

Open Data for *Crime and Place* Research: A Practical Guide in R.

Samuel Langton†

Reka Solymosi‡

This is an unedited draft, not for citation, of a book chapter that has been accepted for publication in E. Groff and C. Haberman (Eds.), *The study of crime and place: A methods handbook* (Temple University Press, 2021). Temple University Press retains all rights and subrights of this document. This pre-print is not published under a CC license.

Keywords: crime, open data, Open Street Map, transport, geography.

Corresponding author: Samuel Langton, University of Leeds, School of Law, Leeds, United Kingdom. Email: s.langton@leeds.ac.uk.

†School of Law, University of Leeds, Leeds, United Kingdom. ORCID: 0000-0002-1322-1553.

‡School of Social Sciences, University of Manchester, Manchester, United Kingdom. ORCID: 0000-0001-8689-1526.

Introduction

Access to data in crime and place research has traditionally been reserved for those who have the means to collect fresh data themselves, pay for access, or obtain data through formal data sharing agreements. Even when access is granted, the usage of these data often comes with conditions that circumscribe how the data can be used through licensing or policy (Kitchin, 2014). Even the public dissemination of findings which emerge from analysis might be subject to restrictions. This can lead to unequal access, controlled usage and curb the diffusion of findings, severely limiting the insight that can be obtained from data.

Open data initiatives provide a response to these shortcomings, broadening access and participation in research, removing the requirement for permissions, formal agreements, and negotiations (Manovich, 2011). Open data can lead to all sorts of novel insight within crime and place research, tapping into constructs and processes which are difficult to capture through surveys, interviews and other traditional measures (Solymosi & Bowers, 2018). As such, it is important that social scientists, researchers, crime analysts, and others interested in making sense of the world around them have the skills and know-how to access, interpret, critique, and analyse open datasets.

This chapter will outline how to develop such skills by providing a framework to approach and meaningfully interpret open data. The chapter also offers a practical hands-on guide to demonstrate how to access, wrangle, and analyse different sources of open data in order to draw conclusions about crime and place.

Background

What is open data?

First, it is useful to clarify what we mean when we refer to ‘open data’. The *Open Data Handbook*, compiled by the Open Knowledge Foundation, states that open data are “data that can be freely used, re-used and redistributed by anyone - subject only, at most, to the requirement to attribute and sharealike” (Dietrich et al., 2009, para. 3). Specifically, open data can be defined along three key domains (Open Knowledge Foundation, 2016).

- **Availability and Access:** the data must be accessible via a public domain, at no more than a reasonable reproduction cost, and through a convenient medium, such as the internet. Ideally, it should be machine-readable and modifiable. For example, it should be downloadable from a website as an unlocked spreadsheet, or JSON file, rather than presented as a summary table in a PDF document.
- **Re-use and Redistribution:** the data must be provided under conditions that allow free use, such as modification, separation, or compilation of the data, and permit re-use and redistribution including the intermixing (merging) with other datasets.
- **Universal Participation:** everyone should be allowed to use, re-use and redistribute the data and its derivatives without restriction. For instance, restrictions that only allow non-commercial or educational usage would *not* constitute open data.

In short, ‘open data’ should be data that are readily and publicly available, in a usable format, and without restrictions on usage, modification and re-distribution.

Open data practices, and the extent to which it is a reality, vary both within and between countries. For example, the United States has a history of making public sector datasets openly available. The United Kingdom can be more strict, releasing data under a licence (‘Open Government Licence’), with some data requiring a fee (Kitchin, 2014). The *Global Open Data Index* (Open Knowledge Foundation, 2020) is one tool for tracking the annual global benchmark for publication of open government data by country. It provides a rank which indicates how well governments across the world follow open data practices across various domains. Many of these have direct relevance to crime and place research, such as administrative boundaries and geographic locations (e.g. coordinates).

The drive towards transparency and scrutiny of public information is further enhanced by an increasing need for replication - a hallmark for open science research practice. Opening up datasets used in criminology publications and wider social sciences fosters and facilitates a culture of replication (Pridemore et al., 2018).

As such, open data fits into the wider discourse around transparency, open review and open scrutiny, which (in time) may even become a policy requirement irrespective of the research field or the source of funding (Vuong, 2017). Increasingly, researchers are making use of online repositories for their papers (e.g. <https://arxiv.org/>), data (e.g. <https://www.ukdataservice.ac.uk/deposit-data>), code (e.g. <https://github.com/>), or whole projects (e.g. <https://osf.io/>) all in the name of transparency and replication. These are guidelines and tools worth considering, not just when conducting analysis on secondary data sources, but also collecting and sharing your own data.

What are types of open data?

Now we have a reasonable idea about what open data is in a broad sense, we can consider the different types of open data which you might come across. Here, we formulate a typology of open datasets based on their origin.

Public sector

The ‘public sector’ refers to organizations owned and operated by local or central government with the core aim of providing services to the public. Much of the open data movement has focused on opening up information generated by these local and national state agencies (Kitchin, 2014). These can include national statistics or administrative data, but also data collected by publicly funded research projects such as victimization and household surveys, amongst others. The *Global Open Data Index*, introduced earlier, focuses on identifying the gaps in governmental organizations, encouraging them to think about how public sector information can become more usable, and ultimately, more impactful (Open Knowledge Foundation, 2020).

Specific to crime and place research, datasets of interest in this particular domain might include data describing urban structure, such as street networks, to assess whether the configuration of roads dictates things like violent crime victimization (Summers & Johnson, 2017), or victimisation surveys to quantify citizens’ perceived safety and security.

Private sector

In contrast to the public sector, ‘private sector’ tends to refer to parts of the economy which are not under direct state control, and often operate for profit. Opening up data generated by the private sector represents a significant challenge, largely due to the proprietary value to its creators (Kitchin, 2014). The aims and objectives of private companies, who have a duty to shareholders and operate in a competitive market environment, differ considerably from local and central government.

That said, some datasets are released openly by private sector organizations, albeit often only a subset of what would be available as a paying customer. The website for ArcGIS, a piece of software maintained by the Environmental Systems Research Institute (ESRI), a private company, host a number of open geographic datasets on their online cloud platform (<https://hub.arcgis.com/search>). Indeed, many papers have made use of data collected or distributed by private organizations to explore crime and place, such as Google Street View images (Langton & Steenbeek, 2017) and Twitter (Malleon & Andresen, 2015). In fact, you will learn how to obtain free Twitter data for such purposes in this book (see Chapter 6, Topic 7).

