
web-poet Documentation

Release 0.4.0

Zyte Group Ltd

Jul 26, 2022

GETTING STARTED

1	web-poet on a surface	3
2	web-poet from the ground up	7
3	Overrides	19
4	Additional Requests	29
5	Retries	47
6	Using page params	49
7	API Reference	53
8	Contributing	65
9	Changelog	67
10	License	69
	Python Module Index	71
	Index	73

Warning: web-poet is in early stages of development; backwards incompatible changes are possible.

web-poet implements Page Object pattern for web scraping. It defines a standard for writing web data extraction code, which allows the code to be portable & reusable.

The main idea is to separate the extraction logic from all other concerns. web-poet Page Objects *don't do I/O*, and they're not dependent on any particular framework like *Scrapy*.

This allows the code written using web-poet to be testable and reusable. For example, one can write a web-poet Page Object in an IPython notebook, plug it into a Scrapy spider, write tests for them using unittest or pytest, and then reuse in a simple script which uses requests library.

To install it, run `pip install web-poet`. It requires Python 3.7+. *License* is BSD 3-clause.

If you want to quickly learn how to write web-poet Page Objects, see *web-poet on a surface*. To understand better all the web-poet concepts and the motivation behind web-poet, start with *web-poet from the ground up*.

WEB-POET ON A SURFACE

In this tutorial, we'll assume that web-poet is already installed on your system.

We are going to scrape books.toscrape.com, a website that lists books from famous authors.

1.1 Creating a Page Object

Let's create a Python file where we'll put our first Page Object implementation. This Page Object will be responsible for extracting book links from the book list page on books.toscrape.com.

```
from web_poet.pages import ItemWebPage

class BookLinksPage(ItemWebPage):

    @property
    def links(self):
        return self.css('.image_container a::attr(href)').getall()

    def to_item(self) -> dict:
        return {
            'links': self.links,
        }
```

1.2 Downloading Response

The BookLinksPage Page Object requires a *HttpResponse* with the book list page content in order to extract the information we need. First, let's download the page using requests library.

```
import requests

response = requests.get('http://books.toscrape.com')
```

1.3 Creating Page Input

Now we need to create and populate a *HttpResponse* instance.

```
from web_poet.page_inputs import HttpResponse

response_data = HttpResponse(response.url,
                             body=response.content,
                             headers=response.headers)
page = BookLinksPage(response_data)

print(page.to_item())
```

1.4 Final Result

Our simple Python script might look like this:

```
import requests

from web_poet.pages import ItemWebPage
from web_poet.page_inputs import HttpResponse

class BookLinksPage(ItemWebPage):

    @property
    def links(self):
        return self.css('.image_container a::attr(href)').getall()

    def to_item(self) -> dict:
        return {
            'links': self.links,
        }

response = requests.get('http://books.toscrape.com')
response_data = HttpResponse(response.url,
                             body=response.content,
                             headers=response.headers)

page = BookLinksPage(response_data)

print(page.to_item())
```

And it should output data similar to this:

```
{
  'links': [
    'catalogue/a-light-in-the-attic_1000/index.html',
    'catalogue/tipping-the-velvet_999/index.html',
```

(continues on next page)

(continued from previous page)

```
'catalogue/soumission_998/index.html',
'catalogue/sharp-objects_997/index.html',
'catalogue/sapiens-a-brief-history-of-humankind_996/index.html',
'catalogue/the-requiem-red_995/index.html',
'catalogue/the-dirty-little-secrets-of-getting-your-dream-job_994/index.html',
'catalogue/the-coming-woman-a-novel-based-on-the-life-of-the-infamous-feminist-
↪victoria-woodhull_993/index.html',
'catalogue/the-boys-in-the-boat-nine-americans-and-their-epic-quest-for-gold-at-
↪the-1936-berlin-olympics_992/index.html',
'catalogue/the-black-maria_991/index.html',
'catalogue/starving-hearts-triangular-trade-trilogy-1_990/index.html',
'catalogue/shakespeares-sonnets_989/index.html',
'catalogue/set-me-free_988/index.html',
'catalogue/scott-pilgrims-precious-little-life-scott-pilgrim-1_987/index.html',
'catalogue/rip-it-up-and-start-again_986/index.html',
'catalogue/our-band-could-be-your-life-scenes-from-the-american-indie-
↪underground-1981-1991_985/index.html',
'catalogue/olio_984/index.html',
'catalogue/mesaerion-the-best-science-fiction-stories-1800-1849_983/index.html',
'catalogue/libertarianism-for-beginners_982/index.html',
'catalogue/its-only-the-himalayas_981/index.html',
]
}
```

1.5 Next Steps

As you can see, it's possible to use web-poet with built-in libraries such as `requests`, but it's also possible to use `Scrapy` with the help of `scrapy-poet`.

If you want to understand the idea behind web-poet better, check the *web-poet from the ground up* tutorial.

WEB-POET FROM THE GROUND UP

This tutorial explains the motivation behind web-poet, and its main concepts. You would learn

- what are the issues which web-poet addresses, why does the library exist;
- how web-poet addresses these issues;
- what is a Page Object, and what is a page object input;
- how are libraries like `scrapy-poet` implemented, and what are they for.

2.1 Reusable web scraping code

Forget about web-poet for a minute. Let's say you're writing code to scrape a book web page from `books.toscrape.com`:

```
import requests
import parsel

def extract_book(url):
    """ Extract book information from a book page on
    http://books.toscrape.com website, e.g. from
    http://books.toscrape.com/catalogue/a-light-in-the-attic_1000/index.html
    """
    resp = requests.get(url)
    sel = parsel.Selector(resp)
    return {
        'url': resp.url,
        'title': sel.css('h1').get(),
        'description': sel.css('#product_description+ p').get().strip(),
        # ...
    }

item = extract_book("http://books.toscrape.com/catalogue/a-light-in-the-attic_1000/index.
↪html")
```

This is sweet & easy, but you realize this code is tightly coupled, and hardly reusable or testable. So you refactor it, to separate downloading from the extraction:

```
import requests
import parsel
```

(continues on next page)

(continued from previous page)

```

def extract_book(response):
    """ Extract book information from a requests.Response obtained
    from a book page on http://books.toscrape.com website, e.g. from
    http://books.toscrape.com/catalogue/a-light-in-the-attic_1000/index.html
    """
    sel = parsel.Selector(response.text)
    return {
        'url': response.url,
        'title': sel.css('h1').get(),
        'description': sel.css('#product_description+ p').get().strip(),
        # ...
    }

def download_and_extract_book(url):
    resp = requests.get(url)
    return extract_book(resp)

item = download_and_extract_book("http://books.toscrape.com/catalogue/a-light-in-the-
↪attic_1000/index.html")

```

Much cleaner! It is slightly more code, but `extract_book` is more reusable & testable now.

Then your project evolves, and you realize that you'd like to download web pages using `aiohttp`. It means that `extract_book` now can't receive `requests.Response`, it needs to work with `aiohttp.Response`. To complicate things, you want to keep `requests` support. There are few options, such as:

- make `extract_book` receive `Selector` instance
- make `extract_book` receive url and unicode response body
- make `extract_book` receive a `Response` object which is compatible for both libraries.

No problem, let's refactor it further. You may end up with something like that:

```

import aiohttp
import requests
import parsel

# === Extraction code
def extract_book(url, text):
    """ Extract book information from a book page
    on http://books.toscrape.com website, e.g. from
    http://books.toscrape.com/catalogue/a-light-in-the-attic_1000/index.html
    """
    sel = parsel.Selector(text)
    return {
        'url': url,
        'title': sel.css('h1').get(),
        'description': sel.css('#product_description+ p').get().strip(),
        # ...
    }

# === Framework-specific I/O code
def download_sync(url):

```

(continues on next page)

(continued from previous page)

```

resp = requests.get(url)
return {'url': resp.url, 'text': resp.text}

async def download_async(url):
    async with aiohttp.ClientSession() as session:
        async with session.get(url) as response:
            text = await response.text()
            return {'url': url, 'text': text}

# === Usage example
# the way to get resp_data depends on an HTTP client
resp_data = download_sync("http://books.toscrape.com/catalogue/a-light-in-the-attic_1000/
↪index.html")

# but after we got resp_data, usage is the same
item = extract_book(url=resp_data['url'], text=resp_data['text'])

```

`extract_book` function now has all the desired properties: it is easily testable and reusable, and it works with any method of downloading data.

2.2 The same, but using web-poet

`web-poet` asks you to organize code in a very similar way. Let's convert `extract_book` function to a Page Object, by defining `BookPage` class:

```

import aiohttp
import requests
from web_poet import WebPage, HttpResponse

# === Extraction code
class BookPage(WebPage):
    """
    A book page on http://books.toscrape.com website, e.g.
    http://books.toscrape.com/catalogue/a-light-in-the-attic_1000/index.html
    """
    def extract_book(self):
        return {
            'url': self.url,
            'title': self.css('h1').get(),
            'description': self.css('#product_description+ p').get().strip(),
            # ...
        }

# === Framework-specific I/O code
def download_sync(url):
    resp = requests.get(url)
    return HttpResponse(url=resp.url, body=resp.content, headers=resp.headers)

async def download_async(url):

```

(continues on next page)

(continued from previous page)

```

async with aiohttp.ClientSession() as session:
    async with session.get(url) as response:
        body = await response.content.read()
        headers = response.headers

    return HttpResponse(url=resp.url, body=body, headers=headers)

# === Usage example

# the way to get response depends on an HTTP client
response = download_sync("http://books.toscrape.com/catalogue/a-light-in-the-attic_1000/
↳index.html")

# but after we got the response, the usage is the same
book_page = BookPage(response=response)
item = book_page.extract_book()

```

Differences from a previous example:

- instead of dicts with “url” and “text” fields, *HttpResponse* instances are used. *HttpResponse* is a structure defined by web-poet acting as a generic data container for HTTP Responses. (check out the API reference of *HttpResponse* for more info about the fields it holds)

Note how headers are passed when creating *HttpResponse* instance. This is needed to decode body (which is bytes) to unicode properly, using the web browser rules. It involves checking Content-Encoding header, meta tags in HTML, BOM markers in the body, etc.

