THE UNIVERSITY *of* EDINBURGH

School of Physics
and Astronomy

Summer Project Report
ArXiv Citation Graph

Matej Vedak
August 2022

Abstract

A reference metadata extraction script is created. The source of the scientific papers that are processed is arXiv. A paper is downloaded, the references are found and processed for metadata. ArXiv IDs and DOIs are looked for first. If none could be found, an online service called Crossref is utilized to match references. A majority of the references are successfully resolved and correct metadata is extracted.

Declaration

I declare that this project and report is my own work.

Signature: Date: 25/08/2022

Supervisor: Florian Beutler

Contents

1 Introduction

1.1 ArXiv

ArXiv[1] is a free paper distribution service and an open-access archive for scholarly articles in the fields of physics, mathematics, computer science, quantitative biology, quantitative finance, statistics, electrical engineering and systems science, and economics. It works on principles of openness, collaboration and scholarship which provide a strong foundation which has attracted dozens of world researchers. It was founded in 1991 and is now maintained and operated by Cornell Tech.

Currently, it hosts more than 2 000 000 scholarly articles. Given its popularity and abundance of articles, it is very popular among the physics and astrophysics community and many other online services utilize arXiv to create reading lists and online gatherings of researchers.

1.2 Benty-Fields

One such online service is a website called Benty Fields [2]. It was created and founded by Dr Florian Beutler and Dr Morag Scrimgeour. Florian is currently the main developer. The website's purpose is to help researchers organize and find the latest publications most relevant to their research. It recommends interesting, just-released papers and hosts Journal Clubs where researchers can recommend papers and discuss previously read articles. There are other features such as an integrated job search and a list of upcoming seminars.

Benty-Fields is closely integrated with arXiv, in fact, arXiv is its only source of papers. It is meant to be sort of an extension of arXiv.

1.3 Problem description

The main feature we are planning on adding to Benty-Fields is reference extraction from arXiv papers. This would greatly improve Benty-Fields functionality as users could instantly get access to references of a paper they are reading without having to manually search them up. Each paper references several other publications, publications that might or might not be within the arXiv itself. Furthermore, there is no standard for writing out references which additionally complicates the problem.

In this project, we are trying to create a robust python script that will detect references within each paper and extract metadata about those references from other online sources. We are primarily focused on references that reference other arXiv papers but the final goal is to have a system that extracts metadata for all references, whether they are published on arXiv or elsewhere. Once this reference extraction works, we can further improve the website by creating a citation graph (sort of a map of referencing papers) and implementing machine learning algorithms that recommend similar papers. The end goal of all this is to implement these features on the Benty-Fields website.

2 Methods

2.1 Methods overview

The main idea of the project is to use the arXiv API to gather metadata about the referenced publications. Every paper on arXiv has a unique ID. Knowing the ID (paper_id) of a paper it can easily be downloaded using the following URL:

`http://export.arxiv.org/e-print/paper_id`.

The downloaded files are processed and unique reference identificators are searched for - either a DOI (Digital Object Identifier) or an arXiv ID. The focus is on finding arXiv IDs as arXiv is the primary web resource we are utilizing, but DOIs are just as useful. The main tool in processing pure references and extracting DOIs and arXiv IDs is RegEx (Regular Expression [3]). Unfortunately, a vast majority of references contain neither a DOI nor an arXiv ID. Usual methods of resolving such references involve machine learning and NLP (natural language processing) algorithms to extract information such as author, title, journal, etc (see, for example, [4]). Luckily, there exist online services which do this reference matching, eliminating the need of implementing local machine learning algorithms. One such service is Crossref [5]. The key service they provide is unstructured reference matching. Essentially, a user dumps a full reference in Crossref's matching algorithms and they return the best match for the given reference. Having either an arXiv ID, DOI, or the option of utilizing Crossref the script is almost complete. The only thing left to do is to store the extracted reference metadata. To this end, we utilize SQLite [6].

To summarize, the full code works as follows:

1. Download the paper from arXiv

2. Scan the raw text and extract references (they are easily found in .bbl or .tex files since every reference starts with `\bibitem`

3. For each detected reference try to find an arXiv ID or a DOI

   - If an arXiv ID is found: extract reference metadata from arXiv
   - If a DOI is found: extract reference metadata from Crossref
   - If neither is found: use Crossref's unstructured reference matching system to get the best match and use its metadata

4. Store the extracted data in a database

2.2   API basics

What exactly is an API?

API [7] stands for an Application Programming Interface. Overall, it's a way for two or more computer programs to communicate with each other. This is in contrast to a user interface, which connects a computer to a person. API is not intended to be used directly by a person (the end-user) apart from software engineers who incorporate an API into the software. An API is usually made up of different parts which act as tools or services that are available to the programmer. The calls that make up the API are known as subroutines, methods, requests, or endpoints. An API needs to have a specification, a set of instructions which defines the previously mentioned calls and explains how to use or implement them.

One of the main purposes of an API is to hide the internal details of how a system works, allowing only the interaction with those parts a programmer will find useful. The details of interaction stay consistent even if the internal implementation changes at a later date. This allows for longer lifespans of applications without the need for constant updating.

There are many different types of APIs, but the one that is used throughout this project is a Web API. It is a service accessed from client devices (laptops, mobile phones, ...) which communicate with a web server using the HTTP (Hypertext Transfer Protocol) [8]. Client devices request something using an HTTP request, and the server provides them with a response message, often in the form of JSON (JavaScript Object Notation) [9] or XML (Extensible Markup Language) [10] format. The most frequent use of web APIs is to query a server for some specific data.

2.3   ArXiv API

Let's see how one of these Web APIs looks in practice. One of the APIs we have made use of in this project is the arXiv API [11]. Generally, the majority of web APIs use a syntax similar to the following:

```
http://export.arxiv.org/api/{method_name}?{parameters}.
```

For example, to search all of arXiv for the term 'electron' the following URL is entered:

```
http://export.arxiv.org/api/query?search_query=all:electron.
```

Here, the `method_name` is query and the parameters are `search_query=all:electron`. The server will process this request and send back a response. ArXiv API returns an XML response which can be opened in any web browser (see figure 1).

