

SAGE Research Methods: Doing Research Online

Authors: Please complete only the white fields below.

Entry Title		<i>[Title must include the topic of focus. For example, “How To...Design and Conduct a Digital Ethnography”.]</i>
		Using the Twitter Academic API for Social Science Research
Authors: Christopher Barrie and Justin Chun-ting Ho		
1	Name	Christopher Barrie
	Author email	christopher.barrie@ed.ac.uk
	Affiliation, country	University of Edinburgh, Edinburgh, United Kingdom
2	Name	Justin Chun-ting Ho
	Author email	justin.chunting.ho@sciencespo.fr
	Affiliation, country	Sciences Po, Paris, France
	SAGE Author ID	<i>[office use only]</i>
Author bio		<i>[Maximum of 200 words]</i>
Academic Level of intended readership		<i>[Click here to select level]</i>
<i>For office use only:</i>		
Title/Spin ID		
Access/Product Code		
ISBN		
DOI		
URI		
URL		
Copyright year		2022
Copyright statement		© SAGE Publications Ltd 2022

Doing Research Online ‘How to’ Guides are text pieces which will support social & behavioural science researchers at all levels, from undergraduates to experienced researchers. Guides will be used by instructors to support their teaching and supervision both in and outside the virtual classroom, and will also be used by students for independent learning. They will be written in an introductory and accessible style, and each ‘How to’ Guide will be enhanced with MCQs to test the readers understanding of the Guide.

‘How to’ Guides provide practical help with using digital research methods. Guides focus on topics such as:

Cutting-edge digital methods and tools (quant, qual and mixed); how to choose between them; when you might want to combine more conventional and cutting-edge methods; how to keep abreast of new developments

Tried and tested methods from focus groups, interviews to online surveys, and how they can be adapted to online use

Key aspects of doing research digitally:

- Designing experiments/studies for online participation
- Translating existing protocol to an online format and begin collecting data immediately.
- Creating research plans and proposals that incorporate best practices for online research.
- How to do online research in a reproducible way
- What constitutes legitimate digital data, and how to clean and analyse it

Guidelines

Please see the full manuscript guidelines document for further information.

Style & Format

Guides should be written in an accessible, jargon-free, and engaging manner. Please aim to be descriptive, explanatory and exploratory where appropriate.

We encourage the use of **multiple headings** to break up the text. Consider throughout the entry whether it is clear which sections pertain to which learning objectives.

Manuscripts should be prepared and submitted using the 'How to ...' Guide template.

Content

Whilst we encourage the **use of examples drawn from real research** in these guides, please bear in mind the **general applicability of the content** at all times. Remember, the aim of these guides is to introduce the user to the topic at hand and to **provide methodological guidance and practical insights that can be applied in the reader's research**.

Explain concepts in an uncomplicated manner, and use examples to illustrate concepts and definitions. Reference **methodological literature** wherever appropriate.

Referencing & Copyright

References should conform to American Psychological Association (APA) style, 7th edition, and should contain the digital object identifier (DOI) where available. SAGE will not accept guides that are incorrectly referenced; please ensure accuracy before submission. For help on reference styling see <https://apastyle.apa.org/style-grammar-guidelines>.

Any **figures/tables** used must be **copyright and permissions free**. If you use a figure or table for which you do not hold the copyright, then you must obtain permission from the copyright holder to use it. A form for requesting permissions is available upon request. You must also include an in-text credit line and a full citation for the figure/table in the reference list.

Queries / Concerns / Deadlines

For further information please consult the extended guidelines – this includes a section on how figures and tables should be supplied and guidance for writing MCQs.

If you have queries that are not addressed in the extended guidelines, please do relay these directly to your editorial contact. In most cases, this will be martin.perchard@sagepub.co.uk.

Your deadline will have been agreed with you previously. If you require a short extension, or if you are no longer able to write this guide, please be in touch with your editorial contact at the earliest opportunity.

Abstract

The abstract should be a concise summary of your 'How To' Guide. What aspect of a research process, or specific methodological and practical challenges, will your guide address? It should be succinct and enticing, and should incorporate key words and concepts discussed in the body of the text. Please do not cite references within the abstract.