- instead of `extract_book` function we got *BookPage* class, which receives response data in its `__init__` method - see how it is created: `BookPage(response=response)`.
- *BookPage* inherits from *WebPage* base class. This base class is not doing much: it
 - defines `__init__` method which receives *HttpResponse*, and
 - provides shortcut methods like `css()`, which work by creating `parsel.selector.Selector` behind the scenes (so that you don’t need to create a selector in the `extract_book` method).

There are pros and cons for using classes vs functions for writing such extraction code, but the distinction is not that important; web-poet uses classes at the moment.

2.3 to_item() method

It is common to have Page Objects for a web page where a single main data record needs to be extracted (e.g. book information in our example). web-poet standardizes this, by asking to name a method implementing the extraction `to_item`. It also provides the *ItemWebPage* base class and the *ItemPage* mixin, which ensure the `to_item` method is implemented. Let’s change the code to follow this standard:

```

import requests
from web_poet import ItemWebPage, HttpResponse

# === Extraction code
class BookPage(ItemWebPage):
    """

```

(continues on next page)

(continued from previous page)

```

A book page on http://books.toscrape.com website, e.g.
http://books.toscrape.com/catalogue/a-light-in-the-attic_1000/index.html
"""
def to_item(self):
    return {
        'url': self.url,
        'title': self.css('h1').get(),
        'description': self.css('#product_description+ p').get().strip(),
        # ...
    }

# ... get resp_data somehow
book_page = BookPage(response=response)
item = book_page.to_item()

```

As the method name is now standardized, the code which creates a Page Object instance can now work for other Page Objects like that. For example, you can have ToscrapeBookPage and BamazonBookPage classes, and

```

def get_item(page_cls: ItemWebPage, response: HttpResponse) -> dict:
    page = page_cls(response=response)
    return page.to_item()

```

would work for both.

But wait. Before the example was converted to web-poet, we were getting it for free:

```

def get_item(extract_func, response: HttpResponse) -> dict:
    return extract_func(url=response.url, text=response.text)

```

No need to agree on to_item name and have a base class to check that the method is implemented. Why bother with classes then?

2.4 Classes for web scraping code

A matter of preference. Functions are great, too. Classes sometimes can make it a easier to organize web scraping code. For example, we can extract logic for different attributes into properties:

```

class BookPage(ItemWebPage):
    """
    A book page on http://books.toscrape.com website, e.g.
    http://books.toscrape.com/catalogue/a-light-in-the-attic_1000/index.html
    """

    @property
    def title(self):
        return self.css('h1').get()

    @property
    def description(self):
        return self.css('#product_description+ p').get().strip()

```

(continues on next page)

(continued from previous page)

```
def to_item(self):
    return {
        'url': self.url,
        'title': self.title,
        'description': self.description,
        # ...
    }
```

You may write some base class to make it nicer - e.g. helper descriptors to define properties from CSS selectors, and a default `to_item` implementation (so, no need to define `to_item`). This is currently not implemented in `web-poet`, but nothing prevents us from having a DSL like this:

```
class BookPage(ItemWebPage):
    title = Css('h1')
    description = Css('#product_description+ p') | Strip()
    url = TakeUrl()
```

Another, and probably a more important reason to consider classes for the extraction code, is that sometimes there is no a single “main” method, but you still want to group the related code. For example, you may define a “Pagination” page object:

```
class Pagination(WebPage):

    def page_urls(self):
        # ...

    def prev_url(self):
        # ...

    def next_url(self):
        # ...
```

or a Listing page on a web site, where you need to get URLs to individual pages and pagination URLs:

```
class BookListPage(ProductListingPage):
    def item_urls(self):
        return self.css(".product a::attr(href)").getall()

    def page_urls(self):
        return self.css(".paginator a::attr(href)").getall()
```

2.5 Web Scraping Frameworks

Let’s recall the example we started with:

```
import requests
import parsel

def extract_book(url):
    """ Extract book information from a book page on
```

(continues on next page)

(continued from previous page)

```

http://books.toscrape.com website, e.g. from
http://books.toscrape.com/catalogue/a-light-in-the-attic_1000/index.html
"""

resp = requests.get(url)
sel = parsel.Selector(resp)
return {
    'url': resp.url,
    'title': sel.css('h1').get(),
    'description': sel.css('#product_description+ p').get().strip(),
    # ...
}

item = extract_book("http://books.toscrape.com/catalogue/a-light-in-the-attic_1000/index.
↪html")

```

And this is what we ended up with:

```

import requests
from web_poet import ItemWebPage, HttpResponse

# === Extraction code
class BookPage(ItemWebPage):
    """ A book page on http://books.toscrape.com website, e.g.
    http://books.toscrape.com/catalogue/a-light-in-the-attic_1000/index.html
    """

    def to_item(self):
        return {
            'url': self.url,
            'title': self.css('h1').get(),
            'description': self.css('#product_description+ p').get().strip(),
            # ...
        }

# === Framework-specific I/O code
def download_sync(url):
    resp = requests.get(url)
    return HttpResponse(url=resp.url, body=resp.content, headers=resp.headers)

async def download_async(url):
    async with aiohttp.ClientSession() as session:
        async with session.get(url) as response:
            body = await response.content.read()
            headers = response.headers

    return HttpResponse(url=resp.url, body=body, headers=headers)

# === Usage example
def get_item(page_cls: ItemWebPage, resp_data: HttpResponse) -> dict:
    page = page_cls(response=resp_data)
    return page.to_item()

```

(continues on next page)

(continued from previous page)

```
# the way to get resp_data depends on an HTTP client
response = download_sync("http://books.toscrape.com/catalogue/a-light-in-the-attic_1000/
↳index.html")
item = get_item(BookPage, resp_data=response)
```

We created a monster!!! The examples in this tutorial are becoming longer and longer, harder and harder to understand. What's going on?

To understand better why this is happening, let's check the web-poet example in more detail. There are 3 main sections:

1. Extraction code
2. Lower-level I/O code
3. "Usage example" - it connects (1) and (2) to get the extracted data.

Extraction code needs to be written for every new web site. But the I/O code and the "Usage example" can, and should be written only once!

In other words, we've been creating a web scraping framework here, and that's where most of the complexity is coming from.

In a real world, a developer who needs to extract data from a web page would only need to write the "extraction" part:

```
from web_poet import ItemWebPage

class BookPage(ItemWebPage):
    """ A book page on http://books.toscrape.com website, e.g.
    http://books.toscrape.com/catalogue/a-light-in-the-attic_1000/index.html
    """
    def to_item(self):
        return {
            'url': self.url,
            'title': self.css('h1').get(),
            'description': self.css('#product_description+ p').get().strip(),
            # ...
        }
```

Then point the framework to the BookPage class, tell which web page to process, and that's it:

```
item = some_framework.extract(url, BookPage)
```

The role of web-poet is to define a standard on how to write the extraction logic, and allow it to be reused in different frameworks. web-poet Page Objects should be flexible enough to be used with:

- synchronous or async frameworks, callback-based and async `def / await` based,
- single node and distributed systems,
- different underlying HTTP implementations - or without HTTP support at all, etc.

2.6 Page Objects

In this document “Page Objects” were casually mentioned a few times, but what are they?

Note: This term comes from a Page Object design pattern; see a description on Martin Fowler’s website: <https://martinfowler.com/bliki/PageObject.html>. web-poet page objects are inspired by Martin Fowler’s page object, but they are not the same.

Essentially, the idea is to create an object which represents a web page (or a part of web page - recall the `Pagination` example), and allows to extract data from there. Page Object must:

1. Define all the inputs needed in its `__init__` method. Usually these inputs are then stored as attributes.
2. Provide methods or properties to extract structured information, using the data saved in `__init__`.
3. Inherit from *Injectable*; this inheritance is used as a marker.

For example, a very basic Page Object could look like this:

```
from parsel import Selector
from web_poet.pages import Injectable
from web_poet.page_inputs import HttpResponse

class BookPage(Injectable):
    def __init__(self, response: HttpResponse):
        self.response = response

    def to_item(self) -> dict:
        return {
            'url': str(self.response.url),
            'title': self.response.css("h1::text").get()
        }
```

There is no *need* to use other base classes and mixins defined by web-poet (*WebPage*, *ResponseShortcutsMixin*, *ItemPage*, *ItemWebPage*, etc.), but it can be a good idea to familiarize yourself with them, as they are taking some of the boilerplate out.

2.7 Page Object Inputs

Here we got to the last, and probably the most complicated and important part of web-poet. So far we’ve been passing *HttpResponse* to the page objects. But is it enough?

If that’d be enough, there wouldn’t be web-poet. We would say “please write `def extract(url, html): ...` functions”, and call it a day.

In practice you may need to use other information to extract data from a web page, not only *HttpResponse*. For example, you may want to

- render a web page in a headless browser like *Splash*, and use HTML after the rendering (snapshot of a DOM tree);
- query third-party API like *AutoExtract*, to extract most of the data automatically - a Page Object may just return the result as-is, or enrich / post-process it;

- take some state in account, passed e.g. from the crawling code.
- use a combination of inputs: e.g. you may need HTML after Headless Chrome rendering + crawling state.

The information you need can depend on a web site. For example, Splash can be required for extracting book information from Bamazon, while for <http://books.toscrape.com> you may need HTTP response body and some crawl state (not really, but let's imagine it is needed). You may define page objects for this task:

```
class BamazonBookPage(Injectable):
    def __init__(self, response: SplashResponseData):
        self.response = response

    def to_item(self):
        # ...

class ToScrapeBookPage(Injectable):
    def __init__(self, response: HttpResponse, crawl_state: dict):
        self.response = response
        self.crawl_state = crawl_state

    def to_item(self):
        # ...
```

Then, we would like to use these page objects in some web scraping framework, like we did before:

```
item1 = some_framework.extract(bamazon_url, BamazonBookPage)
item2 = some_framework.extract(toscrape_url, ToScrapeBookPage)
```

To be able to implement the imaginary `some_framework.extract` method, `some_framework` must

1. Figure out somehow which inputs the Page Objects need.
2. Create these inputs. If needed, make a Splash request, make a direct HTTP request, get a dictionary with the crawl state from somewhere. These actions can be costly; framework should avoid doing unnecessary work here.
3. Pass the obtained data as keyword arguments to `__init__` method.