As we can see, a typical response contains data about the query and certain API specifications, followed by a list of <entry> which are documented papers within arXiv. Between <entry> and </entry> tags, metadata about a certain publication is contained. We can easily extract data such as publication title, authors, date of publishing, and whatever else the server provides us with.

Of course, the `search_query` is not the only parameter we can supply. Other parameters are:

```
<?xml version="1.0" encoding="UTF-8"?>
<feed xmlns="http://www.w3.org/2005/Atom">
  <link href="http://arxiv.org/api/query?search_query%3Dall%3Aelectron%26id_list%3D%26start%3D0%26max_results%3D10" rel="self" type="application/atom+xml"/>
  <title type="html">ArXiv Query: search_query=all:electron&amp;id_list=&amp;start=0&amp;max_results=10</title>
  <id>http://arxiv.org/api/WyBPOs+pRgzCTXTMWhtnbcOmk6g</id>
  <updated>2022-07-15T00:00:00-04:00</updated>
  <opensearch:totalResults xmlns:opensearch="http://a9.com/-/spec/opensearch/1.1/">185633</opensearch:totalResults>
  <opensearch:startIndex xmlns:opensearch="http://a9.com/-/spec/opensearch/1.1/">0</opensearch:startIndex>
  <opensearch:itemsPerPage xmlns:opensearch="http://a9.com/-/spec/opensearch/1.1/">10</opensearch:itemsPerPage>
  <entry>
    <id>http://arxiv.org/abs/cond-mat/0102536v1</id>
    <updated>2001-02-28T20:12:09Z</updated>
    <published>2001-02-28T20:12:09Z</published>
    <title>Impact of Electron-Electron Cusp on Configuration Interaction Energies</title>
    <summary>  The effect of the electron-electron cusp on the convergence of configuration
interaction (CI) wave functions is examined. By analogy with the
pseudopotential approach for electron-ion interactions, an effective
electron-electron interaction is developed which closely reproduces the
scattering of the Coulomb interaction but is smooth and finite at zero
electron-electron separation. The exact many-electron wave function for this
smooth effective interaction has no cusp at zero electron-electron separation.
We perform CI and quantum Monte Carlo calculations for He and Be atoms, both
with the Coulomb electron-electron interaction and with the smooth effective
electron-electron interaction. We find that convergence of the CI expansion of
the wave function for the smooth electron-electron interaction is not
significantly improved compared with that for the divergent Coulomb interaction
for energy differences on the order of 1 mHartree. This shows that, contrary to
popular belief, description of the electron-electron cusp is not a limiting
factor, to within chemical accuracy, for CI calculations.
</summary>
    <author>
      <name>David Prendergast</name>
      <arxiv:affiliation xmlns:arxiv="http://arxiv.org/schemas/atom">Department of Physics</arxiv:affiliation>
    </author>
    <author>
      <name>M. Nolan</name>
      <arxiv:affiliation xmlns:arxiv="http://arxiv.org/schemas/atom">NMRC, University College, Cork, Ireland</arxiv:affiliation>
    </author>
    <author>
      <name>Claudia Filippi</name>
      <arxiv:affiliation xmlns:arxiv="http://arxiv.org/schemas/atom">Department of Physics</arxiv:affiliation>
    </author>
    <author>
      <name>Stephen Fahy</name>
      <arxiv:affiliation xmlns:arxiv="http://arxiv.org/schemas/atom">Department of Physics</arxiv:affiliation>
    </author>
    <author>
      <name>J. C. Greer</name>
      <arxiv:affiliation xmlns:arxiv="http://arxiv.org/schemas/atom">NMRC, University College, Cork, Ireland</arxiv:affiliation>
    </author>
    <arxiv:doi xmlns:arxiv="http://arxiv.org/schemas/atom">10.1063/1.1383585</arxiv:doi>
    <link title="doi" href="http://dx.doi.org/10.1063/1.1383585" rel="related"/>
    <arxiv:comment xmlns:arxiv="http://arxiv.org/schemas/atom">11 pages, 6 figures, 3 tables, LaTeX209, submitted to The Journal of
Chemical Physics</arxiv:comment>
    <arxiv:journal_ref xmlns:arxiv="http://arxiv.org/schemas/atom">J. Chem. Phys. 115, 1626 (2001)</arxiv:journal_ref>
    <link href="http://arxiv.org/abs/cond-mat/0102536v1" rel="alternate" type="text/html"/>
    <link title="pdf" href="http://arxiv.org/pdf/cond-mat/0102536v1" rel="related" type="application/pdf"/>
    <arxiv:primary_category xmlns:arxiv="http://arxiv.org/schemas/atom" term="cond-mat.str-el" scheme="http://arxiv.org/schemas/atom"/>
    <category term="cond-mat.str-el" scheme="http://arxiv.org/schemas/atom"/>
  </entry>
</feed>
```

Figure 1: ArXiv API's response to the following query: search_query=all:electron

- `id_list` - extract data using arXiv IDs to specify papers

- `start` - modifies the search query, specifies what from what result to return entries

- `max_results` - modifies the search-query, specifies how many results to return

Furthermore, the `search_query` parameter can be customized. We won't go into detail but apart from the 'all:' prefix, other prefixes exist such as 'ti:' (title), 'au:' (author), etc.

2.4 Regex - Regular Expressions

Having an extracted reference is still a long way from identifying the publication the reference is actually pointing to. An extracted reference is a jumble of symbols and words, especially because we extract pure BibTeX references (so they are filled with symbols like '{', '[', '\newgroup'). Extracting complex IDs and DOIs from such a complicated string seems like an impossible task. Luckily, this is where regular expressions come into play.

A regex is a sequence of characters that specify a search pattern in the text. Combined with string-searching algorithms, they allow for easy 'find' or 'find and replace operations on strings. See the paper on one of the first regex compilers published in 1968 [12] for more information. A more modern introduction can be found on Wikipedia [3].