Open crowdsourced data

Finally, there is a specific group of open data sources that collate data collectively, generated by large groups of individuals who do not specifically belong to any organization, but instead work together on a collaborative project. These are crowdsourced datasets. Examples include Wikipedia (<https://www.wikipedia.org/>), an online encyclopedia where anyone can contribute, or Flickr (<https://www.flickr.com/>) an online photo gallery where people can upload and tag their photos. Often, the organizations who maintain and monitor these data collection activities are charities or non-governmental organizations that do not operate for profit, but instead provide some form of social good. An example is the online reporting platform FixMyStreet (<https://www.fixmystreet.com/>), where people can report problems such as instances of graffiti, vandalism

or environmental issues. This data has been used to explore signal crimes theory, for instance, offering insight into people's experiences with incivilities (Solymosi et al., 2018). Another topic in this book (see Chapter 6, Topic 6) discusses the merits and pitfalls of crowdsourced data, and what to watch out for when analyzing open data of this type.

Strengths and limitations

The defining characteristics of open data, namely, that of availability, re-usability and universal participation, outlined earlier, represent its greatest strength. But besides from this, there are other advantages. One specific motivation for using open data comes from its potential to address many of the limitations associated with traditional surveying methods, such as social desirability bias, or issues associated with memory and recall (Mayer-Schönberger & Cukier, 2013). This advantage exists largely due to the organic way in which open data is generated. It is often a by-product of other activities, and as such, we can gain an honest insight into people's everyday lives and associated social processes (Solymosi & Bowers, 2018). As noted, open data also means open research, facilitating transparency and reproducibility. That said, there are a number of shortcomings which are worthy of consideration.

Firstly, one of the biggest obstacles to collecting open data is the ability to interpret what the data means (e.g. applying a framework for its analysis), and computational skills in data scraping, wrangling and cleaning, in order to transform it into a usable format for research (Boyd & Crawford, 2012). All sorts of open data remain inaccessible to people who may lack the skills and know-how to acquire them. Although this is certainly an obstacle for data usage more generally, it represents a key challenge to governmental bodies, in particular, on which there is an onus to ensure transparency and facilitate scrutiny. Private companies, especially those who generate open data as a by-product of their primary activity (e.g. Twitter), have little responsibility or pressure to ensure that their data is accessible and usable to the general public. Moreover, in both public and private spheres, licences and conditions on open data are not necessarily concrete, and might be subject to change with little or no notice. This is an important consideration when planning and running long-term research projects which involve open data.

Secondly, despite the merits of open data in terms of remedying pitfalls in traditional survey methods, there are other threats to validity that may emerge. For instance, whilst many online resources offer an 'honest' depiction of society, uncaptured by surveys, researchers should be careful in interpreting people's communication online as completely authentic (Manovich, 2011). This could be a result of individuals willingly managing and 'curating' their online presence (Ellison et al., 2006), or because of wider issues such as government censorship, or cultural norms around particular topics, particularly sensitive topics of interest to researchers of crime, such as sexual assault or drug use.

Relatedly, researchers should be aware of issues over sampling, and ultimately, the generalizability of findings that emerge from the analysis of open data. In the case of crowdsourced data, the sample is not randomly drawn from a population, but rather, it is self-selected, giving way for people willing to discuss or contribute to a particular issue, which introduces a degree of bias (Longley, 2012). Specifically, contributors tend to be men, between the ages of 20-50, with a college or university degree (Budhathoki, 2010; Haklay, 2010). Contributions to resources such as Open Street Map (which we look at in the practical exercise later) are correlated with contextual characteristics such as poverty and population density, and as such, coverage is non-uniformly distributed across urban areas (Mashhadi et al., 2013). These are important to keep in mind (and be transparent about) when reporting findings based on analysis of such data.

What can be done?

Open data can be messy, biased and noisy, but criminology (and social sciences more generally) can benefit immeasurably from its use. Only through open data can public sector bodies be held to account, research be transparent and reproducible, and participation in data analysis universal. In crime and place research, both dependent variables (e.g. police-recorded crime incidents, victimization rates) and independent variables (e.g. demographic characteristics, ambient population estimates) can be sourced from open data, whether public, private or crowdsourced.

So, what can we do to make sure we make good use of these data? With a critical, engaged and considered approach to conducting research with open data, criminology can become a leading force in open and reproducible social science. In sum, researchers and analysts using open data must ask critical questions:

- **Where** does this data come from?
- **Why** was the data collected in the first place?
- **Who** is represented in the data, and who is excluded?
- **What** concepts and constructs can and cannot be operationalized with this data?
- **When** did data collection take place, and how might have that influenced results?

The answers to these questions will put the data sources in context, and ensure that researchers use them appropriately and with due care.

With this in mind, we now move on to a practical exercise in which we will utilize multiple sources of open data to explore police-recorded crime on and around public transport in London, England.

Practical exercise

In this exercise you will acquire a number of different skills, including:

- Accessing open data using three different methods:
 - Direct download.
 - Direct calls to an Application Programming Interface (API).
 - Calls to an API using a wrapper.
- Cleaning, wrangling, and visualizing open spatial data.
- Comparing different sources of open data.
- How to engage with the critical ‘Where’, ‘Why’, ‘Who’, ‘What’ and ‘When’ questions to better understand open data.

Our aim

In this exercise we will explore crime in and around London Underground stations. We know that environmental features are important when it comes to public transport areas being more or less criminogenic. Studies in Sweden, for instance, have shown the importance of environmental and neighborhood characteristics in determining crime concentrations at underground stations in Stockholm, along with the positioning of stations on the line (Ceccato et al., 2013). Similar work has been carried out exploring bus stops in Los Angeles (Loukaitou-Sideris, 1999) and the impact of intensive policing along bus corridors in Merseyside, England (Newton et al., 2004), amongst others.

Here, we will consider the case of London. Specifically, we will examine the question: to what extent does crime cluster in and around London Underground stations? We will use various sources of open data to answer this question. This will allow us to explore the strengths and limitations of public sector and crowdsourced open data sources. Should you struggle to download any of the data during the exercise, or if the data sources have been updated since publishing, each one has been archived and is available on the corresponding GitHub page for this topic (https://github.com/langtonhugh/osm_crim).

Accessing data

We will be using two different types of open data in this exercise: public sector data (police-recorded crime data and local transport authority data) and crowdsourced data (from Open Street Map). To access them, we will use three different methods, namely, (1) direct download, (2) direct request to an API, and (3) request to an API using a wrapper. This will give you a few different ideas about how you might go about accessing other open datasets relevant to your research. Locating, identifying and learning to access open data is a skill in itself.

Direct download

The simplest way that open data can be made available is through direct download from a website. In such a case, you can visit a website, select some parameters, and save a file containing the data you requested locally to your computer.