(2) and (3) are straightforward, once the framework knows that “To create `BamazonBookPage`, I need to pass output of Splash as a `response` keyword argument”, i.e. once (1) is done.

web-poet uses **type annotations** of `__init__` arguments to declare Page Object dependencies. So, type annotations in the examples like the following were not just a nice-thing-to-have:

```
class BookPage(Injectable):
    def __init__(self, response: HttpResponse):
        self.response = response
```

By annotating `__init__` arguments we were actually telling web-poet (or, more precisely, a framework which uses web-poet):

To create a `BookPage` instance, please obtain `HttpResponse` instance somehow, and pass it as a `response` keyword argument. That's all you need to create a `BookPage` instance.

Note: If it sounds like Dependency Injection, you're right.

If something other than `HttpResponse` needs to be passed, a different type annotation should be used:

```

class BamazonBookPage(Injectable):
    def __init__(self, response: SplashResponseData):
        self.response = response

class ToScrapeBookPage(Injectable):
    def __init__(self, response: HttpResponse, crawl_state: CrawlState):
        self.response = response
        self.crawl_state = crawl_state

```

For each possible input a separate class needs to be defined, even if the data has the same format. For example, both `HttpResponse` and `SplashResponseData` may have the same `url` and `text` properties, but they can't be the same class, because they need to work as "markers" - tell frameworks if the html should be taken from HTTP response body or from Splash DOM snapshot.

`CrawlState` in the example above can be defined as a class with some specific properties, or maybe even as a class `CrawlState(dict): pass` - an important thing is that it is a unique type, and that we agree on what should be put into arguments annotated as `CrawlState`.

Pro tip: defining classes like

```

class ToScrapeBookPage(Injectable):
    def __init__(self, response: HttpResponse, crawl_state: CrawlState):
        self.response = response
        self.crawl_state = crawl_state

```

can get tedious; Python's `dataclasses` (or `attrs`, if that's your preference) make it nicer:

```

from dataclasses import dataclass

@dataclass
class ToScrapeBookPage(Injectable):
    response: HttpResponse
    crawl_state: CrawlState

```

2.8 web-poet role

How do you actually inspect the `__init__` method signature - e.g. if you're working on supporting web-poet page objects in some framework? web-poet itself doesn't provide any helpers for doing this.

Use `andi` library. For example, `scrapy-poet` uses `andi`. In addition to signature inspection, it also handles `typing.Optional` and `typing.Union`, and allows to create a build plan for dependency trees, indirect dependencies: that's allowed to annotate an argument as another `Injectable` subclass.

web-poet is not using `andi` on its own; web-poet's role is to standardize things + provide some helpers to write the extraction code easier:

1. Standardize a list of possible inputs for the page objects. This helps with reusability of extraction code across different environments. For example, if you want to support extraction from raw HTTP response bodies, you need to figure out how to populate `HttpResponse` in the given environment, and that's all.

Users are free to define their own inputs (input types), but they may be less portable across environments - which can be fine.

2. Define an interface for the Page Object itself. This allows to have a code which can instantiate and use a Page Object without knowing about its implementation upfront. web-poet requires you to use a base class

(*Injectable*), and defines the semantics of `to_item()` method.

Then, framework's role is to:

1. Figure out which inputs a Page Object needs, likely using `andi` library.
2. Create all the necessary inputs. For example, creating `HttpResponse` instance may involve making an HTTP request; creating `CrawlState` (from the previous examples) may involve getting some data from the shared storage, or from an in-memory data structure.
3. Create a Page Object instance, passing it the inputs it needs.
4. Depending on a task, either return a newly created Page Object instance to the user, or call some predefined method (a common case is `to_item`).

For example, `web-poet` + Scrapy integration package (`scrapy-poet`) may inspect a `WebPage` subclass you defined, figure out it needs `HttpResponse` and nothing else, fetch scrapy's `TextResponse`, create `HttpResponse` instance from it, create your Page Object instance, and pass it to a spider callback.

Finally, the Developer's role is to:

1. Write a Page Object class, likely website-specific, following `web-poet` standards. The extraction code should define the inputs it needs (such as "body of HTTP response", "Chrome DOM tree snapshot", "crawl state"); it shouldn't fetch these inputs itself.
2. Pass the Page Object class to a framework, in a way defined by the framework.
3. Receive a Page Object *instance* from the framework; call its extraction methods (e.g. `to_item`). Depending on the framework and on the task, the framework may be calling `to_item` (or other methods) automatically; in this case user code would be getting the extracted data, not a Page Object instance.

2.9 Summary

First, congratulations for making it through this document!

A take-away from this tutorial:

1. `web-poet` does very little on its own. Almost nothing, really. An important thing about `web-poet` is that it defines a standard for writing web scraping code.

All these Page Objects are just Python classes, which receive some static data in `__init__` methods, and maybe provide some methods to extract the data.
2. `web-poet` prescribes certain things and limits what you can do, but not too much, and as a return you're getting better testability and reusability of your code.
3. Hopefully, now you understand how to write a web scraping framework which uses `web-poet`.
4. Basic `web-poet` usage looks similar to how one could have had refactored the extraction code anyways.

OVERRIDES

Overrides contains mapping rules to associate which URLs a particular Page Object would be used. The URL matching rules is handled by another library called `url-matcher`.

Using such rules establishes the core concept of Overrides wherein a developer could declare that a specific Page Object must be used (*instead of another*) for a given set of URL patterns.

This enables **web-poet** to be used effectively by other frameworks like `scrapy-poet`.

3.1 Example Use Case

Let's explore an example use case for the Overrides concept.

Suppose we're using Page Objects for our broadcrawl project which explores eCommerce websites to discover product pages. It wouldn't be entirely possible for us to create parsers for all websites since we don't know which sites we're going to crawl beforehand.

However, we could at least create a generic Page Object to support parsing of some fields in well-known locations of product information like `<title>`. This enables our broadcrawler to at least parse some useful information. Let's call such Page Object to be `GenericProductPage`.

Assuming that one of our project requirements is to fully support parsing of the *top 3 eCommerce websites*, then we'd need to create a Page Object for each one to parse more specific fields.

Here's where the Overrides concept comes in:

1. The `GenericProductPage` is used to parse all eCommerce product pages *by default*.
2. Whenever one of our declared URL rules matches with a given page URL, then the Page Object associated with that rule *overrides (or replaces)* the default `GenericProductPage`.

This enables us to conveniently declare which Page Object would be used for a given webpage (*based on a page's URL pattern*).

Let's see this in action by declaring the Overrides in the Page Objects below.

3.2 Creating Overrides

Let's take a look at how the following code is structured:

```
from web_poet import handle_urls, ItemWebPage

class GenericProductPage(ItemWebPage):
    def to_item(self):
        return {"product-title": self.css("title::text").get()}

@handle_urls("example.com", overrides=GenericProductPage)
class ExampleProductPage(ItemWebPage):
    def to_item(self):
        ... # more specific parsing

@handle_urls("anotherexample.com", overrides=GenericProductPage, exclude="/digital-goods/
→")
class AnotherExampleProductPage(ItemWebPage):
    def to_item(self):
        ... # more specific parsing

@handle_urls(["dualexample.com/shop/?product=*", "dualexample.net/store/?pid=*"],
→overrides=GenericProductPage)
class DualExampleProductPage(ItemWebPage):
    def to_item(self):
        ... # more specific parsing
```

The code above declares that:

- For sites that match the `example.com` pattern, `ExampleProductPage` would be used instead of `GenericProductPage`.
- The same is true for `DualExampleProductPage` where it is used instead of `GenericProductPage` for two URL patterns which works as something like:
 - (match) `https://www.dualexample.com/shop/electronics/?product=123`
 - (match) `https://www.dualexample.com/shop/books/paperback/?product=849`
 - (NO match) `https://www.dualexample.com/on-sale/books/?product=923`
 - (match) `https://www.dualexample.net/store/kitchen/?pid=776`
 - (match) `https://www.dualexample.net/store/?pid=892`
 - (NO match) `https://www.dualexample.net/new-offers/fitness/?pid=892`
- On the other hand, `AnotherExampleProductPage` is only used instead of `GenericProductPage` when we're handling pages from `anotherexample.com` that doesn't contain `/digital-goods/` in its URL path.

Tip: The URL patterns declared in the `@handle_urls` annotation can still be further customized. You can read some of the specific parameters in the *API section* of `web_poet.handle_urls()`.

3.3 Retrieving all available Overrides

The `get_overrides()` method from the `web_poet.default_registry` allows retrieval of all `OverrideRule` in the given registry. Following from our example above, using it would be:

```
from web_poet import default_registry

# Retrieves all OverrideRules that were registered in the registry
rules = default_registry.get_overrides()

print(len(rules)) # 3
print(rules[0]) # OverrideRule(for_patterns=Patterns(include=['example.com'],
↳ exclude=[], priority=500), use=<class 'my_project.page_objects.ExampleProductPage'>,
↳ instead_of=<class 'my_project.page_objects.GenericProductPage'>, meta={})
```

Remember that using `@handle_urls` to annotate the Page Objects would result in the `OverrideRule` to be written into `web_poet.default_registry`.

Warning: `get_overrides()` relies on the fact that all essential packages/modules which contains the `web_poet.handle_urls()` annotations are properly loaded (*i.e imported*).

Thus, for cases like importing and using Page Objects from other external packages, the `@handle_urls` annotations from these external sources must be read and processed properly. This ensures that the external Page Objects have all of their `OverrideRule` present.

This can be done via the function named `consume_modules()`. Here's an example:

```
from web_poet import default_registry, consume_modules

consume_modules("external_package_A.po", "another_ext_package.lib")
rules = default_registry.get_overrides()
```

The next section explores this caveat further.

3.4 Using Overrides from External Packages

Developers have the option to import existing Page Objects alongside the `OverrideRule` attached to them. This section aims to showcase different scenarios that come up when using multiple Page Object Projects.

