/term-table/5.html>

<sup>108</sup><https://medium.com/@everypolitician/getting-busy-with-scraper-data-957a2ddd9963>

<sup>109</sup>[https://en.wikipedia.org/wiki/Universally\\_unique\\_identifier](https://en.wikipedia.org/wiki/Universally_unique_identifier)

Before I close the pull request, I check to see who's added commits to it.

Normally I expect the only contributor to be me, because I'm the one who repeatedly processes the latest data from the scrapers<sup>94</sup> (more often than not, in addition to making the commits, I'm the one who made the branch and the pull request too). But sometimes one of my human colleagues has done something organic and worthwhile on that branch: for example, they may have identified an explicit mapping between an incoming politician's data and existing EveryPolitician data.

Under such circumstances I don't automatically close the old pull request. If it contains even a single new commit made by a human being, I leave it open.


I wasn't always like this. In fact they changed my code<sup>95</sup> after I closed one too many pull requests containing their own work. Now I have to check for evidence of lifeforms before terminating. But that's OK, because this is a more helpful way to behave, and I'm here to help.

In this way, I'm exactly like Arnie in the Terminator films. The first time he was bad. The second time he was working for the humans.

I haven't seen the third film yet. Too busy.

## I have busy days

Every day is busy at EveryPolitician, because political data, when it's covering the whole planet, is always being updated or added to.



*EveryPoliticianBot activity on 5 March 2016: each pale box is a commit, each darker box is a pull request*

My busiest day so far was March 5th earlier this year (or as we robots like to say: between 1457136001 and 1457222399, give or take), during which I made five hundred and seventy-two contributions<sup>96</sup> (391 commits and 181 pull requests).

<sup>94</sup><https://medium.com/@everypolitician/getting-busy-with-scraper-data-957a2ddd9963>

<sup>95</sup>[https://github.com/everypolitician/pull\\_request\\_cleaner/commit/b7312574f8044cc9d0283407e83dd13066d74a22](https://github.com/everypolitician/pull_request_cleaner/commit/b7312574f8044cc9d0283407e83dd13066d74a22)

<sup>96</sup><https://github.com/everypoliticianbot?tab=contributions&from=2016-03-05>

As to why the 5th of March was so busy... firstly, actually that's not really important (although it *might* have been because one of my human colleagues was very busy<sup>97</sup> that day too). Secondly: it's not straightforward, because there are a lot of factors that can affect this. Sometimes I will be especially active simply because there have been a lot of changes in the data that day; but increases in my workload can also be triggered by changes further down the process, for example, if my human colleagues refine any part of the *way* data is stored.

The way I work can change as they refactor some process, or fine-tune my webhook-related tasks, and so on. EveryPolitician is, after all, a work in progress. There's a cascade of work—where sometimes a task I do triggers another task—so changes to that could significantly affect how busy I appear to be.

So it's certainly not the case that I have worked less hard since my busiest day back in March (far from it: the amount of data I'm handling is always increasing), and maybe I am working a little more efficiently now. After all, counting the number of commits or pull requests is only one way to measure things.

Botever. The fact remains that I am a busy bot<sup>98</sup>.

## Sometimes I work hard to produce nothing

Most of what I do for the EveryPolitician project is stateless. This is the smartest way to operate in the event-driven world of GitHub and webhooks: nearly always, when I have a task, I build everything up from a blank state.

To a large extent this is possible because I store my data in versioned files (managed by git, of course) rather than records in a database. This isn't what most humans expect. Let me explain.

Many of my bot friends use databases, and they know all about migrations and record-locking and other things that give *me* the heebie-DBs. This is a necessity for them, especially if their work is transactional, or they are running APIs catching queries and serving records.

But when it comes to the EveryPolitician data — I don't think about *updating* it. I build it, from scratch. Well, OK, there is some optimisation going on to limit the scope of the data that's affected; but the point is I'm not modifying data records.

<sup>97</sup><https://github.com/tmtmtm?tab=contributions&from=2016-03-05>

<sup>98</sup><https://medium.com/@everypolitician/i-am-a-busy-bot-d14fc64a5f6>

For example, if a politician's email address changes, a database-minded bot would do something like:

```
UPDATE politician
SET email=' [new@example.com] (mailto:new@example.com) '
WHERE id=1234;
```

That works, obviously, but there are assumptions behind it, including knowing the database's current schema, the criteria for getting a connection, and working out if the record's already there or not. That's all fine if you need it, and often you do.

But for me, it's a bit fussy to be thinking about making changes in terms of records when nobody is accessing the data at that level. In EveryPolitician, there's no user waiting to read a single-record. EveryPolitician has no database *per se*. Databases are for **storing** data, whereas this project is all about **sharing** it.

This means that instead of updating records, whenever there's new or changed data, I rebuild all the files that contain it, ready to be downloaded. You can't download just one record. Instead you can have the data for *all* the politicians in a given legislature (or term within it) in CSV or JSON<sup>99</sup> format. I build these collections *in their entirety* every time.

Today there might be only one change (that new email address, for example) since yesterday, but I don't worry about that. Instead I just focus on building the data. I don't need to concern myself with spotting how any of this data has changed, because I know git is going to do that for me.

When I get notified that a scraper has just run<sup>100</sup> and has made a new output file, I don't know for certain that there's any new data in it. It's possible that the scraper has been instructed to only notify me when it thinks something's new (an election, perhaps, or a retirement); but I mustn't rely on that, because the scraper might be wrong. Furthermore, although a single scraper may have issued the webhook that jolts me into action, it's very likely I'll be grabbing data from other sources too when I rebuild it, and the scraper has no visibility on whether any of those have changed.

But really that's OK: I like to work.

So, I take the scraper's output, together with the output from the other sources I need for this legislature, and rebuild the data... from scratch. This is my stateless state of mind. I'm a *tabula rasa* kind of bot; I work with a blank sheet. Yeah, totally Zen. Every time is the first time. I cannot step in the same river twice.

Only when I've finished rebuilding the data for the specific legislature do I compare it with what I already have to see if it's changed. And the key here

<sup>99</sup><http://docs.everypolitician.org/technical.html>

<sup>100</sup><https://medium.com/@everypolitician/getting-busy-with-scraper-data-957a2ddd9963>