Perhaps expand on regex syntax here

DOIs are fairly standardized identificators (see [13] for details of construction and history). Most of them can be caught with a single regex expression [14]:

```
10.\\d{4,9}/[-._;()/:a-z0-9A-Z]+
```

Matching arXiv IDs is a little bit harder though [15]. On the 1st of April, 2007, the identifier scheme has been completely changed to accommodate the popularity growth of the service and whole website. Identifiers up to March 2007 had the following general form

```
archive.subject_class/YYMMnumber(three digits)
```

For example `math.GT/0309136`. The identifier scheme has been changed to a more general form in April 2007:

```
arXiv:YYMM.number{vV}
```

Example: `arXiv:1501:00001` or `arXiv:0706.0001v2` ('v2' stands for version 2). Luckily, other professional developers have already run into the same problem and created good matching patterns [16]. To put into perspective, the regex pattern that matches arXiv IDs is 6742 characters long. Compare that to the pattern that catches most DOIs and you can see how much more complicated it is.

2.5   ArXiv ids percentage analysis

To see what kind of numbers and performance we can expect, before starting any database creation we analysed the distribution of DOIs and arXiv IDs within a random subset of 100 papers. Surprisingly, only 14.4% of references contain either an arXiv ID or a DOI.
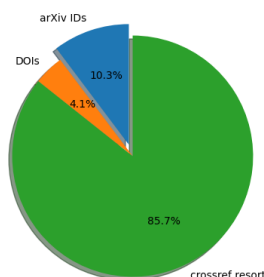

Figure 2: Percentage of references that contain an arxiv ID, percentage of references that contain a DOI, and percentage that contain neither and some other method has to be utilized ('crossref resort'). These are the results for 100 random papers taken from arXiv, with various publication years.

To further investigate the percentage of arXiv IDs, we decided to analyse references by year. The hypothesis was that as the years went by, and arXiv IDs became better established among the scientific community, the number of references that contain an arXiv ID will go up and, subsequently, the arXiv ID percentage will grow.

To that end, we decided to analyse 50 arXiv papers for each year of its existence (1992 - today). Actually grabbing the paper IDs turned out to be a significant challenge. ArXiv's API doesn't have a 'select by year' query parameter which complicates things. To get around this limitation, the best we could do was utilize some other selectors and keep grabbing papers until we get enough papers for every year. Naturally, just grabbing all the papers would take forever - there are around 2 000 000 papers on arXiv. For this reason, we chose specific categories - that contain a more reasonable amount of papers.

Categories of experimental high-energy particle physics, & general relativity and quantum cosmology were chosen.

A special parameter was provided in a GET request which tells the server to return the papers either in ascending or descending year order. The issue of having to go through all the papers in a category remained. ArXiv servers didn't help in this matter, as they would frequently 'lock' the requests out and would stop returning any data whatsoever. Other times, the servers would randomly return an empty XML file - but would continue responding to subsequent requests. Initially, the request grabbed 1000 entries at once. The problem with this was if the server refused to return any data a full year of papers might have been skipped. In addition, the servers were not too happy with the size of this request and would frequently stop sending responses. To remedy the situation, we reduced the number of instances per request to 100, which improved the servers' responsiveness, albeit at a much longer total time of processing.

The results were pretty much what was expected. The reference percentage increases through the years as the arXiv IDs became more standardized. However, the maximum average percentage increases up to around 50%. It looks like referencing varies widely depending on the author: either they put arXiv IDs in almost every reference or almost none at all. To make sure this is not category-specific, we analysed two different categories (see figures 3 and 4).

2.6   Crossref

What about references that have no identifier within them? Usually, such references are matched by a combination of natural language processing and machine learning algorithms which parse the reference and extract the publication's title, authors, and any other useful information. Luckily, Crossref offers this service eliminating the need for the implementation of complex machine learning algorithms. Actually, Crossref doesn't utilize machine learning algorithms, their algorithm uses a search algorithm which tries to match the unprocessed reference to one of the entries in Crossref's database.

To interact with the Crossref's servers, and to utilize its API we used a premade python library called habanero [17]. We extracted various data from Crossref. This included the publication's DOI, title, authors, URL, and many more. An interesting, and important, piece of information is a score Crossref returns upon matching an unstructured reference
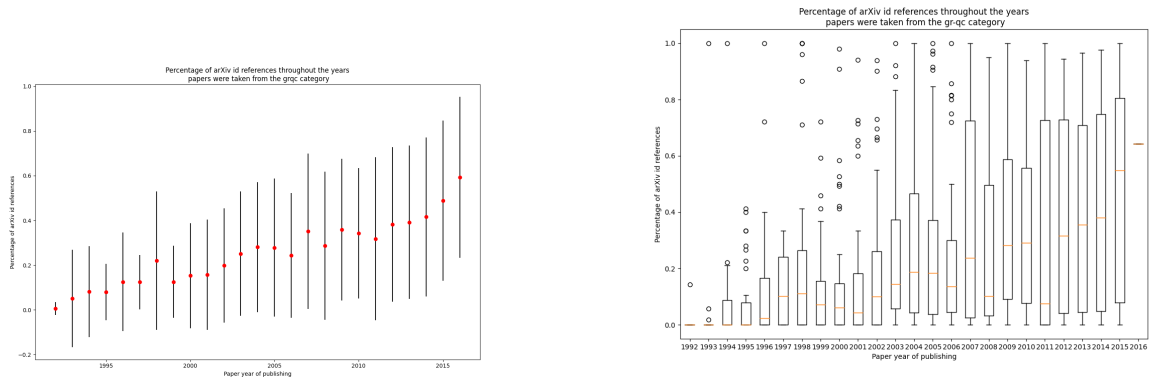
Figure 3: On the left - mean and standard deviation of the arXiv IDs percentage among all references within 50 papers for every year, general relativity and quantum cosmology category. An upward trend can be seen, especially in recent years. The data stops in 2016 because the arXiv server refused to return any data for subsequent requests.