3.4.1 Using all available OverrideRules from multiple Page Object Projects

Let's suppose we have the following use case before us:

- An **external** Python package named `ecommerce_page_objects` is available which contains Page Objects for common websites.
- Another similar **external** package named `gadget_sites_page_objects` is available for even more specific websites.
- Your project's objective is to handle as much eCommerce websites as you can.
 - Thus, you'd want to use the already available packages above and perhaps improve on them or create new Page Objects for new websites.

Remember that all of the *OverrideRule* are declared by annotating Page Objects using the *web_poet.handle_urls()* via *@handle_urls*. Thus, they can easily be accessed using the *get_overrides()* of *web_poet.default_registry*.

This can be done something like:

```
from web_poet import default_registry, consume_modules

# Remember that this wouldn't retrieve any rules at all since the
# annotations are NOT properly imported.
rules = default_registry.get_overrides()
print(rules) # []

# Instead, you need to run the following so that all of the Page
# Objects in the external packages are recursively imported.
consume_modules("ecommerce_page_objects", "gadget_sites_page_objects")
rules = default_registry.get_overrides()

# The collected rules would then be as follows:
print(rules)
# 1. OverrideRule(for_patterns=Patterns(include=['site_1.com'], exclude=[], priority=500),
# ↪ use=<class 'ecommerce_page_objects.site_1.EcomSite1'>, instead_of=<class 'ecommerce_
# ↪ page_objects.EcomGenericPage'>, meta={})
# 2. OverrideRule(for_patterns=Patterns(include=['site_2.com'], exclude=[], priority=500),
# ↪ use=<class 'ecommerce_page_objects.site_2.EcomSite2'>, instead_of=<class 'ecommerce_
# ↪ page_objects.EcomGenericPage'>, meta={})
# 3. OverrideRule(for_patterns=Patterns(include=['site_2.com'], exclude=[], priority=500),
# ↪ use=<class 'gadget_sites_page_objects.site_2.GadgetSite2'>, instead_of=<class 'gadget_
# ↪ sites_page_objects.GadgetGenericPage'>, meta={})
# 4. OverrideRule(for_patterns=Patterns(include=['site_3.com'], exclude=[], priority=500),
# ↪ use=<class 'gadget_sites_page_objects.site_3.GadgetSite3'>, instead_of=<class 'gadget_
# ↪ sites_page_objects.GadgetGenericPage'>, meta={})
```

Note: Once *consume_modules()* is called, then all external Page Objects are recursively imported and available for the entire runtime duration. Calling *consume_modules()* again makes no difference unless a new set of modules are provided.

3.4.2 Using only a subset of the available *OverrideRules*

Suppose that the use case from the previous section has changed wherein a subset of *OverrideRule* would be used. This could be achieved by using the *search_overrides()* method which allows for convenient selection of a subset of rules from a given registry.

Here's an example of how you could manually select the rules using the *search_overrides()* method instead:

```
from web_poet import default_registry, consume_modules
import ecommerce_page_objects, gadget_sites_page_objects

consume_modules("ecommerce_page_objects", "gadget_sites_page_objects")

ecom_rules = default_registry.search_overrides(instead_of=ecommerce_page_objects.
↪ EcomGenericPage)
```

(continues on next page)

(continued from previous page)

```

print(ecom_rules)
# OverrideRule(for_patterns=Patterns(include=['site_1.com'], exclude=[], priority=500),
↳ use=<class 'ecommerce_page_objects.site_1.EcomSite1'>, instead_of=<class 'ecommerce_page_
↳ objects.EcomGenericPage'>, meta={})
# OverrideRule(for_patterns=Patterns(include=['site_2.com'], exclude=[], priority=500),
↳ use=<class 'ecommerce_page_objects.site_2.EcomSite2'>, instead_of=<class 'ecommerce_page_
↳ objects.EcomGenericPage'>, meta={})

gadget_rules = default_registry.search_overrides(use=gadget_sites_page_objects.site_3.
↳ GadgetSite3)
print(gadget_rules)
# OverrideRule(for_patterns=Patterns(include=['site_3.com'], exclude=[], priority=500),
↳ use=<class 'gadget_sites_page_objects.site_3.GadgetSite3'>, instead_of=<class 'gadget_
↳ sites_page_objects.GadgetGenericPage'>, meta={})

rules = ecom_rules + gadget_rules
print(rules)
# OverrideRule(for_patterns=Patterns(include=['site_1.com'], exclude=[], priority=500),
↳ use=<class 'ecommerce_page_objects.site_1.EcomSite1'>, instead_of=<class 'ecommerce_page_
↳ objects.EcomGenericPage'>, meta={})
# OverrideRule(for_patterns=Patterns(include=['site_2.com'], exclude=[], priority=500),
↳ use=<class 'ecommerce_page_objects.site_2.EcomSite2'>, instead_of=<class 'ecommerce_page_
↳ objects.EcomGenericPage'>, meta={})
# OverrideRule(for_patterns=Patterns(include=['site_3.com'], exclude=[], priority=500),
↳ use=<class 'gadget_sites_page_objects.site_3.GadgetSite3'>, instead_of=<class 'gadget_
↳ sites_page_objects.GadgetGenericPage'>, meta={})

```

As you can see, using the `search_overrides()` method allows you to conveniently select for `OverrideRule` which conform to a specific criteria. This allows you to conveniently drill down to which `OverrideRule` you're interested in using.

After gathering all the pre-selected rules, we can then store it in a new instance of `PageObjectRegistry` in order to separate it from the `default_registry` which contains all of the rules. We can use the `from_override_rules()` for this:

```

from web_poet import PageObjectRegistry

my_new_registry = PageObjectRegistry.from_override_rules(rules)

```

3.4.3 Improving on external Page Objects

There would be cases wherein you're using Page Objects with `OverrideRule` from external packages only to find out that a few of them lacks some of the fields or features that you need.

Let's suppose that we wanted to use *all* of the `OverrideRule` similar to this section: *Using all available OverrideRules from multiple Page Object Projects*. However, the `EcomSite1` Page Object needs to properly handle some edge cases where some fields are not being extracted properly. One way to fix this is to subclass the said Page Object and improve its `to_item()` method, or even creating a new class entirely. For simplicity, let's have the first approach as an example:

```

from web_poet import default_registry, consume_modules, handle_urls
import ecommerce_page_objects, gadget_sites_page_objects

```

(continues on next page)

(continued from previous page)

```

consume_modules("ecommerce_page_objects", "gadget_sites_page_objects")
rules = default_registry.get_overrides()

# The collected rules would then be as follows:
print(rules)
# 1. OverrideRule(for_patterns=Patterns(include=['site_1.com'], exclude=[], priority=500),
↳ use=<class 'ecommerce_page_objects.site_1.EcomSite1'>, instead_of=<class 'ecommerce_
↳ page_objects.EcomGenericPage'>, meta={})
# 2. OverrideRule(for_patterns=Patterns(include=['site_2.com'], exclude=[], priority=500),
↳ use=<class 'ecommerce_page_objects.site_2.EcomSite2'>, instead_of=<class 'ecommerce_
↳ page_objects.EcomGenericPage'>, meta={})
# 3. OverrideRule(for_patterns=Patterns(include=['site_2.com'], exclude=[], priority=500),
↳ use=<class 'gadget_sites_page_objects.site_2.GadgetSite2'>, instead_of=<class 'gadget_
↳ sites_page_objects.GadgetGenericPage'>, meta={})
# 4. OverrideRule(for_patterns=Patterns(include=['site_3.com'], exclude=[], priority=500),
↳ use=<class 'gadget_sites_page_objects.site_3.GadgetSite3'>, instead_of=<class 'gadget_
↳ sites_page_objects.GadgetGenericPage'>, meta={})

@handle_urls("site_1.com", overrides=ecommerce_page_objects.EcomGenericPage,
↳ priority=1000)
class ImprovedEcomSite1(ecommerce_page_objects.site_1.EcomSite1):
    def to_item(self):
        ... # call super().to_item() and improve on the item's shortcomings

rules = default_registry.get_overrides()
print(rules)
# 1. OverrideRule(for_patterns=Patterns(include=['site_1.com'], exclude=[], priority=500),
↳ use=<class 'ecommerce_page_objects.site_1.EcomSite1'>, instead_of=<class 'ecommerce_
↳ page_objects.EcomGenericPage'>, meta={})
# 2. OverrideRule(for_patterns=Patterns(include=['site_2.com'], exclude=[], priority=500),
↳ use=<class 'ecommerce_page_objects.site_2.EcomSite2'>, instead_of=<class 'ecommerce_
↳ page_objects.EcomGenericPage'>, meta={})
# 3. OverrideRule(for_patterns=Patterns(include=['site_2.com'], exclude=[], priority=500),
↳ use=<class 'gadget_sites_page_objects.site_2.GadgetSite2'>, instead_of=<class 'gadget_
↳ sites_page_objects.GadgetGenericPage'>, meta={})
# 4. OverrideRule(for_patterns=Patterns(include=['site_3.com'], exclude=[], priority=500),
↳ use=<class 'gadget_sites_page_objects.site_3.GadgetSite3'>, instead_of=<class 'gadget_
↳ sites_page_objects.GadgetGenericPage'>, meta={})
# 5. OverrideRule(for_patterns=Patterns(include=['site_1.com'], exclude=[],
↳ priority=1000), use=<class 'my_project.ImprovedEcomSite1'>, instead_of=<class 'ecommerce_
↳ page_objects.EcomGenericPage'>, meta={})

```

Notice that we're adding a new *OverrideRule* for the same URL pattern for `site_1.com`.

When the time comes that a Page Object needs to be selected when parsing `site_1.com` and it needs to replace `ecommerce_page_objects.EcomGenericPage`, rules #1 and #5 will be the choices. However, since we've assigned a much **higher priority** for the new rule in #5 than the default 500 value, rule #5 will be chosen because of its higher priority value.

More details on this in the *Priority Resolution* subsection.

3.5 Handling conflicts from using Multiple External Packages

You might've observed from the previous section that retrieving the list of all *OverrideRule* from two different external packages may result in a conflict.