In the United Kingdom, police-recorded crime data in England and Wales can be accessed this way using an online web portal (<https://data.police.uk/>) under Open Government Licence. Visiting this website, you will see a welcome message, and six tabs across the top which should read “Home”, “Data”, “API”, “Changelog”, “Contact”, and “About”.

Before we download any data, we can learn more about it by clicking on the “About” tab. This brings up information that is important to review carefully in order to answer the where, why, who, what and when questions posed earlier. Take a moment to read through this information, and take notes on what you think might be relevant for your analysis. For example, if we want to map crimes, we will want to explore if there is any type of anonymization that might take place before the data are released (to protect the privacy of the victims). If you read the “About” page, you might find the following note:

Location anonymization. The latitude and longitude locations of Crime and ASB incidents published on this site always represent the approximate location of a crime — not the exact place that it happened.

This indicates that although we get a latitude and longitude coordinate with each crime event, it may only be approximate. This may have implications for our findings!

To then download some data, move on to clicking on the “Data” tab. This should open a page entitled “Data downloads” under which you can see another five tabs: “Custom download”, “Archive”, “Boundaries”, “Open data”, and “Statistical data”. By staying on this “Custom download” page you can select what sort of data you want to download. We can select the time period and police force of interest, and the type of information required (e.g. crimes, stop and search, outcomes). We are also informed that the data are downloaded in comma-separated values format (.csv file extension), which meets our machine-readable, easy-to-manipulate data format requirements.

For this exercise, we are going to use British Transport Police data, a force which operates on railways and light-rail systems across the country, for the month of January in 2020. Select the time period using the dropdown menus, the force using the tickboxes, and then generate and download the file.

Save this file locally (in your working directory) to a subfolder named “data”. You can load it in with the `read_csv()` command from the `readr` package. Remember that you will need to load it first using `library()`.

```
library(readr)

btp_df <- read_csv("data/2020-01-btp-street.csv")
```

Note: If you didn’t manage to follow along with the download instructions, you can also access this dataset from our GitHub page. Note that the line break below is purely for easy formatting - this will need removing in your own script so that the URL runs across one line.

```
btp_df <- read_csv("https://github.com/langtonhugh/osm_crim/raw/master/
data/2020-01-btp-street.csv")
```

You should now have a data frame with 5694 crimes in, ready for you to explore!

Direct request to an API

Another way of downloading open data is through an Application Programming Interface (API). This is a tool which defines an interface for a programme to interact with a software component. For example, it defines the kind of requests or calls which can be made, and how these calls and requests can be carried out. Here, we are using the term ‘API’ to denote tools created by an open data provider to give access to different

subsets of their content. Such APIs facilitate scripted and programmatic extraction of content, as permitted by the API provider (Olmedilla et al., 2016). APIs can take many different forms and be of varying quality and usefulness (Foster et al., 2016). For the purposes of accessing open data from the web, we are specifically talking about *RESTful* APIs. The ‘REST’ stands for Representational State Transfer. These APIs work directly over the web, which means users can play with the API with relative ease in order to understand how it works (Foster et al., 2016).

Here, we will make a direct request to the API created by London’s local government organization responsible for transport: Transport for London (TfL). TfL oversee the London Underground. They provide access to their open data through a unified API. Much like how we found the “About” page when downloading police data directly, to answer our ‘Where’, ‘Why’, ‘Who’, ‘What’ and ‘When’ questions, we must find a similar document for TfL. Details about the TfL API are provided through their open data page (<https://tfl.gov.uk/info-for/open-data-users/unified-api>), which explains what data is available, and how the API is designed.

They further provide a documentation page with examples of how you can use HTTP (i.e. web link) requests to make calls to the API (see: <https://api-portal.tfl.gov.uk/docs>). If you follow this link you should see a page titled “Our Unified API” and some examples of calls. To draw conclusions about crimes in and around London Underground stations, we will need some spatial data about the stations themselves.

In the documentation, there is a section called ‘API area’ which offers some guidance on the types of data available. Scroll down and find the example called *Stops*. You will see that this call returns information on stops for buses or London Underground lines, and there are two examples (bus route 24 and the Bakerloo line). Copy their demonstration URL for the bakerloo line (<https://api.tfl.gov.uk/line/bakerloo/stoppoints>) and paste it into your web browser. When you visit this page, you should see something like Figure 1.

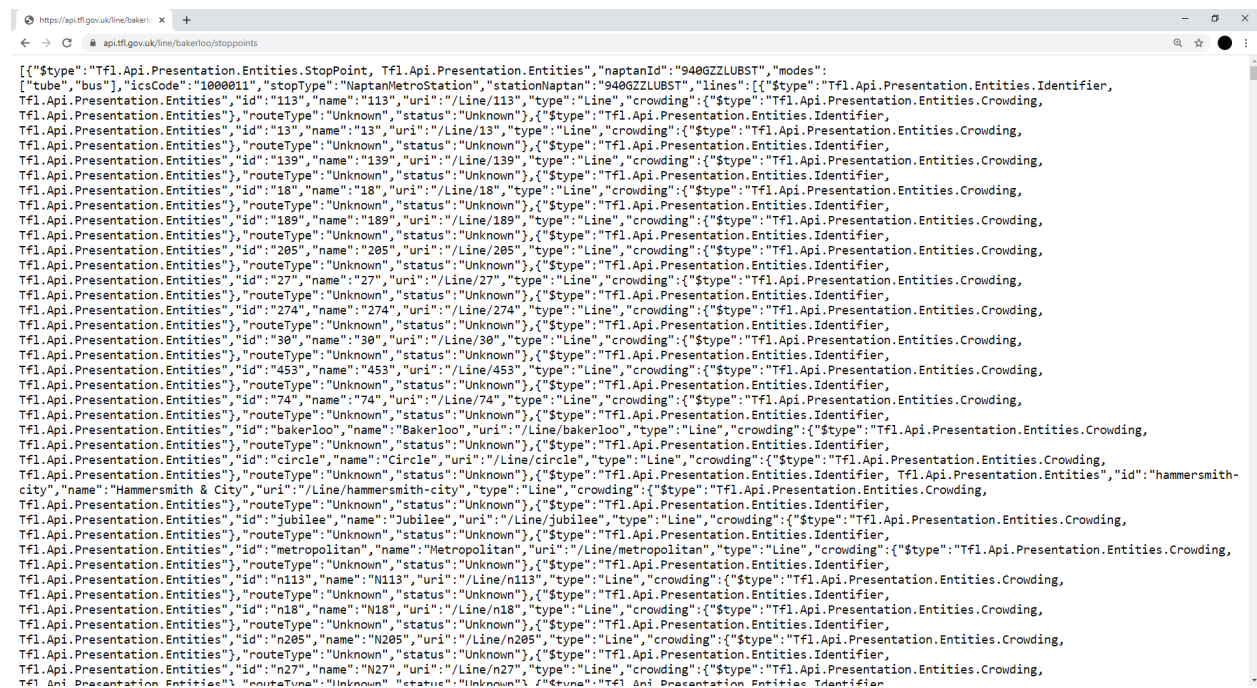


Figure 1: example of data from Transport for London

Although this contains the information we requested, it is not very legible for our human eyes. It is in a format called *JSON*. This stands for JavaScript Object Notation. JSON is an open standard way of storing and exchanging data, and will most likely be the format in which data are returned from most API calls. As you can see, it is not intuitive to read in its current format *but* it is machine-readable friendly, which again aligns well with our requirements for open data.