[Insert here: Maximum of 250 words]

In this 'How To' Guide you will learn how to access the Twitter Academic Research Product Track API and collect tweet data of different kinds. This API was recently introduced and allows users far more access to Twitter data than ever before. In order to make requests to (get data from) this API, you will be using a dedicated R package called "academictwitterR," which the authors helped to develop. We will learn how to count tweets matching specific criteria—that is, words or phrases, from particular regions, written in particular languages, over particular time periods. We will also learn how to collect the actual tweets themselves as well as information on the tweeters. On the way, you will learn to consider several methodological questions when collecting data of this kind. Chiefly, we will learn how to denominate our tweet counts by a relevant baseline in order to gauge over-time trends; we will think about how to sample from relevant users, and we will consider the population of individuals or phenomena of interest to which we might make generalizations on the basis of our Twitter-based findings.

Learning Outcomes

For this guide learning outcomes have been proposed by an editor. Each guide may include up to 5 learning outcomes; therefore, additional learning outcomes can be included if there is space.

It will be beneficial to refer back to these learning objectives when writing your guide – each learning objective must be addressed and it should be evident which sections meet specific objectives.

Learning outcomes must be objective and measurable. Guidance on writing effective outcomes can be found at the following links:

- [Writing learning outcomes](#)

- Bloom's Taxonomy Action Verbs

By the end of this guide, readers should be able to . . .

- gain authorization for the Twitter Academic Research Product Track API
- search and collect counts of tweets
- search and collect tweet content
- search and collect Twitter user data
- [Learning objective 5]

Main text

*Insert your 'How To' Guide text here. The main body of the text – excluding the abstract, learning outcomes, questions, and reference list - should be between **2000 – 4000 words**. **This blank template stands at 1250 words**. Please ensure you have demonstrably fulfilled all the learning objectives.*

Sub-headings add structure to the body of your guide, enhance online discoverability and make your guide easier to read on screen. Consider using sub-headings to signpost which learning outcome you are addressing.

*Each main section with a heading must be followed by a **Section Summary**. Each Section Summary should consist of 2-3 bullet points, written out as full sentences, succinctly encapsulating the preceding section.*

Introduction

Build on the abstract to describe what methodological issues will be discussed in this guide, what the student reader will gain from reading the guide, which real-life research examples will be drawn upon, etc.

In what follows, we will first discuss the first steps you need to complete before collecting and analyzing data from Twitter. Twitter data is not just the text of tweets themselves—we can get information on users, on overall trends in tweets, and on tweet statistics and content (how many likes, retweets, and the type of information they contain). We will walk you through this step by step, explaining, first, how to gain authorization for accessing the Twitter API, second, how to gather Twitter data of different types and, third, what methodological considerations to keep in mind along the way.

Before anything, we need to get authorized for accessing the Twitter API...

Gain Authorization for Twitter API

In order to gain authorization you first need a Twitter account. You can set this up as you would a normal Twitter account. Once you have done this, there remain a few steps before you can become recognized on the Twitter Developer platform and gain authorization for using the Academic Research Product Track API.

First, Twitter will ask for some evidence of your academic profile. They will then ask for details of the proposed research project. Here, questions include the name of the project, how Twitter data will be used, the method involved, how you will distribute results, and whether data will be shared with government entities. This need not be an exhaustive project proposal—indicating the skeleton of the project.

Users should also be aware, here, that irrespective of the non-commercial and/or academic nature of the project, there are certain restrictions to the sharing of raw Twitter data. At present, to comply with the Twitter Terms of Service, users are only permitted to share the so-called “Tweet IDs” of individual tweets when distributing data for replication purposes or otherwise. There are also limits to the number of tweet IDs one can share publicly—1.5m at the time of writing. Though users should check for any updated guidance at the time of conducting their research.

Once Twitter is satisfied of your academic credentials and the nature of the project, you are ready to go! On the Twitter Developer Platform you will be given something called a “Bearer Token,” which you can associate with your Academic Project on the platform. This is all you need to then start collecting data with the R package “`academictwitterR`.”

Given we’re going to be using programmatic code to collect data, and storing the scripts we use, it’s important to think about how we save our authorization credentials. It is generally considered bad form to “hard-code” authorization credentials (i.e., include them in the script itself). Happily, `academictwitterR` offers a way around this.