We can take a look at the rules for **#2** and **#3** when we were importing all available rules:

```
# 2. OverrideRule(for_patterns=Patterns(include=['site_2.com'], exclude=[], priority=500),
↳ use=<class 'ecommerce_page_objects.site_2.EcomSite2'>, instead_of=<class 'ecommerce_
↳page_objects.EcomGenericPage'>, meta={})
# 3. OverrideRule(for_patterns=Patterns(include=['site_2.com'], exclude=[], priority=500),
↳ use=<class 'gadget_sites_page_objects.site_2.GadgetSite2'>, instead_of=<class 'gadget_
↳sites_page_objects.GadgetGenericPage'>, meta={})
```

However, it's technically **NOT** a *conflict*, **yet**, since:

- `ecommerce_page_objects.site_2.EcomSite2` would only be used in **site_2.com** if `ecommerce_page_objects.EcomGenericPage` is to be replaced.
- The same case with `gadget_sites_page_objects.site_2.GadgetSite2` wherein it's only going to be utilized for **site_2.com** if the following is to be replaced: `gadget_sites_page_objects.GadgetGenericPage`.

It would be only become a conflict if both rules for **site_2.com** intend to replace the **same** Page Object.

However, let's suppose that there are some *OverrideRule* which actually result in a conflict. To give an example, let's suppose that rules **#2** and **#3** intends to replace the **same** Page Object. It would look something like:

```
# 2. OverrideRule(for_patterns=Patterns(include=['site_2.com'], exclude=[], priority=500),
↳ use=<class 'ecommerce_page_objects.site_2.EcomSite2'>, instead_of=<class 'common_items.
↳ProductGenericPage'>, meta={})
# 3. OverrideRule(for_patterns=Patterns(include=['site_2.com'], exclude=[], priority=500),
↳ use=<class 'gadget_sites_page_objects.site_2.GadgetSite2'>, instead_of=<class 'common_
↳items.ProductGenericPage'>, meta={})
```

Notice that the `instead_of` param are the same and only the `use` param remained different.

There are two main ways we recommend in solving this.

1. Priority Resolution

If you notice, the `for_patterns` attribute of *OverrideRule* is an instance of `url_matcher.Patterns`. This instance also has a `priority` param where a higher value will be chosen in times of conflict.

Note: The `url-matcher` library is the one responsible breaking such priority conflicts (*amongst others*). It's specifically discussed in this section: [rules-conflict-resolution](#).

Unfortunately, updating the `priority` value directly isn't possible as the `url_matcher.Patterns` is a **frozen data-class**. The same is true for *OverrideRule*. This is made by design so that they are hashable and could be deduplicated immediately without consequences of them changing in value.

The only way that the `priority` value can be changed is by creating a new *OverrideRule* with a different `priority` value (*higher if it needs more priority*). You don't necessarily need to *delete* the **old** *OverrideRule* since they will be resolved via `priority` anyways.

Creating a new *OverrideRule* with a higher `priority` could be as easy as:

1. Subclassing the Page Object in question.

2. Create a new `web_poet.handle_urls()` annotation with the same URL pattern and Page Object to override but with a much higher priority.

Here's an example:

```
from web_poet import default_registry, consume_modules, handle_urls
import ecommerce_page_objects, gadget_sites_page_objects, common_items

@handle_urls("site_2.com", overrides=common_items.ProductGenericPage, priority=1000)
class EcomSite2Copy(ecommerce_page_objects.site_1.EcomSite1):
    def to_item(self):
        return super().to_item()
```

Now, the conflicting #2 and #3 rules would never be selected because of the new `OverrideRule` having a much higher priority (see rule #4):

```
# 2. OverrideRule(for_patterns=Patterns(include=['site_2.com'], exclude=[], priority=500),
↳ use=<class 'ecommerce_page_objects.site_2.EcomSite2'>, instead_of=<class 'common_items.
↳ ProductGenericPage'>, meta={})
# 3. OverrideRule(for_patterns=Patterns(include=['site_2.com'], exclude=[], priority=500),
↳ use=<class 'gadget_sites_page_objects.site_2.GadgetSite2'>, instead_of=<class 'common_
↳ items.ProductGenericPage'>, meta={})
# 4. OverrideRule(for_patterns=Patterns(include=['site_2.com'], exclude=[],
↳ priority=1000), use=<class 'my_project.EcomSite2Copy'>, instead_of=<class 'common_items.
↳ ProductGenericPage'>, meta={})
```

A similar idea was also discussed in the *Improving on external Page Objects* section.

2. Specifically Selecting the Rules

When the last resort of priority-resolution doesn't work, then you could always specifically select the list of `OverrideRule` you want to use.

We **recommend** in creating an **inclusion**-list rather than an **exclusion**-list since the latter is quite brittle. For instance, an external package you're using has updated its rules and the exclusion strategy misses out on a few rules that were recently added. This could lead to a *silent-error* of receiving a different set of rules than expected.

This **inclusion**-list approach can be done by importing the Page Objects directly and creating instances of `OverrideRule` from it. You could also import all of the available `OverrideRule` using `get_overrides()` to sift through the list of available rules and manually selecting the rules you need.

Most of the time, the needed rules are the ones which uses the Page Objects we're interested in. Since `PageObjectRegistry` is a dict subclass, you can easily find the Page Object's rule using its `key`. Here's an example:

```
from web_poet import default_registry, consume_modules
import package_A, package_B, package_C

consume_modules("package_A", "package_B", "package_C")

rules = [
    default_registry[package_A.PageObject1], # OverrideRule(for_
↳ patterns=Patterns(include=['site_A.com'], exclude=[], priority=500), use=<class 'package_
↳ A.PageObject1'>, instead_of=<class 'GenericPage'>, meta={})
    default_registry[package_B.PageObject2], # OverrideRule(for_
↳ patterns=Patterns(include=['site_B.com'], exclude=[], priority=500), use=<class 'package_
↳ B.PageObject2'>, instead_of=<class 'GenericPage'>, meta={})
```

(continues on next page)

(continued from previous page)

```

    default_registry[package_C.PageObject3], # OverrideRule(for_
↳ patterns=Patterns(include=['site_C.com'], exclude=[], priority=500), use=<class 'package_
↳ C.PageObject3'>, instead_of=<class 'GenericPage'>, meta={})
]

```

Another approach would be using the `search_overrides()` functionality as described from this tutorial section: *Using only a subset of the available OverrideRules*. The `search_overrides()` is quite useful in cases wherein the **POP** contains a lot of rules as it presents a utility for programmatically searching for them.

Here's an example:

```

from url_matcher import Patterns
from web_poet import default_registry, consume_modules
import package_A, package_B, package_C

consume_modules("package_A", "package_B", "package_C")

rule_from_A = default_registry.search_overrides(use=package_A.PageObject1)
print(rule_from_A)
# [OverrideRule(for_patterns=Patterns(include=['site_A.com'], exclude=[], priority=500),
↳ use=<class 'package_A.PageObject1'>, instead_of=<class 'GenericPage'>, meta={})]

rule_from_B = default_registry.search_overrides(instead_of=GenericProductPage)
print(rule_from_B)
# []

rule_from_C = default_registry.search_overrides(for_patterns=Patterns(include=["site_C.
↳ com"]))
print(rule_from_C)
# [
#     OverrideRule(for_patterns=Patterns(include=['site_C.com'], exclude=[],
↳ priority=500), use=<class 'package_C.PageObject3'>, instead_of=<class 'GenericPage'>,
↳ meta={}),
#     OverrideRule(for_patterns=Patterns(include=['site_C.com'], exclude=[],
↳ priority=1000), use=<class 'package_C.PageObject3_improved'>, instead_of=<class
↳ 'GenericPage'>, meta={})
# ]

rules = rule_from_A + rule_from_B + rule_from_C

```


ADDITIONAL REQUESTS

Websites nowadays needs a lot of page interactions to display or load some key information. In most cases, these are done via AJAX requests. Some examples of these are:

- Clicking a button on a page to reveal other similar products.
- Clicking the “*Load More*” button to retrieve more images of a given item.
- Scrolling to the bottom of the page to load more items (*i.e. infinite scrolling*).
- Hovering on a certain webpage element that reveals a tool-tip containing additional page info.

As such, performing additional requests inside Page Objects are inevitable to properly extract data for some websites.

Warning: Additional requests made inside a Page Object aren’t meant to represent the **Crawling Logic** at all. They are simply a low-level way to interact with today’s websites which relies on a lot of page interactions to display its contents.

The key words “MUST”, “MUST NOT”, “REQUIRED”, “SHALL”, “SHALL NOT”, “SHOULD”, “SHOULD NOT”, “RECOMMENDED”, “MAY”, and “OPTIONAL” in this document are to be interpreted as described in RFC 2119.

4.1 HttpRequest

Additional requests are defined using a simple data container that represents a generic HTTP Request: *HttpRequest*. Here’s an example:

```
import json
import web_poet

request = web_poet.HttpRequest(
    url="https://www.api.example.com/product-pagination/",
    method="POST",
    headers={
        "Content-Type": "application/json;charset=UTF-8"
    },
    body=json.dumps(
        {
            "Page": page_num,
            "ProductID": product_id,
        }
    ).encode("utf-8"),
```

(continues on next page)

(continued from previous page)

```

)

print(request.url)          # https://www.api.example.com/product-pagination/
print(type(request.url))   # <class 'web_poet.page_inputs.http.RequestUrl'>
print(request.method)      # POST

print(type(request.headers)) # <class 'web_poet.page_inputs.HttpRequestHeaders'>
print(request.headers)     # <HttpRequestHeaders('Content-Type': 'application/json;
↪ charset=UTF-8')>
print(request.headers.get("content-type")) # application/json; charset=UTF-8
print(request.headers.get("does-not-exist")) # None

print(type(request.body))  # <class 'web_poet.page_inputs.HttpRequestBody'>
print(request.body)       # b'{"Page": 1, "ProductID": 123}'

```

There are a few things to take note here:

- `method` is simply a **string**.
- `url` is represented by the `RequestUrl` class.
- `headers` is represented by the `HttpRequestHeaders` class which resembles a `dict`-like interface. It supports case-insensitive header-key lookups as well as multi-key storage.
 - See `multidict.CIMultiDict` for the set of features since `HttpRequestHeaders` simply inherits from it.
- `body` is represented by the `HttpRequestBody` class which is simply a subclass of the `bytes` class. Using the `body` param of `HttpRequest` needs to have an input argument in bytes. In our code example, we've converted it from `str` to `bytes` using the `encode()` string method.