In reality, you will not usually make these calls by pasting them into a web browser, but it is useful to see how it works. More likely, we will actually make these calls from within the R environment. That said, it might be useful (and interesting) to examine how basic queries are constructed using this documentation. For example, we now know that to request the stops from the “Bakerloo” line, we need the URL:

```
"https://api.tfl.gov.uk/line/bakerloo/stoppoints"
```

Let’s say that instead, we want data for the Jubilee Line. What do you think that URL will look like? Indeed, all you need to do is replace “bakerloo” with “jubilee”.

In R, we can use the `fromJSON()` function from the `jsonlite` package to parse all this information into a data frame (with rows and columns). It will then become more familiar and usable. All we need to do is input the URL from the TfL API into the function with a bit of help from `readLines()`, a function in base R for reading text from URLs.

To keep things focused, let’s request data about stop points on the Jubilee line, one of the busiest on the network. Note how this is simply an amended version of the example provided by TfL in the API documentation.

```
library(jsonlite)

api_call <- fromJSON(readLines("https://api.tfl.gov.uk/line/jubilee/stoppoints"))
```

This gives us an object (`api_call`) which contains all the information returned by the TfL API. JSON is slightly different to traditional data frames with rows and columns, which we are probably more familiar with, because the data is nested. For instance, `api_call` is classed as a data frame, but now try viewing the object using `View(api_call)`. You will notice that some of the columns are actually lists, rather than character or factor vectors, which is what we might usually expect. This demonstrates an important challenge faced by researchers when using open data, because dealing with data in this format can be messy and complicated. It is not always a neatly formatted data frame like the CSV file from the open police data portal. That said, we can transform this data into something more familiar from within R, as we will see later on.

Request to an API using a wrapper

Finally, we will look at how open data can be obtained using API wrappers. Often, developers who work with APIs will share their code, and release them in the form of a package or module, so that other people can use it. This is called a *wrapper* because it uses code that ‘wraps’ around the API to make it a neater, more usable package. Wrappers remove (or at least lower) many of the obstacles to accessing open data noted earlier. The wrapper can take many forms, such as a Python module or an R package. It could even be a web interface that provides a graphical user interface (GUI) for accessing the API in question.

To demonstrate this, we will be accessing data from Open Street Map, a database of geospatial information built by a community of mappers, enthusiasts and members of the public, who contribute and maintain data about all sorts of environmental features, such as roads, green spaces, restaurants and railway stations, amongst many other things, all over the world. As such, it is a prime example of ‘crowdsourced’ open data. You can view the information contributed to Open Street Map using their online mapping platform (<https://www.openstreetmap.org/>). The result of people’s contributions is a database of spatial information rich in local knowledge which provides invaluable information about places and their features, without being subject to strict terms on usage.

Open Street Map has two main types of wrappers available for its API, a web-based GUI called Overpass Turbo (<https://overpass-turbo.eu/>), and an R package called `osmdata`.

If we load the package `osmdata` we can use its functions to query the Open Street Map API, rather than the API query being made directly (like we did for TfL, using the URL). Once again, we will want to identify documentation which can help us understand and critique our data, and learn how to query it. When it comes to Open Street Map, you can read all about it on their community page: <https://www.openstreet>

map.org/about. Using the wrapper in R, we can refer to the package documentation and the associated vignette online: <https://cran.r-project.org/web/packages/osmdata/vignettes/osmdata.html>.

Unlike the TfL API, `osmdata` is an international database, and has lots of data that we might not necessarily need. As such, the first thing we need to specify is our study region. To do this, we need to specify a bounding box. You can think of the bounding box as a box drawn around the area that we are interested in (in this case, London, England) which tells the Open Street Map API that we want everything *inside* the box, but nothing *outside* the box.

So, how can we name a bounding box specification to define the study region? One way to do this is through a search term. Here, we want to select Greater London, so we can use the search term “greater london united kingdom” within the `getbb()` function. Using this function we can also specify what format we want the data to be in. In this case, we want a spatial object, specifically an `sf` polygon object (from the `sf` package), which we name `bb_sf`.

```
library(osmdata)
library(sf)

bb_sf <- getbb(place_name = "greater london united kingdom", format_out = "sf_polygon")
```

One downside of using search terms to define the bounding box, as with any search term, is that inconsistencies in terminology, wordings and spelling might lead to unexpected outputs (or none at all). Another way to obtain the bounding box is to manually specify the latitude and longitude coordinates. This requires a bit of pre-existing knowledge about your study region. A quick internet search, or exploration of the interactive Open Street Map platform (<https://www.openstreetmap.org/>), will give you an idea of the coordinates for your study region. For Greater London, we can specify our bounding box as follows, saving the coordinates to the object `bb_gl` for later use.

```
bb_gl <- c(-0.51037, 51.28676, 0.33401, 51.69187) # xmin, ymin, xmax, ymax
```

We can now move on to query data from the Open Street Map API using the `opq()` function. The function name is short for ‘Overpass query’, which is how users can query the Open Street Map API using search criteria.

Besides specifying what area we want to query with our bounding box object(s) in the `opq()` function, we must also define the feature which we want to pull from the API. Features in Open Street Map are defined through ‘keys’ and ‘values’. Keys are used to describe a broad category of features (e.g. highway, amenity), and values are more specific descriptions (e.g. cycleway, bar). These are tags which contributors to Open Street Map have defined. A useful way to explore these is by using the comprehensive Open Street Map Wiki page on map features (https://wiki.openstreetmap.org/wiki/Map_Features).