You can install the package in R with:

```
install.packages("academictwitterR")
```

To use the package, it first needs to be loaded with:

```
library(academictwitterR)
```

What we then need to do is run the `set_bearer()` function. This will open your “.Renvron” file in the home directory. The bearer token can then be inserted here in the following way (note: the token in the image below is not real):



```
1 TWITTER_BEARER=AAAAAAAAAAAAAAAAAPwXWFF1LLDVC6G0PFo4shkDVg02DwVxGQIVKvhPVE3vdV
2
```

Note: you then have to restart R in order for this environment variable to be recognized. Once you have restarted R, you are ready to start using `academictwitterR` to collect Twitter data.

Section summary

- ...
- ...

Search and Collect Counts of Tweets

A good place to start when thinking about collecting Twitter data is by checking out how many tweets match your query. By query, we're here talking about the type of information we're requesting from the Twitter API. Think of this like a search query you would enter into any search engine but with options to specify more accurate parameters.

There are a couple of reasons it makes sense to count the tweets matching your query before doing any serious collection.

First, by looking at over-time trends in tweets that match your criteria, you can very quickly see if there are any abnormalities in the data you collect. As a silly example, let's say we were interested in tweets about the card game "Top Trumps" and we wanted to see how people were discussing this over time. Well, an obvious issue here is that, at a certain point in time, another individual by the name of Trump achieved prominence in the public imagination. We would immediately see from our counts of tweets that either interest in playground card games really soared over time or something else was going on in the data... And we might then need to tweak our query.

Second, by looking at how many tweets match our search query, we can estimate how much time it's going to take us to actually collect all the tweets themselves. On a reasonably fast Internet connection, we can access and download 500 tweets per 2-4 seconds. So if we find that our search criteria match many millions, tens of millions or hundreds of millions of tweets, we know we'll be waiting a long time to capture all of this information.


So how do we capture information on tweet counts? Let's take the following code as an example:

```
tweetcounts <- count_all_tweets(  
  query = "Hogmanay",  
  start_tweets = "2019-12-27T00:00:00Z",  
  end_tweets = "2020-01-05T00:00:00Z",  
  bearer_token = get_bearer(),  
  granularity = "hour",  
  n = 500  
)
```

Here, we are using the function `count_all_tweets()` to collect the information. The query is a simple one—the word “Hogmanay,” which is a Scots word for the last day of the year or New Year’s Eve. We’d expect, then, for there to be many tweets around the time of December 31. And here we’re specifying as “start_tweets” and “end_tweets” the two dates defining the start and end of the period in which we’re interested. And the final part of our query concerns the so-called “granularity” of our request. This can be “minutes,” “hours” or “days” and refers to the time unit over which we are asking Twitter to count up our tweets—i.e., if, as here, we specify hours, Twitter will tell us how many tweets there were every hour over this period matching our search criteria.¹

And if we were to plot this data, we would see something like this:

¹ Note that the other two arguments here – “bearer_token” and “get_bearer” do not need to be changed. In fact, as long as you have stored your authorization token with the `set_bearer()` function then this line isn’t strictly necessary. The `n` here just refers to the upper limit of tweet counts to be fetched; i.e., for this seven day period there are 168 hours (7*24). This means we just need an upper limit of above this number in order to capture all data over this period—500 is specified therefore as an arbitrary upper limit above 168.



, which happily confirms what we thought we might see: a spike in the use of the term at the time of New Year's Eve celebrations.

Section summary

- Use `count_all_tweets()` to count tweets matching your search criteria
- Change the “granularity” to alter the time unit over which you’re counting

Search and Collect Tweet Content

In many cases, the analysis of twitter data does not stop at the over-time trends, we are interested in the content as well as a range of information of the tweets. To search and collect the tweet content, we can use the following code:

```
tweetcontent <- get_all_tweets(  
  query = "Hogmanay",  
  start_tweets = "2019-12-31T00:00:00Z",  
  end_tweets = "2020-01-01T00:00:00Z",  
  bearer_token = get_bearer(),  
  n = 500
```

)

Here, we are using the function `get_all_tweets()` to collect the tweets. The query is the same as the above example. Similarly, “start_tweets” and “end_tweets” are the two dates defining the start and end of the period in which we’re interested. Note that we have specified ‘n = 500’ in our example, meaning that we are only collecting the latest 500 tweets within the search period. If you would like to collect everything within the search period, you can specify ‘n = Inf’ instead.