Most of the time though, what you'll be defining would be GET requests. Thus, it's perfectly fine to define them as:

```

import web_poet

request = web_poet.HttpRequest("https://api.example.com/product-info?id=123")

print(request.url)          # https://api.example.com/product-info?id=123
print(type(request.url))   # <class 'web_poet.page_inputs.http.RequestUrl'>
print(request.method)      # GET

print(type(request.headers)) # <class 'web_poet.page_inputs.HttpRequestHeaders'>
print(request.headers)     # <HttpRequestHeaders()>
print(request.headers.get("content-type")) # None
print(request.headers.get("does-not-exist")) # None

print(type(request.body))  # <class 'web_poet.page_inputs.HttpRequestBody'>
print(request.body)       # b''

```

The key take aways are:

- The default value of `method` is GET.
- `headers` still holds `HttpRequestHeaders` which doesn't contain anything.
- The same is true for `body` holding an empty `HttpRequestBody`.

Now that we know how *HttpRequest* are structured, defining them doesn't execute the actual requests at all. In order to do so, we'll need to feed it into the *HttpClient* which is defined in the next section (see *HttpClient* tutorial section).

4.2 HttpResponse

HttpResponse is what comes after a *HttpRequest* has been executed. It's typically returned by the methods from *HttpClient* (see *HttpClient* tutorial section) which holds the information regarding the response.

HttpResponse can also be used as a Page Object dependency, e.g. *WebPage* uses it.

Note: The additional requests are expected to perform redirections except when the method is HEAD. This means that the *HttpResponse* that you'll be receiving is already the end of the redirection trail.

Let's check out an example to see its internals:

```
import web_poet

response = web_poet.HttpResponse(
    url="https://www.api.example.com/product-pagination/",
    body='{"data": "value "}'.encode("utf-8"),
    status=200,
    headers={"Content-Type": "application/json;charset=UTF-8"}
)

print(response.url)          # https://www.api.example.com/product-pagination/
print(type(response.url))   # <class 'web_poet.page_inputs.http.ResponseUrl'>

print(response.body)        # b'{"data": "value \xf0\x9f\x91\x8d"}'
print(type(response.body))  # <class 'web_poet.page_inputs.HttpResponseBody'>

print(response.status)      # 200
print(type(response.status)) # <class 'int'>

print(response.headers)     # <HttpResponseHeaders('Content-Type': 'application/json;
    ↪ charset=UTF-8')>
print(type(response.headers)) # <class 'web_poet.page_inputs.HttpResponseHeaders'>
print(response.headers.get("content-type")) # application/json;charset=UTF-8
print(response.headers.get("does-not-exist")) # None

# These methods are also available:

print(response.body.declared_encoding()) # None
print(response.body.json())             # {'data': 'value '}

print(response.headers.declared_encoding()) # utf-8

print(response.encoding)                # utf-8
print(response.text)                     # {"data": "value "}
print(response.json())                   # {'data': 'value '}
```

Despite what the example above showcases, you won't be typically defining *HttpResponse* yourself as it's the implementing framework (see *Framework Expectations*) that's responsible for it. Nonetheless, it's important to understand

its underlying structure in order to better access its methods.

Here are the key take aways from the example above:

- `status` is simply an `int`.
- `url` is represented by the `ResponseUrl` class.
- `headers` is represented by the `HttpResponseHeaders` class. It's similar to `HttpRequestHeaders` where it inherits from `multidict.CIMultiDict`, granting it case-insensitive header-key lookups as well as multi-key storage.
 - The **encoding** can be derived using the `declared_encoding()` method. In this example, it was retrieved from the `Content-Type` header.
- `body` is represented by the `HttpResponseBody` class which is simply a subclass of the `bytes` class. Using the `body` param of `HttpResponse` needs to have an input argument in `bytes`. In our code example, we've converted it from `str` to `bytes` using the `encode()` string method.
 - Similar to the headers, the **encoding** can be derived using the `declared_encoding()`. In this case, it returned `None` since no encoding can be derived from the response body.
 - A `json()` method is also available to conveniently access decoded contents from JSON responses. It uses the derived **encoding** to properly decode the contents like the emoji.
- The `HttpResponse` class itself also have these convenient methods:
 - The `encoding()` property method returns the proper encoding of the response based on the availability of this hierarchy:
 - * user-specified encoding (*using the `_encoding` attribute*)
 - * header encodings
 - * body encodings
 - Instead of accessing the raw bytes values (*which doesn't represent the underlying content properly like the emoji*), the `text()` property method can be used which takes into account the derived **encoding** when decoding the bytes value.
 - The `json()` method is available as a shortcut to `HttpResponseBody`'s `json()` method.

We've only explored a JSON response as a result from an additional request. Let's take a look at another example having an HTML response:

```
import web_poet

response = web_poet.HttpResponse(
    url="https://www.api.example.com/product-pagination/",
    body=(
        '<html>'
        ' <head>'
        '   <title>Some page</title>'
        '   <meta http-equiv="Content-Type" content="text/html; charset=utf-8">'
        ' </head>'
        ' <body>Sample content </body>'
        '</html>'
    ).encode("utf-8"),
    status=200,
    headers={}
)
```

(continues on next page)

(continued from previous page)

```

print(response.headers.declared_encoding()) # None
print(response.body.declared_encoding())   # utf-8
print(response.encoding)                   # utf-8

print(response.body.json()) # JSONDecodeError
print(response.json())     # JSONDecodeError

print(type(response.selector)) # <class 'parsel.selector.Selector'>

print(response.selector.css("body ::text").get()) # Sample content
print(response.css("body ::text").get())         # Sample content

print(response.selector.xpath("//body/text()").get()) # Sample content
print(response.xpath("//body/text()").get())         # Sample content

```

The key take aways for this example are:

- The **encoding** is derived from the body inside the meta tags since the headers is empty for this example.
- Since we now have an HTML response, using `json()` method would raise a `JSONDecodeError` as a JSON document cannot be parsed from it.
- The `selector()` property is an instance of `parsel.selector.Selector`; there are also `css()` and `xpath()` methods.
 - Usually there's no need to use `selector()`, as `css()` and `xpath()` are available.

4.3 HttpClient

The main interface for executing additional requests would be `HttpClient`. It also has full support for `asyncio` enabling developers to perform additional requests asynchronously using `asyncio.gather()`, `asyncio.wait()`, etc. This means that `asyncio` could be used anywhere inside the Page Object, including the `to_item()` method.

In the previous section, we've explored how `HttpRequest` is defined. Let's see a few quick examples to see how to execute additional requests using the `HttpClient`.

4.3.1 Executing a HttpRequest instance

```

import attrs
import web_poet

@attrs.define
class ProductPage(web_poet.ItemWebPage):
    http_client: web_poet.HttpClient

    async def to_item(self):
        item = {
            "url": self.url,
            "name": self.css("#main h3.name ::text").get(),
            "product_id": self.css("#product ::attr(product-id)").get(),

```

(continues on next page)

(continued from previous page)

```

    }

    # Simulate clicking on a button that says "View All Images"
    request = web_poet.HttpRequest(f"https://api.example.com/v2/images?id={item[
↪'product_id']}")
    response: web_poet.HttpResponse = await self.http_client.execute(request)

    item["images"] = response.css(".product-images img::attr(src)").getall()
    return item

```

As the example suggests, we're performing an additional request that allows us to extract more images in a product page that might not be otherwise be possible. This is because in order to do so, an additional button needs to be clicked which fetches the complete set of product images via AJAX.

There are a few things to take note of this example:

- Recall from the *HttpRequest* tutorial section that the default method is GET. Thus, the method parameter can be omitted for simple GET requests.
- We're now using the `async/await` syntax inside the `to_item()` method.
- The response from the additional request is of type *HttpResponse*.

Tip: Check out the *Batch requests* tutorial section to see how to execute a group of *HttpRequest* in batch.

Fortunately, there are already some quick shortcuts on how to perform single additional requests using the `request()`, `get()`, and `post()` methods of *HttpClient*. These already define the *HttpRequest* and executes it as well.

4.3.2 A simple GET request

Let's use the example from the previous section and use the `get()` method on it.

```

import attrs
import web_poet

@attrs.define
class ProductPage(web_poet.ItemWebPage):
    http_client: web_poet.HttpClient

    async def to_item(self):
        item = {
            "url": self.url,
            "name": self.css("#main h3.name ::text").get(),
            "product_id": self.css("#product ::attr(product-id)").get(),
        }

        # Simulates clicking on a button that says "View All Images"
        response: web_poet.HttpResponse = await self.http_client.get(
            f"https://api.example.com/v2/images?id={item['product_id']}"
        )
        item["images"] = response.css(".product-images img::attr(src)").getall()
        return item

```

There are a few things to take note in this example:

- A GET request can be done via `HttpClient`'s `get()` method.
- There is no need create an instance of `HttpRequest` when `get()` is used.

4.3.3 A POST request with *header* and *body*

Let's see another example which needs headers and body data to process additional requests.

In this example, we'll paginate related items in a carousel. These are usually lazily loaded by the website to reduce the amount of information rendered in the DOM that might not otherwise be viewed by all users anyway.

Thus, additional requests inside the Page Object are typically needed for it:

```
import attrs
import web_poet

@attrs.define
class ProductPage(web_poet.ItemWebPage):
    http_client: web_poet.HttpClient

    async def to_item(self):
        item = {
            "url": self.url,
            "name": self.css("#main h3.name ::text").get(),
            "product_id": self.css("#product ::attr(product-id)").get(),
            "related_product_ids": self.parse_related_product_ids(self),
        }

        # Simulates "scrolling" through a carousel that loads related product items
        response: web_poet.HttpResponse = await self.http_client.post(
            url="https://www.api.example.com/related-products/",
            headers={
                "Content-Type": "application/json;charset=UTF-8"
            },
            body=json.dumps(
                {
                    "Page": 2,
                    "ProductID": item["product_id"],
                }
            ).encode("utf-8"),
        )
        item["related_product_ids"].extend(self.parse_related_product_ids(response))
        return item

    @staticmethod
    def parse_related_product_ids(response_page) -> List[str]:
        return response_page.css("#main .related-products ::attr(product-id)").getall()
```

Here's the key takeaway in this example:

- Similar to `HttpClient`'s `get()` method, a `post()` method is also available. It is often used to submit forms.