We can select what features we want using the `add_osm_feature()` function, specifying our key as ‘public transport’ and our value as ‘station’. We also want to specify what sort of object (what class) to get our data into, and as we are still working with spatial data, we stick to the `sf` format, for which the function is `osmdata_sf()`. Here, we specify our bounding box as the manual coordinates contained in `bb_gl`.¹

```
osm_stat_sf <- opq(bbox = bb_gl) %>% # select bounding box
  add_osm_feature(key = 'public_transport', value = 'station') %>% # select features
  osmdata_sf() # specify class
```

The resulting object `osm_stat_sf` contains lots of information. We can view the contents of the object by simply executing the object name into the **Console**.

```
osm_stat_sf
```

This confirms details like the bounding box coordinates, but also provides information on the features collected from the query. As one might expect, most information relating to public transport station locations

¹If you use the bounding box obtained through `getbb()` one can subsequently trim down the outputs from `add_osm_feature()` using `trim_osmdata()`. For instance, we could add `trim_osmdata(bb_poly = bb_sf)` to our initial query.

has been recorded using points (i.e. two-dimensional vertices, coordinates) of which we have 1736 at the time of writing. We also have around fifty polygons. For now, let's extract the point information.

```
osm_stat_sf <- osm_stat_sf$osm_points
```

We now have an `sf` object with all the public transport stations in our study region mapped by Open Street Map volunteers, along with ~130 variables of auxiliary data, such as the `fare_zone` the station is in, what `amenity` it may have and whether it has `toilets`, amongst many others. Of course, it is up to the volunteers whether they collect all these data, and in many cases, they have not added information. Nevertheless, when the details are recorded, they provide rich insight and local knowledge that we may otherwise be unable to obtain.

One additional step needed is to select the stations which are relevant to us (i.e. they fall along the Jubilee Line). The variable `line` can help us with this. First, let's explore this variable by using the `table()` function.

```
table(osm_stat_sf$line)
```

You will see that "Jubilee" appears several times under different line names. The main category ("Jubilee") is the most common, but many other combinations exist (e.g. "Central;Jubilee") for stations which sit on multiple lines. To get round this, and select all stations with "Jubilee" somewhere in the name, we can use the `grepl()` function from base R. Note that we also make use of `filter()` from within `dplyr`, so we first need to load that package.

```
library(dplyr)
```

```
osm_jub_sf <- osm_stat_sf %>%  
  filter(grepl("Jubilee", line))
```

This gives us Open Street Map data on all 27 stations on the Jubilee underground line. Now we have all our open datasets loaded into the R environment!

Cleaning, wrangling, and visualizing our data

Obtaining open data from a direct download or API is often only half the battle. It doesn't necessarily mean that the data is in the shape needed to conduct analysis. The first thing to do once we have acquired open data is investigate it using data cleaning, wrangling, and visualization. We've already loaded the `osmdata` and `dplyr` packages, but here, we will also make use of `ggplot2`, `sf`, `tidyr` and `patchwork`, so load these now. Remember, if you do not have these packages installed, use the `install.packages()` function prior to loading each one with `library()`.

```
library(ggplot2)  
library(sf)  
library(tidyr)  
library(patchwork)
```

From JSON to data frame

We usually want to work with data that is in the format of a data frame, that is, all your observations are rows and all your variables are columns. While this is the case for our police data and our Open Street Map data, recall that the TfL data came in the nested JSON format.


Fortunately for us, R is more than capable of dealing with nested data. Moreover, we can see that the information we are interested in (namely, the location of underground stations) has already been successfully parsed into columns called `lon` (longitude) and `lat` (latitude), along with an identification column of the station names called `commonName`. We can extract these columns, and create an `sf` (spatial) point object using these coordinates, all in one chunk of code, using the piping operator (`%>%`). Note that we define the

Coordinate Reference System (CRS) which the data has come in (WGS 84), and then transform it into something more appropriate (the British National Grid, a projected CRS used in the United Kingdom).

```
tfl_jub_sf <- api_call %>%  
  select(commonName, lat, lon) %>%  
  st_as_sf(coords = c(x = "lon", y = "lat"), crs = 4326) %>%  
  st_transform(27700)
```

We can now easily visualize this data using the `ggplot2` package, which is compatible with `sf` objects.

```
ggplot(tfl_jub_sf) +  
  geom_sf()
```

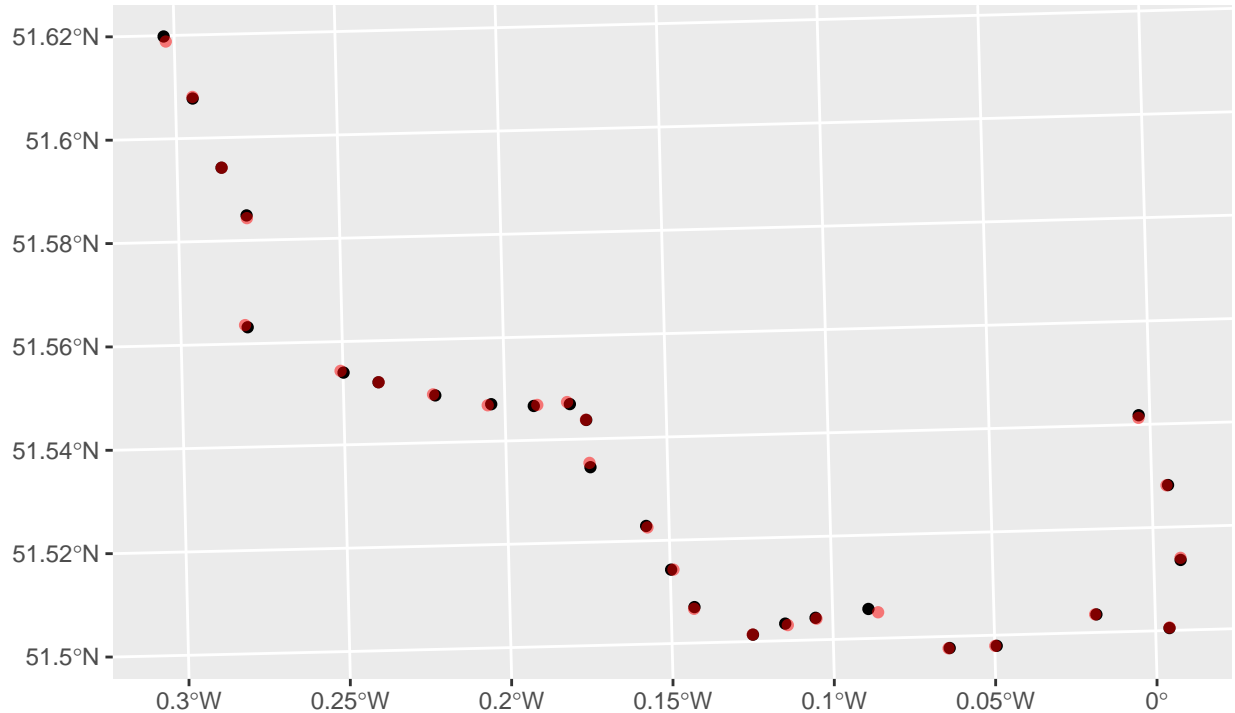


There we have it: with just a few lines of code in R, we have queried the TfL API, a public sector source of open data, and plotted a basic map of stations on the Jubilee underground line.