The returned data will be stored in the object ‘tweetcontent’. In addition to the tweet text, the Twitter API will also return an array of information about the tweets, including the author id, language, geographical information (if available), retweet count, reply count, like count, and so on. Note the authorship information is presented in the form of user id (a numerical identifier), instead of user handle (ie the user name that we use when we @mention). We will talk about this more in the next section.

Instead of searching for a single word, in some cases you might want to build a more sophisticated query. To do so, you can use “operators” to modify your query. For example, you specify “Hogmanay Edinburgh” to search for tweets that contain Hogmanay AND Edinburgh (note empty space means AND in the query), or “Hogmanay OR (New Year)” to search for tweets that contain Hogmanay, or those containing New and Year at the same time.

Alternatively, you can also specify “Hogmanay -Edinburgh” to search for tweets that mention Hogmanay but not Edinburgh. It is impossible to introduce every operator in this short guide, if you would like to know more, you can visit Twitter’s documentation [here](#) for further information or the dedicated [vignette](#) for building complex queries.

Section summary

- Use the function `get_all_tweets()` to collect the tweets
- Use operators to build sophisticated query

Search and Collect Twitter User Data

On top of tweets, “academictwitteR” also incorporates a range of `get_user_*` functions to obtain user-level data. To work with these functions, we first need to know the user id. A user id is a numerical identifier to identify a certain user. Unlike user handles (the id you use when you @ mention someone), the user id cannot be modified by the user. For example, @edhogmanay is the official account of Edinburgh's Hogmanay event organiser—though in future they may alter their account name. To get its user id, we can use the following code:

```
get_user_id(  
  usernames = "edhogmanay",  
  bearer_token = get_bearer(),
```

```
)
```

The function returns “80868086”, this is the unique identifier that we can use in other `get_user_*` functions. For example, to obtain information about a user, we can use the `get_user_profile()` function:

```
get_user_profile("80868086",  
                 bearer_token = get_bearer()  
)
```

This will return the screen name, user name, the creation date of the account, its location, and public metrics such as followers count, following count, and tweet count.

In some cases, we might have a list of users in mind (for example, the major political parties in a given country) and we are interested in what they are trying to communicate on Twitter. We can fetch all the tweets published by the user(s) with the `get_user_timeline()` function. For example, to fetch all the tweets published by `@edhogmanay`, we can use the following code:

```
get_user_timeline(  
  "80868086",  
  start_tweets = "2020-01-01T00:00:00Z",  
  end_tweets = "2020-12-31T23:59:59Z",  
  bearer_token = get_bearer(),  
  n = Inf  
)
```

If you are interested in the following/follower relationship, you can also use the `get_user_following()` and `get_user_followers()` functions to retrieve them.

What if we’re interested in particular types of tweets by particular types of users? Let’s say we’re interested in tweets about the movement “Black Lives Matters” but we want to know what major media organizations are saying about the movement. We can do this by returning to the `get_all_tweets()` function, but now specify both user- and tweet-level queries:

```
tweetsblm <- get_all_tweets(  
  query = "BLM",  
  users = c("CNN", "FoxNews"),  
  start_tweets = "2016-01-01T00:00:00Z",  
  end_tweets = "2020-01-05T00:00:00Z",  
  bearer_token = get_bearer(),  
  file = "data/blmtweets.rds",
```

```
data_path = "data/json_data/",  
n = Inf  
)
```

Here we are asking the Twitter API to return tweets for two major news organizations—CNN and Fox News—that include the term “BLM.” And there are a couple of final things to note here. We’re collecting tweets over a long period of time, meaning that there may be many thousands of tweets being returned (not the case here, but it would be if we were querying, say, tens or hundreds of media outlets). This means we need to store our data sensibly. We are doing this in two ways here. First, the “file” argument means that tweets will be sequentially stored as a serialized “.rds” file, which is an R-specific file extension that allows for efficient data storage. Second, we are specifying a “data_path.” As you collect tweets, you will see this folder being populated with a series of JSON files. These are the files returned by the Twitter API for each API call.² It is a good idea to store data like this when making large API calls (i.e., calls that will return lots of data). This is because sometimes your Internet connection may be cut off, or the Twitter servers may be overloaded. By storing files for each individual API call, we can mitigate any data loss that may result. For

Section summary

- Use `get_user_id()` to obtain user id from user handle
- Use `get_user_profile()` to obtain information about a user
- Use `get_user_timeline()` to fetch the tweets published by a specific user
- Use `get_user_following()` and `get_user_followers()` to fetch the following/followers of a specific user
- Use `get_all_tweets()` with both user- and tweet-level parameters.