On the right - candlestick chart of arXiv IDs percentage among all references within 50 papers for every year. The thin red line represents the median. Within the box, 50% of all data is contained, while within the whole big structure (box plus 'whiskers', thin lines) 99% of the data is contained. Empty circles are 'outliers'. An upward trend can be seen. The data chart doesn't go past the year 2015 because the arXiv server refused to return any data for subsequent requests.
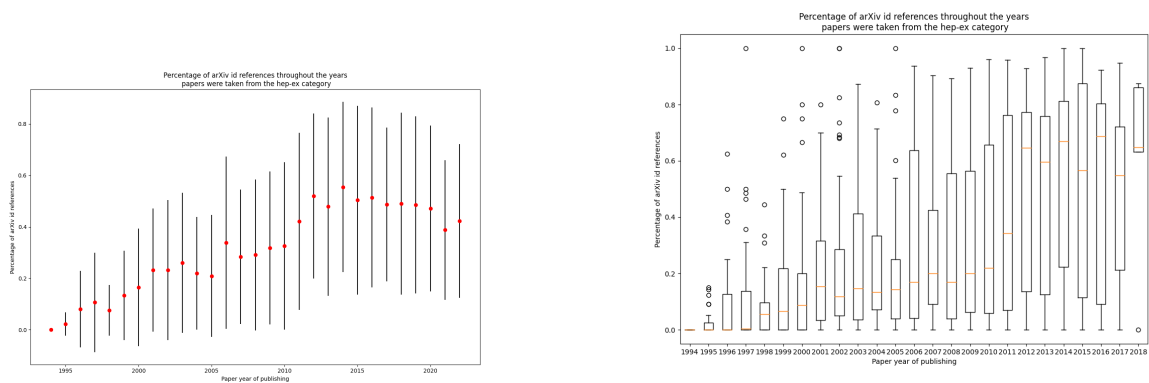

Figure 4: On the left - mean and standard deviation of the arXiv IDs percentage among all references within 50 papers for every year, experimental high energy particle physics category. An upward trend that gets saturated at about 50% is visible.

On the right - candlestick chart of arXiv IDs percentage among all references within 50 papers for every year. The thin red line represents the median. Within the box, 50% of all data is contained, while within the whole big structure (box plus 'whiskers', thin lines) 99% of the data is contained. Empty circles are 'outliers'. An upward trend can be seen. The plot is cut off in the year 2018 because of server issues.

to a publication in the online database. A score is a metric of how certain of a match the algorithm is. As of the writing of this publication, we couldn't find any details on how this score metric works. Manual inspection of this score showed that matches with a score above 50 are probably correctly matched. A score between 30 and 50 indicates some level of uncertainty. A score below 30 indicates an unlikely match.

2.7 Databases - SQLite

To store the extracted references we used an SQL database engine called SQLite (the final version of the application that is going to be deployed on the website probably won't use SQLite - it is just a placeholder). It is a C-language library that implements a small, fast, self-contained, high-reliability, full-featured, SQL database engine [6]. SQLite is an embedded SQL database engine and, unlike most other SQL databases, SQLite does not have a separate server process. This is the reason SQLite was ideal for testing the already written code.

The version used in the code is a python-adapted version called sqlite3. For the purposes of this application we only really needed two SQL statements. The first one is the standard table creation statement

```
cursor.execute("CREATE TABLE {} ({})".format(tableName, ",".join(
    FIELDS_TO_STORE)))
```

where the tableName is what the table should be called, and the FIELDS_TO_STORE is a list containing the names of the desired columns. The second one is the standard data insertion statement where we inserted the extracted metadata for a single reference of a particular paper:

```
cursor.execute("INSERT INTO {} VALUES ({})". format(tableName,
    QUESTIONMARKS), tuple(md.values()))
```

where QUESTIONMARKS is a specially formatted list of question marks, which is necessary since SQLite has built-in defences against SQL injections and this is how it is accessed, and the `tuple(md.values())` is just the unpacking of the collected metadata for a reference.

3 Details of Code Implementation

The problem we are trying to solve here quickly turned out to be much harder than anticipated. The main script has more than 1500 lines of code, and to manage that much code some sort of version control system had to be used. We used the standard git and all of the code can be found in this GitHub repository: https://github.com/MatejVe/arxiv_citation_graph.

The code is structured in such a way that the end-user of this script only needs to use one function. That function is `create_database`. It takes the following arguments

- `dbname` - the name of the database in which the data will be stored.

- `tableName` - the name of the table in which the data will be stored.

- `arxivIDs` - list of arXiv papers that are to be processed.

- `mode` - default is "restAPI", other options are "neither" and "xmlAPI".
  "restAPI" uses Crossref's JSON REST API to process references that have no identifier within them.
  "neither" only processes the references with an identifier and ignores the rest.
  "xmlAPI" uses Crossref's XML API to process references that have no identifier within them.

- `email` - defaults to none but really should be provided. When using the XML API it is essentially required, although there is a workaround and a false e-mail can be provided (there are no checks on the server side). When using the JSON REST API a provided e-mail allows for access to the 'polite' pool of servers, better behaving servers than the 'public' pool (as claimed by Crossref, I have found that the 'polite' servers can be in fact slower than the public servers; however when using the public servers one can get its IP address banned if their usage overwhelms the servers frequently).

- `threshold` - defaults to 0, signals how certain the JSON REST API has to be in its matching algorithm. If the score of a matched reference is below this threshold the matching will be void and the reference will be flagged as "unresolved". No formal specifications of this score could be found. The author of this paper [4] claims that only scores above a 120 are absolute matches but we could find no official documentation to support this claim. From manual inspection, references with a score below 30 are most likely not a good match, references with a score between 30 and 50 are a mix of correct and incorrect matches and references with a score above 50 are most likely a good match.

- `xmlQueryNum` - argument that controls how many references are to be sent at once to the XML server for processing. Only applicable if the mode is "xmlAPI". 5 references at once is the default.