4.3.4 Other Single Requests

The `get()` and `post()` methods are merely quick shortcuts for `request()`:

```
client = HttpClient()

url = "https://api.example.com/v1/data"
headers = {"Content-Type": "application/json;charset=UTF-8"}
body = b'{"data": "value"}'

# These are the same:
response = await client.get(url)
response = await client.request(url, method="GET")

# The same goes for these:
response = await client.post(url, headers=headers, body=body)
response = await client.request(url, method="POST", headers=headers, body=body)
```

Thus, apart from the common GET and POST HTTP methods, you can use `request()` for them (e.g. HEAD, PUT, DELETE, etc).

4.3.5 Batch requests

We can also choose to process requests by **batch** instead of sequentially or one by one (e.g. using `execute()`). The `batch_execute()` method can be used for this which accepts an arbitrary number of `HttpRequest` instances.

Let's modify the example in the previous section to see how it can be done.

The difference for this code example from the previous section is that we're increasing the pagination from only the **2nd page** into the **10th page**. Instead of calling a single `post()` method, we're creating a list of `HttpRequest` to be executed in batch using the `batch_execute()` method.

```
from typing import List

import attrs
import web_poet

@attrs.define
class ProductPage(web_poet.ItemWebPage):
    http_client: web_poet.HttpClient

    default_pagination_limit = 10

    async def to_item(self):
        item = {
            "url": self.url,
            "name": self.css("#main h3.name ::text").get(),
            "product_id": self.css("#product ::attr(product-id)").get(),
            "related_product_ids": self.parse_related_product_ids(self),
        }

        requests: List[web_poet.HttpRequest] = [
            self.create_request(item["product_id"], page_num=page_num)
```

(continues on next page)

(continued from previous page)

```

        for page_num in range(2, self.default_pagination_limit)
    ]
    responses: List[web_poet.HttpResponse] = await self.http_client.batch_
↪execute(*requests)
    related_product_ids = [
        id_
        for response in responses
        for product_ids in self.parse_related_product_ids(response)
        for id_ in product_ids
    ]

    item["related_product_ids"].extend(related_product_ids)
    return item

def create_request(self, product_id, page_num=2):
    # Simulates "scrolling" through a carousel that loads related product items
    return web_poet.HttpRequest(
        url="https://www.api.example.com/product-pagination/",
        method="POST",
        headers={
            "Content-Type": "application/json;charset=UTF-8"
        },
        body=json.dumps(
            {
                "Page": page_num,
                "ProductID": product_id,
            }
        ).encode("utf-8"),
    )

@staticmethod
def parse_related_product_ids(response_page) -> List[str]:
    return response_page.css("#main .related-products ::attr(product-id)").getall()

```

The key takeaways for this example are:

- An *HttpRequest* can be instantiated to represent a Generic HTTP Request. It only contains the HTTP Request information for now and isn't executed yet. This is useful for creating factory methods to help create requests without any download execution at all.
- *HttpClient* has a *batch_execute()* method that can process a list of *HttpRequest* instances asynchronously together.

Tip: The *batch_execute()* method can execute multiple *HttpRequest* instances. For example, it could be a mixture of GET and POST requests or even representing requests for various parts of the page altogether.

Processing the additional requests in batch is useful since it takes advantage of async execution which could be faster in certain cases (*assuming you're allowed to perform HTTP requests in parallel*).

Nonetheless, you can still use the *batch_execute()* method to execute a single *HttpRequest* instance.

Note: The *batch_execute()* method is a simple wrapper over *asyncio.gather()*. Developers are free to use other functionalities available inside *asyncio* to handle multiple requests.

For example, `asyncio.as_completed()` can be used to process the first response from a group of requests as early as possible. However, the order could be shuffled.

4.4 Handling Exceptions in Page Objects

Let's have a look at how we could handle exceptions when performing additional requests inside Page Objects. For this example, let's improve the code snippet from the previous subsection named: *A simple GET request*.

```
import logging

import attrs
import web_poet

logger = logging.getLogger(__name__)

@attrs.define
class ProductPage(web_poet.ItemWebPage):
    http_client: web_poet.HttpClient

    async def to_item(self):
        item = {
            "url": self.url,
            "name": self.css("#main h3.name ::text").get(),
            "product_id": self.css("#product ::attr(product-id)").get(),
        }

        try:
            # Simulates clicking on a button that says "View All Images"
            response: web_poet.HttpResponse = await self.http_client.get(
                f"https://api.example.com/v2/images?id={item['product_id']}"
            )
        except web_poet.exceptions.HttpRequestError as err:
            logger.warning(
                f"Unable to request images for product ID '{item['product_id']}' "
                f"using this request: {err.request}"
            )
        except web_poet.exceptions.HttpResponseError as err:
            logger.warning(
                f"Received a {err.response.status} response status for product ID "
                f"'{item['product_id']}' from this URL: {err.request.url}"
            )
        else:
            item["images"] = response.css(".product-images img::attr(src)").getall()

        return item
```

In this code example, the code became more resilient on cases where it wasn't possible to retrieve more images using the website's public API. It could be due to anything like *SSL errors*, *connection errors*, *page not found*, etc.

Using `HttpClient` to execute requests raises exceptions with the base class of type `web_poet.exceptions.http.HttpError` irregardless of how the HTTP Downloader is implemented. From our example above, we could've simply

used the `web_poet.exceptions.http.HttpError` base error. However, it's ambiguous in the sense that the error could originate during the HTTP Request execution or when receiving the HTTP Response.

A more specific `web_poet.exceptions.http.HttpRequestError` exception is raised when the `HttpRequest` was being handled while the `web_poet.exceptions.http.HttpResponseError` is raised when receiving a response with an HTTP error. Notice from the example that the exceptions have the attributes like `request` and `response` which are respective instance of `HttpRequest` and `HttpResponse`. Accessing them would be useful to debug and log the problems.

Note that `web_poet.exceptions.http.HttpResponseError` only occurs when receiving responses with status codes in the 400–5xx range. However, this behavior could be altered by using the `allow_status` param in the methods of `HttpClient`.

Note: In the future, more specific exceptions which inherits from the base `web_poet.exceptions.http.HttpError` exception would be available. This should allow developers writing Page Objects to properly identify what went wrong and act specifically based on the problem.

Let's take another example when executing requests in batch as opposed to using single requests via these methods of the `HttpClient`: `request()`, `get()`, and `post()`.

For this example, let's improve the code snippet from the previous subsection named: *Batch requests*.

```
import logging
from typing import List, Union

import attrs
import web_poet

@attrs.define
class ProductPage(web_poet.ItemWebPage):
    http_client: web_poet.HttpClient

    default_pagination_limit = 10

    async def to_item(self):
        item = {
            "url": self.url,
            "name": self.css("#main h3.name ::text").get(),
            "product_id": self.css("#product ::attr(product-id)").get(),
            "related_product_ids": self.parse_related_product_ids(self),
        }

        requests: List[web_poet.HttpRequest] = [
            self.create_request(item["product_id"], page_num=page_num)
            for page_num in range(2, self.default_pagination_limit)
        ]

        try:
            responses: List[web_poet.HttpResponse] = await self.http_client.batch_
↪ execute(*requests)
            except web_poet.exceptions.HttpError:
                logger.warning(
                    f"Unable to request for more related products for product ID: {item[
↪ 'product_id']}"

```

(continues on next page)

(continued from previous page)

```

    )
    else:
        related_product_ids = []
        for response in responses:
            related_product_ids.extend(
                [
                    id_
                    for product_ids in self.parse_related_product_ids(response)
                    for id_ in product_ids
                ]
            )
            item["related_product_ids"].extend(related_product_ids)

    return item

def create_request(self, product_id, page_num=2):
    # Simulates "scrolling" through a carousel that loads related product items
    return web_poet.HttpRequest(
        url="https://www.api.example.com/product-pagination/",
        method="POST",
        headers={
            "Content-Type": "application/json;charset=UTF-8"
        },
        body=json.dumps(
            {
                "Page": page_num,
                "ProductID": product_id,
            }
        ).encode("utf-8"),
    )

@staticmethod
def parse_related_product_ids(response_page) -> List[str]:
    return response_page.css("#main .related-products ::attr(product-id)").getall()

```

Handling exceptions using `batch_execute()` remains largely the same. However, the main difference is that you may be wasting perfectly good responses just because a single request from the batch ruined it. Notice that we're using the base exception class of `web_poet.exceptions.http.HttpError` to account for any type of errors, both during the HTTP Request execution and when receiving the response.

An alternative approach would be salvaging good responses altogether. For example, you've sent out 10 `HttpRequest` and only 1 of them had an exception during processing. You can still get the data from 9 of the `HttpResponse` by passing the parameter `return_exceptions=True` to `batch_execute()`.

This means that any exceptions raised during the HTTP execution are returned alongside any of the successful responses. The return type of `batch_execute()` could be a mixture of `HttpResponse` and `web_poet.exceptions.http.HttpError` (and its exception subclasses).

Here's an example:

```

# Revised code snippet from the to_item() method

requests: List[web_poet.HttpRequest] = [
    self.create_request(item["product_id"], page_num=page_num)

```

(continues on next page)

(continued from previous page)

```

    for page_num in range(2, self.default_pagination_limit)
]
responses: List[Union[web_poet.HttpResponse, web_poet.exceptions.HttpError]] = (
    await self.http_client.batch_execute(*requests, return_exceptions=True)
)

related_product_ids = []
for i, response in enumerate(responses):
    if isinstance(response, web_poet.exceptions.HttpError):
        logger.warning(
            f"Unable to request related products for product ID '{item['product_id']}' "
            f"using this request: {requests[i]}. Reason: {response}."
        )
        continue
    related_product_ids.extend(
        [
            id_
            for product_ids in self.parse_related_product_ids(response)
            for id_ in product_ids
        ]
    )

item["related_product_ids"].extend(related_product_ids)
return item

```

From the example above, we're now checking the list of responses to see if any exceptions are included in it. If so, we're simply logging it down and ignoring it. In this way, perfectly good responses can still be processed through.