Visualise to check and compare our datasets

We can now plot the Open Street Map station points, which are already spatial, over these station locations pulled from the TfL API, to see how the two datasets compare.

```
ggplot() +  
  geom_sf(data = tfl_jub_sf) +  
  geom_sf(data = osm_jub_sf, col = "red", alpha = 0.5)
```



As we can see, the two data sources appear to be broadly comparable in terms of their spatial patterning. However, the exact location of stations sometimes differs between Open Street Map and TfL. Determining a single pinpoint location for underground stations is not straightforward, especially for larger stations which have multiple entrances, platforms, and walkway tunnels. For operational purposes, TfL might have a specific definition for the location of a station, such as its centroid, which has been systematically applied to create the coordinates. The data pulled from Open Street Map might have been compiled by a volunteer contributor with their own perception of how to define the location of a station, such as the ticket booth or the main entrance. This definition *might* be more theoretically meaningful than that used by TfL, since it could reflect a realistic ‘on the ground’ perception of the station’s built environment, and thus opportunities for crime. On the other hand, if the locations from Open Street Map were collated by multiple different contributors, each with their own idea of where each station lies on the line, we might conclude that the measure is unreliable. These kinds of questions demonstrate why critically engaging with open data sources is so crucial, but also challenging. We return to this at the end of the exercise.

Given the differences observed between the two open data sources, what might the impact be when studying crime in and around London Underground stations? For that, we can turn to our British Transport Police data. In its raw form, this data is a standard data frame, but we can make it spatial using some of the functions from the `sf` package that we used earlier, after first dropping crimes that have incomplete coordinates.

```
btp_sf <- btp_df %>%
  drop_na(Longitude, Latitude) %>%
  st_as_sf(coords = c(x = "Longitude", y = "Latitude"), crs = 4326) %>%
  st_transform(27700)
```

We now have the recorded location of all crimes recorded by the British Transport Police during January, 2020. For the purposes of this demonstration, we can define crimes occurring as ‘in and around’ Jubilee line

tube stations by creating a 50 meter buffer around each station, for each source of data, and counting the number of points falling within each. It is worth noting that this definition is somewhat arbitrary, and is subject to the spatial inaccuracy in open police-recorded crime data (see Tompson et al., 2015), but it does serve to facilitate this demonstration.

```
# Ensure each has the same CRS (BNG, 27700).
osm_jub_sf <- st_transform(osm_jub_sf, st_crs(tfl_jub_sf))

# Create buffers to define 'in and around'.
osm_buff_sf <- st_buffer(osm_jub_sf, dist = 50)
tfl_buff_sf <- st_buffer(tfl_jub_sf, dist = 50)

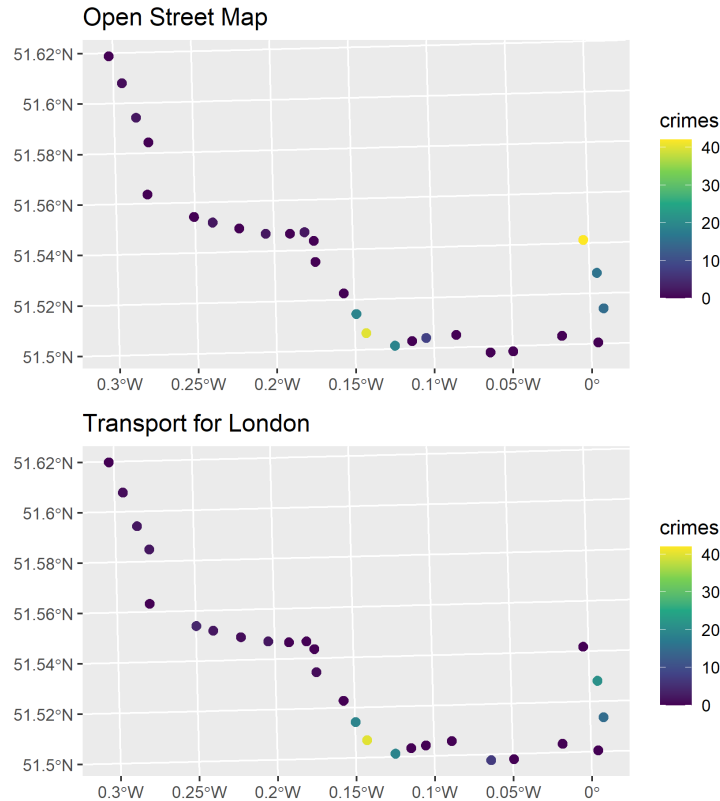
# Count number of crimes recorded within each buffer.
osm_jub_sf <- osm_buff_sf %>%
  mutate(crimes = lengths(st_intersects(osm_buff_sf, btp_sf)))
tfl_jub_sf <- tfl_buff_sf %>%
  mutate(crimes = lengths(st_intersects(tfl_buff_sf, btp_sf)))
```

We can then visualize these counts and compare the two sources of data by coloring in our buffers according to the crime count. Note that we arrange the plots using syntax available using the `patchwork` package, and set the scales to be comparable using the `lim` argument within `scale_color_viridis_c()`.

```
p1 <- ggplot(data = osm_jub_sf) +
  geom_sf(mapping = aes(color = crimes), size = 2) +
  labs(title = "Open Street Map") +
  scale_color_viridis_c(lim = c(0,42))

p2 <- ggplot(data = tfl_jub_sf) +
  geom_sf(mapping = aes(color = crimes), size = 2) +
  labs(title = "Transport for London") +
  scale_color_viridis_c(lim = c(0,42))

p1 / p2
```



A number of factors emerge from this investigation. First, there is a disparity between the total number of crimes captured by the buffers in each dataset. Using the Open Street Map data, a total of 168 crimes were recorded across all stations on the line. Using the TfL data, only 136 were recorded. Thus, the volume of crime observed in and around Jubilee stations during January will differ depending on the open data source being used. This could have important implications for how community and policing resources are distributed between underground lines. Second, the spatial patterning of crime hotspots on the line differs considerably between the two data sources. Using the Open Street Map data, the final station on the line, Stratford, has the greatest concentration of crimes, whereas the TfL data suggests that the most problematic station is Green Park, in the city centre. Despite the minor differences in the locations of each station, the conclusions drawn are starkly different, and would determine the consistency of findings with existing research examining crime at underground stations (Ceccato et al., 2013).