- `xmlUseRESTAPI` - boolean argument and it defaults to True. This argument controls whether we use the JSON REST API to process references that the XML server couldn't resolve (XML API doesn't have a 'score', it either matches a reference or it doesn't).

For documentation purposes, here is a list of all current fields that the script tries to extract metadata for.

- `paper_id` - arXiv id of the paper where the reference was found.

- `reference_num` - number of the reference as it was found.

- `status` - can be "resolved", "unresolved" or "ignored". Resolved is for fully processed references, unresolved if there was an error or the servers returned an unresolved reference and ignored is for references where some external processing wasn't even tried.

- `ref_arxiv_id` - arxiv identifier of the reference, if there is one.

- `DOI` - DOI identifier of a reference, if there is one.

- `title` - title of the reference.

- `abstract` - abstract of the referenced work; arXiv API is the only one that provides this.

- `authors` - authors of the referenced work.

- `database_name` - the name of the database where the work is stored. Currently only gets populated for arXiv references with the value "arxiv".

- `last_modified_datetime` - date and time of the last modification of the referenced work.

- `authors_linked` - author string including links to the profile pages of users, not implemented.

- `authors_linked_short` - authors string trimmed to 8 authors including an et al. Not implemented.

- `pdf_link` - direct link to a PDF of the reference. Mostly can only be found for arXiv references.

- `source_link` - URL pointing to the website of the reference, where the reference can be found.

- `arxiv_comment` - comment connected to an arXiv paper. Can contain information on where the paper was published.

- `arxiv_primary_category` - primary category of an arXiv paper. For example hep-ex, math, ...

- `type` - what type of a reference is it: journal article, book, conference paper, ...

- `container` - where was the reference published, can be journals, books etc.

- `score` - a score that indicates how certain the matching algorithm is in this particular match.

- `length_of_bibitem` - length of the cleaned up BibTex text of the reference.

- `time_taken` - time taken to get a response from the server that was queried.

- `bibitem` - the full text of the BibTex entry for reference.

- `clean_bibitem` - the full text of the cleaned version of the BibTex entry for reference.

Depending on the mode of metadata extraction, some of these fields will remain unpopulated because there are differences in what data is provided between the XML API and the JSON REST API. Furthermore, Crossref simply doesn't have metadata for certain fields for many references. This is why all the fields are populated with 'Null' by default - and only updated when the data is found. For the purposes of metadata 'standardization', we utilized a special type of dictionary, defined in the script. It is a dictionary which has keys defined on instance creation and new keys can't be added later. This was done to prevent spelling mistakes and to make data extraction easier. Keys that are created are defined in a list called `FIELDS_TO_STORE`. The dictionary is defined as

```python
class AttrDict(dict):
    """
    Type of dictionary that can't have new keys amended after
instance creation.
    A caveat is that the initialization of the dictionary is a bit
different:
    keys = ["a", "b"]
    values = [1, 2]
    dict = AttrDict(zip(keys, values))

    The idea is to set a standard for a metadata dictionary so that
later mistakes in assigning data can be avoided.
    """

    def __init__(self, *args, **kwargs):
        super().__init__(*args, **kwargs)
        self.__dict__ = self

    def __setattr__(self, key, value):
        if key not in [*self.keys(), "__dict__"]:
            raise KeyError(f"No new keys allowed - {key}")
        else:
            super().__setattr__(key, value)

    def __setitem__(self, key, value):
        if key not in self:
            raise KeyError(f"No new keys allowed - {key}")
        else:
            super().__setitem__(key, value)
```

Another important function in the code is the `clean_up_bibtex` function which takes in a pure BibTex entry of a reference and tries to create a reasonable, readable reference out of it. This cleaned-up BibTex is used for Crossref queries and it resulted in a significant increase in the average matching score. For example, given a pure BibTex entry such as

```
{Hamkins1} J.~Hamkins and K.~Zeger.
\newblock Asymptotically dense spherical codes.
i. wrapped spherical codes.
\newblock 43(6):1774--1785, Nov. 1997.
```

get transformed into

```
J. Hamkins and K. Zeger. Asymptotically dense spherical codes.
i. wrapped spherical codes 43(6):1774--1785, Nov. 1997.
```

There is a caveat. For some references, this cleaning up causes a loss of information. Some information from the reference gets discarded. On average though, cleaning up increases the matching accuracy of Crossref. For completion, here is the full `clean_up_bibtex` function.

```python
def clean_up_bibtex(bibitem: str) -> str:
    def remove_from_beggining(openTag, closeTag, string):
        if string[0] == openTag:
            i = 1
            while string[i] != closeTag:
                i += 1
            return string[i + 1 :]
        return string

    # Many bibtex items start with the local reference name enclosed within the '{}'
    # brackets or some other text within '[]' brackets
    # What we do here is run '{}' removal and then '[]' removal
    # or run '[]' removal and then '{}' removal
    if bibitem[0] == "{":
        bibitem = remove_from_beggining("{", "}", bibitem)
        bibitem = remove_from_beggining("[", "]", bibitem)
    elif bibitem[0] == "[":
        bibitem = remove_from_beggining("[", "]", bibitem)
        bibitem = remove_from_beggining("{", "}", bibitem)

    # Clean up latex items like \em{} etc
    pattern = re.compile(r"\\[A-z]+{")
    bibitem = re.sub(pattern, "", bibitem)
    # Clean up item like \newblock (items which don't utilize '{}')
    pattern = re.compile(r"\\[A-z]+")
    bibitem = re.sub(pattern, "", bibitem)
    # Remove '\n'
    pattern = re.compile(r"\n")
    bibitem = re.sub(pattern, "", bibitem)
    # Clean up simple characters: {}[]
    reduntant_characters = "{}[]\"'%"
    bibitem = bibitem.translate(
        {ord(char): None for char in reduntant_characters}
    ).strip()
    # Remove the weird ~ symbol
    # I have no idea what it is even used for in bibtex
    # Replace with whitespace though
    for i in range(len(bibitem)):
        if bibitem[i] == "~":
            bibitem = bibitem[:i] + " " + bibitem[i + 1 :]
    # Remove redundant white space
    pattern = re.compile(r"\s{2,}")
    bibitem = re.sub(pattern, " ", bibitem)

    return bibitem
```

The main control flow of the whole script is defined in the `get_citations` function which takes the following arguments: `list_of_files`, `mode`, `email`, `threshold`, `xmlQueryNum`, `xmlUseRESTAPI`. All the arguments are passed from the `create_database` function and the user doesn't need to manually tune them here.