Methodological Considerations

Among social media platforms, Twitter is known for its comparatively open data policy. This is especially true after the introduction of Academic Research Product Track API in 2021, which has provided an easy way for researchers to access the full archive of Twitter data. Not surprisingly, Twitter has become the single most popular platform in studies that use social media data (Özkula et al, 2022).

While Twitter data has been proven valuable in a wide range of research, we also need to be mindful of its potential biases caused by the selection of search query or user, platform affordances, and demographic (mis)representations. We do not intend this guide as a comprehensive overview of these issues and would recommend that readers consult Sen et al

² Each API call requests tweets for a “page” of Twitter, with maximum 500 tweets returned per “page.” This is just a feature of the Twitter API and not something most users need to engage with. Large requests will cycle through many thousands of pages.

(2021). That said, it's worth outlining a number of considerations that users will want to keep in mind.

In different ways, all of these considerations concern sampling. And here we need to keep in mind some of the concerns that animate more conventional forms of data collection. In short, we need to know why the sample (of users or tweets) we are collecting accurately represents the phenomenon or group of individuals/organizations in which we're interested.

Let's take the news media as an example. And let's say that we want to know how news media in the UK talk about climate change. We first need to consider who makes it into our sample—our sampling protocol. Specifically, we need to ask how we are selecting our news organizations, and whether our list of accounts cover the gamut of editorial outlooks and geographical locales. This is very similar to how we might design a survey of a specific subsample of the population. The difference is that our sample, instead of being survey respondents, are Twitter users.

We also need to consider how we're going to capture tweets about climate change. Are we just going to search for the phrase “climate change” or do we want to consider other words or phrases? The answer to this depends on how specific we want our analysis to be—climate change as a term or as a phenomenon generally. If the latter, we would want to expand our search queries. This is similar to the considerations we would make when designing survey questions. The difference is that, instead of asking questions, we are specifying search queries to an API. And these can be more or less specific, just as survey questions can be more or less specific.

Next, we might want to consider how different news media organizations use the platform in question. While this is less true of the UK (at the time of writing) some news media have been removed from Twitter in the US, or have boycotted in protest at perceived censorship. If these accounts are therefore not present, or use the platform differently, we then have to ask ourselves: do we simply not sample information from these accounts? Or do we try seek other sources of data for these organizations? The answer to this last question will depend on the level to which we wish to make generalizable inferences. If we are only interested in climate reporting on Twitter, it would be reasonable to include only accounts active on the platform; if we are interested in climate reporting generally, we would want to think about other sources of information. We can think of this problem as a type of non-response. When fielding a survey, some respondents may refuse to give an answer. We can interpret their non-response in different ways, of course—and there is no easy fix. But in essence we are dealing with a similar phenomenon.

Finally, if we are wishing to sample Twitter data over long periods of time, we need to think about how we are going to make valid inferences about over-time trends. When analyzing over-time trends in a phenomenon of interest, we often need to take into account cyclical, seasonal or other trends that might confound our inferences. Take an example from economics. Here, we might be interested in price fluctuations over time. But there might be cyclical trends in these

data: that is, the price of a good might increase every summer because demand increases in warm-weather months. In order to accurately detect any fluctuation separate from these cyclical over-time increases, we would need to take account of these. This is sometimes called “detrending” or, more generally, “normalizing.” A not dissimilar problem obtains with Twitter. If we were interested in the salience of immigration-related discussion online, we might choose to count tweets that include the term “immigration.” But there’s a problem here. The number of Twitter users is going up every day. In other words, there is an over-time trend in the number of users of the platform. This may lead you to infer that immigration talk is increasing when, in fact, it’s just the number of users of the platform that is increasing. So what do we do? Well, if we’re interested in a particular subsample of users, then we can just divide the immigration tweets by the total number of all tweets by these users. If, on the other hand, we’re just counting tweets and not getting tweets of individual users, then we might think of another, more innocuous, word we can count tweets for. This word should be chosen as something that is often mentioned and unlikely to display any cyclical trend: that is, it should “ubiquitous and innocuous” (Dalmaijer et al. 2021). We can then divide our immigration term counts by the counts for this word.