We can explore Stratford station as a potential crime hotspot using some of the skills we have learnt so far. One useful first step might to visualize the built environment around the station, along with the crime buffers from Open Street Map and TfL respectively, and the crime locations themselves. This way, we might find out why the TfL buffer is failing to capture any crimes, while the Open Street Map buffer is capturing an apparent crime hotspot. In doing so, it is worth remembering that crime locations in open police data are approximations of the real geography in order to ensure anonymity.

```
# Retrieve Stratford buffer from Open Street Map.
stratford_osm_sf <- osm_jub_sf %>%
  filter(name == "Stratford")

# Do the same using TfL.
stratford_tfl_sf <- tfl_jub_sf %>%
  filter(commonName == "Stratford Underground Station")

# Create a mini study area for our visualization, by getting the bounding box for the
```

```

# Open Street Map (or TfL) buffer. Note that we transform to WGS 84 first in preparation
# for the Open Street Map API.
bb_sf <- stratford_osm_sf %>%
  st_transform(crs = 4326) %>%
  st_bbox()

# Retrieve the coordinates of the bounding box, and create a numeric list for use with
# osmdata package.
strat_bb <- c(bb_sf[[1]], bb_sf[[2]], bb_sf[[3]], bb_sf[[4]])

# Pull building footprints in this bounding box from the Open Street Map API.
strat_fp_sf <- opq(strat_bb) %>%
  add_osm_feature(key = 'building') %>%
  osmdata_sf()


# Get the polygons only.
strat_fp_poly_sf <- strat_fp_sf$osm_polygons

# Transform back to British National Grid to match our other data.
strat_fp_poly_sf <- st_transform(strat_fp_poly_sf, 27700)

# Clip the crime data to the Open Street Map buffer. Remember that the TfL buffer
# did not capture any crimes.
strat_osm_btp_sf <- st_intersection(stratford_osm_sf, btp_sf)

# Visualize the building footprints, Open Street Map buffer, TfL buffer, and crime
# locations. Note that the single point for the crimes is due to the approximation!
# There are actually multiple crimes. Colors have been manually selected
# from the viridis package to be easier to read by those with colorblindness.
ggplot() +
  geom_sf(data = strat_fp_poly_sf, fill = "black") +
  geom_sf(data = stratford_osm_sf, fill = "#440154FF", alpha = 0.3) +
  geom_sf(data = stratford_tfl_sf, fill = "#FDE725FF", alpha = 0.3) +
  geom_sf(data = strat_osm_btp_sf, col = "#3CBC75FF", fill = "white") +
  theme(axis.text = element_text(size = 6))

```



This quick check has shed some light on what is happening. The buffer generated from the TfL station location (in yellow) covers a different part of the station to the Open Street Map buffer (in blue). The latter covers most of the main station building, while the former centers around the platforms to the north. Because the British Transport Police crime locations have been geocoded to a location in the main station building (in green) the TfL buffer has missed them. If we had used a larger buffer, it is likely that the TfL location would have captured the crime cluster, but then we would increase the risk of picking up crimes outside of the station. In many cases, there will not be an obvious ‘right or wrong’ answer to determine which open data source is most appropriate, but investigations like this one will help inform the decision. We might consider investigating other stations before deciding on which data to use, and ensure that findings are reported transparently with reference to these discussions.

Engaging with the creators of the data would also be an important step in answering the critical questions (‘Where’, ‘Why’, ‘Who’, ‘What’ and ‘When’) posed at the beginning of this chapter. For TfL, this might mean investigating the criteria used by their analysts to define the coordinates for stations. For Open Street Map, its online transparency means that we could find out which user(s) recorded the Jubilee station locations, and directly contact this community of volunteers to find out how station locations were recorded and why.

Discussion

Summary

This chapter has sought to introduce open data as a novel, growing and invaluable tool in crime and place research through a review of key fundamentals and a substantive demonstration within R. We have defined open data, and outlined some key types of open information available, along with their respective strengths and weaknesses. An important component of this review has been to encourage an engaged and critical research approach to using open data. Open data is both a rich resource of dependent and independent variables for researching the geography of crime, and an interesting research topic in and of itself as a tool for providing insight overlooked by traditional data sources. By way of an example, we accessed different sources of open data in order to explore crime concentrations at London Underground stations on the Jubilee line: police-recorded crime data via a direct download from a public sector website, transport data via directly querying a public sector API, hosted by Transport for London, and the equivalent transport data retrieved from a crowdsourced database, Open Street Map, using an API wrapper. Minor discrepancies in the geographic location of underground stations in each open data source painted two starkly different stories about the volume and concentration of crime occurring on the line. The demonstration highlighted both the power of open data and the need to critically engage with sources.

Future of open data

As it stands, open data represents a key resource in crime and place research. It is a fundamental component in the movement towards transparent and reproducible scientific research. The skills required to effectively access and use open data, such as those demonstrated in this chapter, are increasingly being taught and deployed in social science. The number and quality of open data resources is also improving, many of which will help illuminate key topics in spatial criminology. For instance, information about building construction in Greater London is currently being collected through a crowdsourced platform, *Colouring London* (Hudson et al., 2019), opening prospect for studies to investigate historical urban development and crime concentrations in the capital, unrestricted by fees or strict terms, for the first time. But, as the discussions and demonstrations in this chapter have shown, there is still a long way to go. As with any emerging domain, open data sources vary considerably in their degree of completeness and reliability. Whilst this can offer insight, for instance, in terms of public participation and sense of community, it can also represent a significant obstacle to empirical analysis. A key challenge for the future will be to ensure that students and research practitioners ask critical questions of their data before blindly delving into analysis and public dissemination. Moreover, the longevity of open data licences remains an unknown parameter. With so much open data being distributed by the public sector, the accessibility of open information is subject to fluctuations in the amount of resource (and goodwill) available to safe-guard continuity. One way of justifying the continued collection and distribution of open public sector data is simply to use the data for public good, such as informing interventions which help reduce crime victimization.

Software

As noted in the tutorial, a number of R packages were used to create the material in this chapter. Packages *readr* (Wickham et al., 2018) and *jsonlite* (Ooms, 2014) were used for reading and parsing data, *osmdata* for querying the Open Street Map API (Padgham et al., 2017), *sf* for handling and visualizing spatial data (Pebesma, 2018), *dplyr* and *tidyr* for wrangling data (Wickham et al., 2020; Wickham, 2020), *ggplot2* for visualization (Wickham, 2016) and *patchwork* for arranging plots (Pedersen, 2020). The document itself was prepared using *rmarkdown* (Xie et al., 2018).