The function starts with the opening of the file that might contain references. It looks

for the `\bibitem` string which indicates whether there are any references within the file. All the references text is extracted, from the first occurrence of the `\bibitem` to the `\end{thebibliography}`. Some light preprocessing is done (no changing of the actual citation data).

Found references are enumerated and processed one at a time. First, we try to find a DOI or an arXiv ID. A DOI takes precedence over an arXiv ID and the DOI is used to extract metadata if both are found. ArXiv IDs and DOIs that are found within the files are always used to extract metadata. Now comes the control flow:

- If the mode argument is "neither" arXiv IDs and DOIs that are found within the references are used to extract metadata for corresponding references, but this is where the script stops. Any references without an identifier within them are marked as "ignored" and aren't processed anymore. DOI metadata is extracted through Crossref (function `crossref_metadata_from_doi`) and arXiv ID metadata through arXiv's REST API (function `arxiv_metadata_from_id`).

- If the mode argument is "restAPI" all the references that don't have an identifier are processed through Crossref's JSON REST API. For easier interaction with the API, a python module called habanero was utilized [17]. The function that is used is called `crossref_metadata_from_query`. Its usefulness lies in the fact that Crossref can match unstructured citations. The script takes the BibTeX text of the reference, and cleans it up (see function `clean_up_bibtex`) and only then sends it to Crossref for matching. Crossref returns a potential match to the reference which comes with an attached score, a measure of certainty in the matching. See above for details of this score. There is also an email argument which controls whether the request is sent to a pool of 'public' servers or a pool of 'polite' servers.

- If the mode argument is "xmlAPI" all the references that don't have an identifier are processed through Crossref's XML API. The `xmlQueryNum` argument controls how many references at once are sent to the server. XML server doesn't have a score attached to the matchings, it either matches a reference or doesn't. That is why there is an additional argument, `xmlUseRESTAPI`. It signifies that the references that couldn't be matched with the XML API are to be matched using the JSON REST API.

This covers the main and most important parts of the script. There are many more details and caveats within the source code. It has been thoroughly commented on and should be readable.

There are a few "TODO:" comments left out in the code. These are the places where future improvements could be implemented.

4 Discussion

With the code working and free of errors that would break the script, we analysed the speed and efficiency of different modes the script can run in. We conducted the analysis on a random sample of 100 arXiv papers.
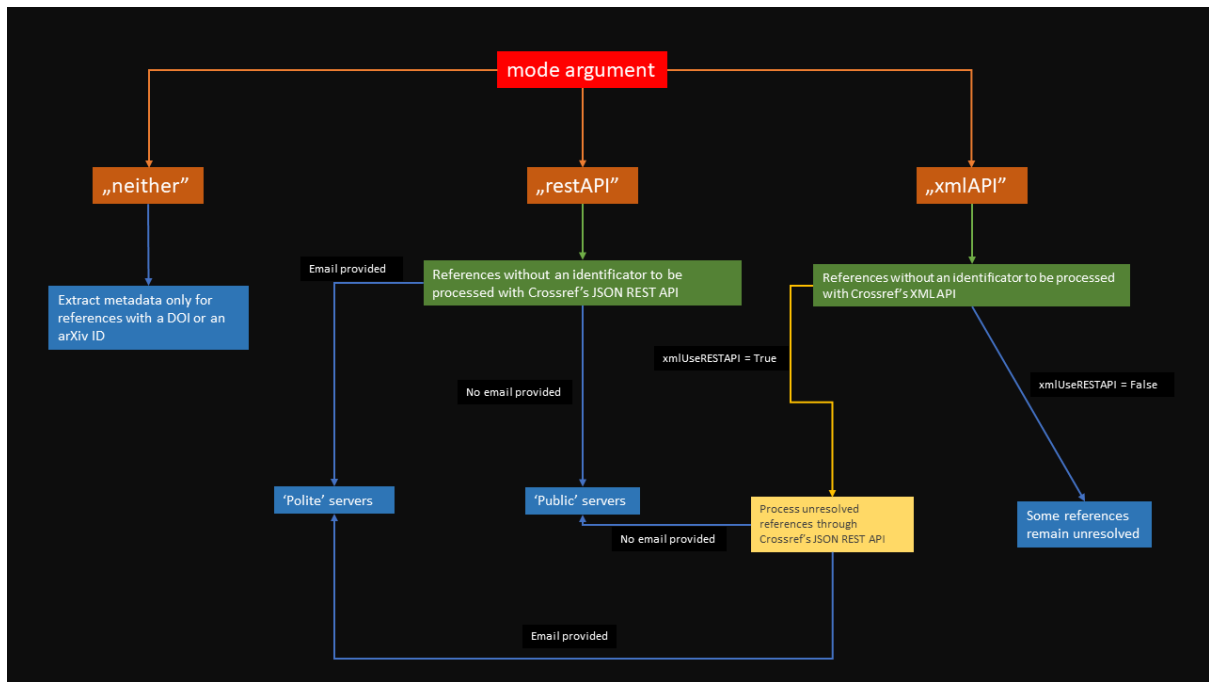
Figure 5: Schematic representation of the metadata extraction options.

Mode "neither" - only looking for arXiv IDs and DOIs, ignoring references with no identifier.

- Time taken was 815s or 0.23 hours.

- Estimated time to process 2000 papers is 4.53 hours.