By way of more general conclusion, it is important to note that, when sampling data from Twitter, you must keep in mind the distinction between offline and online, and between one platform and another. It may well be that Twitter accurately reflects the attitudes or temperature of a given (offline) population of interest, but we also know that Twitter users have distinct demographic attributes from the general population (Sloan et al. 2015). And it may be that Twitter accurately reflects the attitudes or temperature of a given (online) population of interest, but we also know that user demographics differ from one platform to another (Auxier and Anderson 2021).

Section summary

- Potential biases can arise from the selection of search query or user, platform affordances, and demographic (mis)representations

Conclusion

*Includes a round-up of the methodological issues discussed in your guide. **What can readers learn from this guide and apply to their own research?***

Through this guide, we hope you learnt enough to go out and start using the Twitter API in your own research. The guide has outlined how to count tweets, collect tweets, and gather information about Twitter users. All of these data types, and ways of collecting them, may be used separately or together in your own research. But it is vital to remember that Twitter data comes with its own store of methodological considerations. And these can be thought of in not dissimilar terms to considerations animating more conventional forms of data collection.

We must always have in mind from where we are gathering our sample, what our sample represents, and what might the threats be to the generalizability of any sample.

Multiple Choice Quiz Questions

[Insert three to five multiple choice quiz questions here. Each question should have only three possible answers (A, B, or C). Please indicate the correct answer by writing CORRECT after the relevant answer.]

Multiple Choice Quiz Questions should test readers' understanding of your entry, and should not require any previous knowledge. Guidance for writing MCQs can be found in the extended manuscript guidelines.

1. What function do we use to capture tweet *content*?
 - a. `get_all_tweets()` CORRECT
 - b. `count_all_tweets()`
 - c. `get_user_profile()`
 2. Why do we specify a data path in `get_all_tweets()`?
 - a. To mitigate data loss CORRECT
 - b. To indicate where our bearer token is stored
 - c. It is required by Twitter Terms of Service
 3. Why might tweet counts be an unreliable measure of online issue salience?
 - a. Because it may just reflect a larger user pool
 - b. Because the Twitter API is unreliable
 - c. Because some people don't have Twitter
 4. ...
 5. ...
-

Further Reading

[Insert list of up to six further readings here]

- Barrie, Christopher and Ho, Justin Chun-ting. (2021). academictwitterR: an R package to access the Twitter Academic Research Product Track v2 API endpoint. *Journal of Open Source Software*, 6(62), 3272, <https://doi.org/10.21105/joss.03272>
- ...
- ...

Web Resources

[Insert links to up to six relevant web resources here]

- <https://github.com/cjbarrie/academictwitterR>
- <https://developer.twitter.com/en/docs/twitter-api/tutorials>
- ...

References

- Auxier, B., and Anderson, M. (2021). Social media use in 2021. *PEW Research Center*, at <https://www.pewresearch.org/internet/2021/04/07/social-media-use-in-2021/>.
- Dalmajer, E. S., Van Rheede, J., Sperr, E. V., amd and Tkotz, J. (2021). Banana for scale: Gauging trends in academic interest by normalising publication rates to common and innocuous keywords. *arXiv:2102.06418* (12 February), at <https://arxiv.org/abs/2102.06418>.
- Özkula, S. M., Reilly, P. J., & Hayes, J. (2022). Easy data, same old platforms? A systematic review of digital activism methodologies. *Information, Communication & Society*, 1–20. <https://doi.org/10.1080/1369118X.2021.2013918>

Sen, I., Flöck, F., Weller, K., Weiß, B., & Wagner, C. (2021). A Total Error Framework for Digital Traces of Human Behavior on Online Platforms. *Public Opinion Quarterly*, 85(S1), 399–422. <https://doi.org/10.1093/poq/nfab018>

Sloan L, Morgan J, Burnap P, Williams M. (2021). Who Tweets? Deriving the Demographic Characteristics of Age, Occupation and Social Class from Twitter User Meta-Data. *PLOS ONE* 10(3): e0115545.
<https://doi.org/10.1371/journal.pone.0115545>

References should conform to American Psychological Association (APA) style, 7th edition, and should contain the digital object identifier (DOI) where available. SAGE will not accept guides that are incorrectly referenced. Please ensure accuracy before submission. For help on reference styling see <https://apastyle.apa.org/style-grammar-guidelines>.

- ...
- ...
- ...