References

- Boyd, D., & Crawford, K. (2012). Critical questions for big data: Provocations for a cultural, technological, and scholarly phenomenon. *Information, Communication & Society*, 15(5), 662–679.
- Budhathoki, N. R. (2010). *Participants' motivations to contribute geographic information in an online community* [PhD thesis]. University of Illinois at Urbana-Champaign.
- Ceccato, V., Uittenbogaard, A., & Bamzar, R. (2013). Security in stockholm's underground stations: The importance of environmental attributes and context. *Security Journal*, 26(1), 33–59.
- Dietrich, D., Gray, J., McNamara, T., Poikola, A., Pollock, P., Tait, J., Zijlstra, T., & others. (2009). Open data handbook. In *Open Knowledge International*. <https://opendatahandbook.org/guide/en/what-is-open-data/>
- Ellison, N., Heino, R., & Gibbs, J. (2006). Managing impressions online: Self-presentation processes in the online dating environment. *Journal of Computer-Mediated Communication*, 11(2), 415–441.
- Foster, I., Ghani, R., Jarmin, R. S., Kreuter, F., & Lane, J. (2016). *Big data and social science: A practical guide to methods and tools*. crc Press.
- Haklay, M. (2010). How good is volunteered geographical information? A comparative study of open-streetmap and ordnance survey datasets. *Environment and Planning B: Planning and Design*, 37(4), 682–703.
- Hudson, P., Dennett, A., Russell, T., & Smith, D. (2019). Colouring london—a crowdsourcing platform for geospatial data related to london's building stock. Proceedings from 27th annual gis research uk conference. *Newcastle University, Newcastle, 23-26 April*.
- Kitchin, R. (2014). *The data revolution: Big data, open data, data infrastructures and their consequences*. SAGE Publications.
- Langton, S. H., & Steenbeek, W. (2017). Residential burglary target selection: An analysis at the property-level using google street view. *Applied Geography*, 86, 292–299.
- Longley, P. A. (2012). Geodemographics and the practices of geographic information science. *International Journal of Geographical Information Science*, 26(12), 2227–2237.
- Loukaitou-Sideris, A. (1999). Hot spots of bus stop crime: The importance of environmental attributes. *Journal of the American Planning Association*, 65(4), 395–411.
- Malleson, N., & Andresen, M. A. (2015). The impact of using social media data in crime rate calculations: Shifting hot spots and changing spatial patterns. *Cartography and Geographic Information Science*, 42(2), 112–121.
- Manovich, L. (2011). Trending: The promises and the challenges of big social data. *Debates in the Digital Humanities*, 2, 460–475.
- Mashhadi, A., Quattrone, G., & Capra, L. (2013). Putting ubiquitous crowd-sourcing into context. *Proceedings of the 2013 Conference on Computer Supported Cooperative Work*, 611–622.
- Mayer-Schönberger, V., & Cukier, K. (2013). *Big data: A revolution that will transform how we live, work, and think*. Houghton Mifflin Harcourt.
- Newton, A. D., Johnson, S. D., & Bowers, K. J. (2004). Crime on bus routes: An evaluation of a safer travel initiative. *Policing: An International Journal of Police Strategies & Management*.
- Olmedilla, M., Martínez-Torres, M. R., & Toral, S. (2016). Harvesting big data in social science: A methodological approach for collecting online user-generated content. *Computer Standards & Interfaces*, 46, 79–87.
- Ooms, J. (2014). The jsonlite package: A practical and consistent mapping between json data and r objects. *arXiv:1403.2805 [stat.CO]*. <https://arxiv.org/abs/1403.2805>
- Open Knowledge Foundation. (2016). *Open definition version 2.1*. <https://opendefinition.org/od/2.1/en/>

- Open Knowledge Foundation. (2020). *Global open data index*. <https://index.okfn.org/>
- Padgham, M., Rudis, B., Lovelace, R., & Salmon, M. (2017). Osmdata. *The Journal of Open Source Software*, 2(14). <https://doi.org/10.21105/joss.00305>
- Pebesma, E. (2018). Simple Features for R: Standardized Support for Spatial Vector Data. *The R Journal*, 10(1), 439–446. <https://doi.org/10.32614/RJ-2018-009>
- Pedersen, T. L. (2020). *Patchwork: The composer of plots*. <https://CRAN.R-project.org/package=patchwork>
- Pridemore, W. A., Makel, M. C., & Plucker, J. A. (2018). Replication in criminology and the social sciences. *Annual Review of Criminology*, 1, 19–38.
- Solymosi, R., & Bowers, K. (2018). The role of innovative data collection methods in advancing criminological understanding. *The Oxford Handbook of Environmental Criminology*, 210, 210–237.
- Solymosi, R., Bowers, K. J., & Fujiyama, T. (2018). Crowdsourcing subjective perceptions of neighbourhood disorder: Interpreting bias in open data. *The British Journal of Criminology*, 58(4), 944–967.
- Summers, L., & Johnson, S. D. (2017). Does the configuration of the street network influence where outdoor serious violence takes place? Using space syntax to test crime pattern theory. *Journal of Quantitative Criminology*, 33(2), 397–420.
- Tompson, L., Johnson, S., Ashby, M., Perkins, C., & Edwards, P. (2015). UK open source crime data: Accuracy and possibilities for research. *Cartography and Geographic Information Science*, 42(2), 97–111.
- Vuong, Q. H. (2017). Open data, open review and open dialogue in making social sciences plausible. *Nature: Scientific Data Updates*. <https://ssrn.com/abstract=3086667>
- Wickham, H. (2016). *Ggplot2: Elegant graphics for data analysis*. Springer-Verlag New York. <https://ggplot2.tidyverse.org>
- Wickham, H. (2020). *Tidyr: Tidy messy data*. <https://CRAN.R-project.org/package=tidyr>
- Wickham, H., François, R., Henry, L., & Müller, K. (2020). *Dplyr: A grammar of data manipulation*. <https://CRAN.R-project.org/package=dplyr>
- Wickham, H., Hester, J., & François, R. (2018). *Readr: Read rectangular text data*. <https://CRAN.R-project.org/package=readr>
- Xie, Y., Allaire, J. J., & Golemund, G. (2018). *R markdown: The definitive guide*. Chapman; Hall/CRC. <https://bookdown.org/yihui/rmarkdown>