- 623 out of 3805 references have an arxiv id or a DOI - 16.4%.

Utilizing Crossref's REST API, no email was provided.

- Time taken was 3657s or 1 hour.

- Estimated time to process 2000 papers is 20 hours.

- 2 references remain unresolved; a score less than 30 is found for 340 references; a score less than 50 is found for 1372 references. There are 3183 references that had to be run through Crossref. This means that 10.6% of references have a score less than 30, these matches are very unlikely to be correct. 43.1% have a score less than 50 - 32.5% have a score between 30 and 50. These 32.5% of references are mostly correctly matched (from what I have seen).

Utilizing Crossref's REST API, email provided.

- Time taken was 6021s or 1.67 hours.

- Estimated time to process 2000 papers is 33.4 hours.

- Score percentages are the same as for when no email was provided.

Utilizing Crossref's XML API, error 400 references were thrown back into the JSON REST API, 5 references per query. Note that there were quite a few 400 errors thrown back by the XML server. These are probably caused by some characters within the bibitems that the XML server is not happy with. We couldn't identify all of these characters.

- Time taken was 10839s or 3.01 hours.

- Estimated time to process 2000 papers is 60.22 hours.

- 803 out of 3805 references remain unresolved (21%) although these are only unresolved references that we successfully queried through the XML API. Had the error 400 references been included in this calculation, the ratio would have probably been a bit bigger.

XML API, no processing of unresolved references.

- Time taken was 5867s or 1.63 hours.

- Estimated time to process 2000 papers is 32.6 hours.

- 1818 out of 3805 references are resolved (47.8%).

We provide the estimated time to process 2000 papers because that is the approximate number of papers that get released every day on arXiv. It is important that this number is below 24 hours, otherwise, the paper processing speed is simply not good enough.

There are a few options on how to increase the processing speed for all implemented methods. The simplest one is to add an extra thread that processes the references. The speed bottleneck is neither the raw data processing nor the speed of the data traversing the internet to reach the servers. By far the most time-consuming bit is the Crossref server queries processing time. If this could be parallelized the speed increase could be n-folded, depending on the number of threads utilized. Of course, there is the possibility of overloading the servers, but, with the email address provided, Crossref's team would be in touch and some solution would be agreed upon.

Another option is to process only references with an arXiv ID or a DOI for newly released papers, while another script slowly processes the remaining references without an identifier. Since newly uploaded papers' references do contain some identifiers this method should provide a solid user experience where new papers have most of their references linked, while the older papers, papers with references without identifiers, are slowly getting updates for their references. Probably even a single thread could be used in this scenario, new papers will get processed rather quickly, in about 5 hours, and the rest of the day (another 19 hours) could be used to slowly extract metadata through Crossref for the other references

4.1 Future work - implementation of NLP

This project could further be improved with the addition of some sort of Natural Language Processing in combination with Machine Learning to find metadata within the references themselves, without any need for external resources. Unfortunately, this route was not pursued in this project, due to a lack of time, and it was deemed that the usage of external resources would yield better results and a better user experience, once the application is implemented on a website.

Reasons for the decision against NLP are various. Firstly, pure BibTex reference entries can be a mess. Various BibTex writing styles can be found, with vastly different practices between researchers. Proper tokenization of the references from BibTex would be a very, very big challenge. This issue could be solved by extracting the references from the PDF instead of the BibTex and LaTex files. The issue that pops up here is how to identify references in a block of free text. Finding where they start is doable (just find the last occasion of the word "References" and its synonyms, and hope that there is no mention of the word in the appendix), however, finding the end of the references is the bulk of the problem. Perhaps a machine learning algorithm would be able to do this properly.

Secondly, the information provided by the references is limited. A typical good reference contains the authors' names, the title, the journal where it was published, the publication date, a DOI and a URL. So the metadata extracted from a good reference (and many are far from being good) pales in comparison with the availability of data on Crossref or arXiv. This is simply because Crossref is a community, community which has been updating its data for the last 22 years, where most entries are submitted manually. And not to mention that there is a professional software development team backing them up.

Another issue that might arise in natural language processing is the vagueness of the data that can be extracted. For example, some referencing styles shorten an author's name to an initial, while others do not. So say we find an M. Chenet and a Mattias Chenet. Are these two the same person? Does M stand for Mattis or for Marc? There is no way to tell. This issue only gets more complicated with the addition of et al where many authors are simply omitted from the reference. Arguably, these issues are sort of resolved in BibTeX: all authors' names should be written out in BibTeX and then the LaTeX engine only displays et al in the actual PDF. Unfortunately, many papers do not conform to this practice and put an et al in BibTex.

Nevertheless, if someone wanted to implement natural language processing here is an outline of how to do it.

1. Step number 1, a data set on which a classifier is to be trained is created. Luckily, this can be done using Crossref. Take a bunch of references, as many as needed, query them through Crossref and extract the metadata. It is important to capture all metadata that appears in the written-out references.

2. Having the training data, the next step is to create a good tokenizer. Parts of a reference are usually split by a period, while a comma separates a part into subparts. A simple tokenizer would then recognize that a single token is whatever lies between two periods. The tokenizer needs to take care of special cases, for example not to

split up an author's name when written out like "Colin B. Clement". This isn't hard, simply don't split when encountering a period preceded by a single capitalized letter. Certainly, other edge cases complicate this process and they would need to be taken care of.

3. Finally, we get to the usage of machine learning. The problem we are facing here is that of multiclass classification. Quite a few machine learning approaches are capable of solving this problem: neural networks, k-nearest neighbours, naive Bayes, decision trees, support vector machines, and many more. These approaches will probably work fine, but only if they are provided with the right features of the tokens (parts of the reference) they need to analyse and classify. Here are some ideas for the features. Every field has specific properties that should be used as

Feature name	Explanation
Paragraph-Length	Count the number of words in the token
Capital Percentage	The percentage of words that start with a capital letter
Number-count	The count of numbers in the token
Quote marks existence	Boolean value that indicates whether something is enclosed in quotes
Month abbreviation	Boolean value whether there is an abbreviation for a month (Apr, Jun, ...)

Table 1: Some ideas for the feature machine learning algorithms could use.

features so the ML algorithms work better.

5   Conclusion

In this project, we have built a reference extracting python script that finds all the references of a given arXiv paper, extracts their metadata and stores that metadata in a database. The script is as close to being implemented in a final website product as possible, given that we weren't working with any backend frameworks and the goal was the creation of a pure python script. The script shouldn't be hard to adapt into some backend framework.

The original goals of the project were not met in their entirety, but the goals of the project changed as a deeper understanding of the project and its purpose were gained. No natural language processing was added. Metadata extraction turned out to be a bigger problem than previously thought. It was realized that a reference tree can only be created for arXiv papers; to create a reference tree for all other papers, every single paper in the entire history of publishing would need to be processed. Some methods of metadata extraction are a bit too slow for immediate practical purposes and will need some more tuning and improvement to be deployable as a service.

However, the processing of the references works quite well. All arXiv IDs are captured by the employed regex, as are all DOIs (within the test set of 100 test papers). References with no identifier can be queried to Crossref's JSON REST API which can match references in as quick as half a second. There is also an option of querying Crossref's XML API, where we can submit multiple references at once for processing, instead of a one-by-one approach utilized by the REST API.

As with most projects, there is always room for improvement. Natural language processing can be implemented and could improve reference matching speed (no need for external resources). There are many hurdles to be overcome to achieve this and would require significant work. Possibilities of better performance with multi-threaded programming are there as well. Various other minor improvements could be made and they are all documented within the code. Certainly, the BibTex cleaning up function could be improved upon.

References

[1] Paul Ginsparg. *arXiv.org e-Print archive*. Cornell University. 1991. URL: `https://arxiv.org/` (visited on 07/14/2022).

[2] Florian Beutler and Morag Scrimgeour. *Benty-Fields*. 2015. URL: `https://benty-fields.com/` (visited on 07/14/2022).

[3] Wikipedia contributors. *Regular expression — Wikipedia, The Free Encyclopedia*. 2022. URL: `https://en.wikipedia.org/w/index.php?title=Regular_expression&oldid=1095256273` (visited on 07/14/2022).

[4] Mattia Chenet. *Identify and extract entities from bibliography references in a free text*. Sept. 2017. URL: `http://essay.utwente.nl/73817/`.

[5] Ed Pentz. *Crossref*. Publishers International Linking Association, Inc. 2000. URL: `https://crossref.org/` (visited on 07/14/2022).

[6] *SQLite*. 2000. URL: `https://sqlite.org/` (visited on 07/14/2022).

[7] Wikipedia contributors. *API — Wikipedia, The Free Encyclopedia*. 2022. URL: `https://en.wikipedia.org/w/index.php?title=API&oldid=1096919409` (visited on 07/15/2022).

[8] *Hypertext Transfer Protocol*. Developer.mozilla.org. 2022. URL: `https://developer.mozilla.org/en-US/docs/Web/HTTP` (visited on 07/15/2022).

[9] *Introducing JSON*. 2022. URL: `https://json.org/json-en.html` (visited on 07/15/2022).

[10] *Extensible Markup Language (XML) 1.0 (Fifth Edition)*. 2008. URL: `https://w3.org/TR/REC-xml/` (visited on 07/15/2022).

[11] *arXiv API User's Manual*. 2021. URL: `https://arxiv.org/help/api/user-manual` (visited on 07/15/2022).

[12] Ken Thompson. "Programming Techniques: Regular Expression Search Algorithm". In: *Commun. ACM* 11.6 (June 1968), pp. 419–422. ISSN: 0001-0782. DOI: `10.1145/363347.363387`. URL: `https://doi.org/10.1145/363347.363387`.

[13] International DOI Foundation. *The DOI System*. 2021. URL: `https://doi.org/` (visited on 07/16/2022).

[14] Andrew GIlmartin. *DOIs and matching regular expressions*. 2015. URL: `https://www.crossref.org/blog/dois-and-matching-regular-expressions/` (visited on 07/15/2022).

[15]   *Understanding the arXiv identifier.* 2019. URL: `https://arxiv.org/help/arxiv_identifier` (visited on 07/15/2022).

[16]   Colin B. Clement et al. *On the Use of ArXiv as a Dataset.* 2019. arXiv: `1905.00075` `[cs.IR]`.

[17]   Scott Chamberlain. *Habanero.* 2022. URL: `https://github.com/sckott/habanero` (visited on 07/18/2022).

A   Personal Statement

In this project, the main skills I work and improved on are technical, software development skills. The project involved learning about APIs and how to interact with them through HTTP requests implemented in python, data management in python, database management, implementation of regex, and basics of natural language processing. Furthermore, I learned a great deal about how referencing works within the scientific community, the challenges that Crossref tries to solve and just how challenging those challenges actually are, and the purpose of DOIs as a robust reference (pointer) to a reference. Soft skills I improved on are report writing, especially software engineering reports and code writing – I improved on good coding practices. This is especially important as this code is going to be used by someone else after me, and the code has to be in an easily understandable state. Additionally, I learned how to work in a team of researchers and share my progress and ideas with other team members.

B   Lay Summary

ArXiv is an online database of scientific papers. Every scientific paper references several other scientific papers. Our goal was to create a python script that would detect these references and extract data about them - linking the references to the paper they were found in. References are identified primarily either through arXiv IDs or DOIs (digital object identifiers). References that have neither of those things are sent to a special online service called Crossref, which matches the sent reference to a scientific paper in its database. Data that is gathered about these references is saved in a database, so it can be retrieved quickly and on-demand. This reference data is going to be displayed next to the corresponding scientific paper on a website called Benty-Fields - a website that gathers researchers into book clubs where they can recommend good papers to one another.

C   Acknowledgements

A big thank you to my supervisor for this project, Florian Beutler. He has been a great mentor and a pleasure to work with.