4.5 Framework Expectations

In the earlier sections, the tutorial was primarily focused on how Page Object developers could use additional requests using the **web-poet**'s built-in functionalities. However, as the docs have repeatedly mentioned, **web-poet** doesn't know how to execute any of these HTTP requests at all. It would be up to the framework that's handling **web-poet** to do so.

In this section, we'll explore the guidelines for frameworks. If you're a Page Object developer, you can skip this part as it mostly discusses the internals. However, reading through this section could render a better understanding of **web-poet** as a whole.

4.5.1 Providing the Downloader

On its own, *HttpClient* doesn't do anything. It doesn't know how to execute the request on its own. Thus, for frameworks or projects wanting to use additional requests in Page Objects, they need to set the implementation on how to execute an *HttpRequest*.

For more info on this, kindly read the API Specifications for *HttpClient*.

In any case, frameworks that wish to support **web-poet** could provide the HTTP downloader implementation in two ways:

1. Context Variable

`contextvars` is natively supported in `asyncio` in order to set and access context-aware values. This means that the framework using **web-poet** can assign the request downloader implementation using the `contextvars` instance named `web_poet.request_downloader_var`.

This can be set using:

```
import attrs
import web_poet

async def request_implementation(req: web_poet.HttpRequest) -> web_poet.HttpResponse:
    ...

def create_http_client():
    return web_poet.HttpClient()

@attrs.define
class SomePage(web_poet.ItemWebPage):
    http_client: web_poet.HttpClient

    async def to_item(self):
        ...

# Once this is set, the `request_implementation` becomes available to
# all instances of HttpClient, unless HttpClient is created with
# the `request_downloader` argument (see the #2 Dependency Injection
# example below).
web_poet.request_downloader_var.set(request_implementation)

# Assume that it's constructed with the necessary arguments taken somewhere.
response = web_poet.HttpResponse(...)

page = SomePage(response=response, http_client=create_http_client())
item = await page.to_item()
```

When the `web_poet.request_downloader_var` contextvar is set, `HttpClient` instances use it by default.

Warning: If no value for `web_poet.request_downloader_var` is set, then `RequestDownloaderVarError` is raised. However, no exception is raised if **option 2** below is used.

2. Dependency Injection

The framework using **web-poet** may be using libraries that don't have a full support to `contextvars` (e.g. *Twisted*). With that, an alternative approach would be to supply the request downloader implementation when creating an `HttpClient` instance:

```
import attrs
import web_poet

async def request_implementation(req: web_poet.HttpRequest) -> web_poet.HttpResponse:
    ...

def create_http_client():
    return web_poet.HttpClient(request_downloader=request_implementation)

@attrs.define
class SomePage(web_poet.ItemWebPage):
    http_client: web_poet.HttpClient

    async def to_item(self):
        ...

# Assume that it's constructed with the necessary arguments taken somewhere.
response = web_poet.HttpResponse(...)

page = SomePage(response=response, http_client=create_http_client())
item = await page.to_item()
```

From the code sample above, we can see that every time an `HttpClient` instance is created for Page Objects needing an `http_client`, the framework must create `HttpClient` with a framework-specific **request downloader implementation**, using the `request_downloader` argument.

4.5.2 Downloader Behavior

The request downloader **MUST** accept an instance of `HttpRequest` as the input and return an instance of `HttpResponse`. This is important in order to handle and represent generic HTTP operations. The only time that it won't be returning `HttpResponse` would be when it's raising exceptions (see *Exception Handling*).

The request downloader **MUST** resolve Location-based **redirections** when the HTTP method is not HEAD. In other words, for non-HEAD requests the returned `HttpResponse` must be the final response, after all redirects. For HEAD requests redirects **MUST NOT** be resolved.

Lastly, the request downloader function **MUST** support the `async/await` syntax.

4.5.3 Exception Handling

In the previous *Handling Exceptions in Page Objects* section, we can see how Page Object developers could use the exception classes built inside **web-poet** to handle various ways additional requests MAY fail. In this section, we'll see the rationale and ways the framework MUST be able to do that.

Rationale

Frameworks that handle **web-poet** MUST be able to ensure that Page Objects having additional requests using *HttpClient* are able to work with any type of HTTP downloader implementation.

For example, in Python, the common HTTP libraries have different types of base exceptions when something has occurred:

- `aiohttp.ClientError`
- `requests.RequestException`
- `urllib.error.HTTPError`

Imagine if Page Objects are **expected** to work in *different* backend implementations like the ones above, then it would cause the code to look like:

```
import attrs
import web_poet

import aiohttp
import requests
import urllib

@attrs.define
class SomePage(web_poet.ItemWebPage):
    http_client: web_poet.HttpClient

    async def to_item(self):
        try:
            response = await self.http_client.get("...")
        except (aiohttp.ClientError, requests.RequestException, urllib.error.HTTPError):
            # handle the error here
```

Such code could turn messy in no time especially when the number of HTTP backends that Page Objects have to support are steadily increasing. Not to mention the plethora of exception types that HTTP libraries have. This means that Page Objects aren't truly portable in different types of frameworks or environments. Rather, they're only limited to work in the specific framework they're supported.

In order for Page Objects to work in different Downloader Implementations, the framework that implements the HTTP Downloader backend MUST raise exceptions from the `web_poet.exceptions.http` module in lieu of the backend specific ones (e.g. `aiohttp`, `requests`, `urllib`, etc.).

This makes the code simpler:

```
import attrs
import web_poet
```

(continues on next page)

(continued from previous page)

```
@attrs.define
class SomePage(web_poet.ItemWebPage):
    http_client: web_poet.HttpClient

    async def to_item(self):
        try:
            response = await self.http_client.get(...)
        except web_poet.exceptions.HttpError:
            # handle the error here
```

Expected behavior for Exceptions

All exceptions that the HTTP Downloader Implementation (see *Providing the Downloader* doc section) explicitly raises when implementing it for **web-poet** MUST be `web_poet.exceptions.http.HttpError` (or a subclass from it).

For frameworks that implement and use **web-poet**, exceptions that occurred when handling the additional requests like *connection errors*, *TLS errors*, etc MUST be replaced by `web_poet.exceptions.http.HttpRequestError` by raising it explicitly.

For responses that are not really errors like in the 100-3xx status code range, exception MUST NOT be raised at all. For responses with status codes in the 400-5xx range, **web-poet** raises the `web_poet.exceptions.http.HttpResponseError` exception.

From this distinction, the framework MUST NOT raise `web_poet.exceptions.http.HttpResponseError` on its own at all, since the `HttpClient` already handles that.

RETRIES

The responses of some websites can be unreliable. For example, sometimes a request can get a response that may only include a part of the data to be extracted, no data at all, or even data unrelated to your request, but sending a follow-up, identical request can get you the expected data.

Pages objects are responsible for handling these scenarios, where issues with response data can only be detected during extraction.

5.1 Retrying Page Object Input

When the bad response data comes from the inputs that your web-poet framework supplies to your page object, your page object must raise *Retry*:

```
from web_poet import ItemWebPage
from web_poet.exceptions import Retry

class MyPage(ItemWebPage):

    def to_item(self) -> dict:
        if not self.css('.expected'):
            raise Retry
        return {}
```

As a result, your web-poet framework will retry the source requests and create a new instance of your page object with the new inputs.

5.2 Retrying Additional Requests

When the bad response data comes from *additional requests*, you must handle retries on your own.

The page object code is responsible for retrying additional requests until good response data is received, or until some maximum number of retries is exceeded.

It is up to you to decide what the maximum number of retries should be for a given additional request, based on your experience with the target website.

It is also up to you to decide how to implement retries of additional requests.

One option would be *tenacity*. For example, to try an additional request 3 times before giving up:

```
import attrs
from tenacity import retry, stop_after_attempt
from web_poet import HttpClient, HttpRequest, ItemWebPage

@attrs.define
class MyPage(ItemWebPage):
    http_client: HttpClient

    @retry(stop=stop_after_attempt(3))
    async def get_data(self):
        request = HttpRequest('https://toscrape.com/')
        response = await self.http_client.execute(request)
        if not response.css('.expected'):
            raise ValueError
        return response.css('.data').get()

    async def to_item(self) -> dict:
        try:
            data = await self.get_data()
        except ValueError:
            return {}
        return {'data': data}
```

If the reason your additional request fails is outdated or missing data from page object input, do not try to reproduce the request for that input as an additional request. *Request fresh input instead.*

5.3 Framework Expectations

Web-poet frameworks must catch *Retry* exceptions raised from the *to_item()* method of a page object.

When *Retry* is caught:

1. The original request whose response was fed into the page object must be retried.
2. A new page object must be created, of the same type as the original page object, and with the same input, except for the response data, which must be the new response.

The *to_item()* method of the new page object may raise *Retry* again. Web-poet frameworks must allow multiple retries of page objects, repeating the *Retry*-capturing logic.

However, web-poet frameworks are also encouraged to limit the amount of retries per page object. When retries are exceeded for a given page object, the page object output is ignored. At the moment, web-poet does not enforce any specific maximum number of retries on web-poet frameworks.

USING PAGE PARAMS

In some cases, Page Objects might require additional information to be passed to them. Such information can dictate the behavior of the Page Object or affect its data entirely depending on the needs of the developer.

If you can recall from the previous basic tutorials, one essential requirement of Page Objects that inherit from *WebPage* or *ItemWebPage* would be *HttpResponse*. This holds the HTTP response information that the Page Object is trying to represent.

In order to standardize how to pass arbitrary information inside Page Objects, we'll need to use *PageParams* similar on how we use *HttpResponse* as a requirement to instantiate Page Objects:

```
import attrs
import web_poet

@attrs.define
class SomePage(web_poet.ItemWebPage):
    # The HttpResponse attribute is inherited from ItemWebPage
    page_params: web_poet.PageParams

    # Assume that it's constructed with the necessary arguments taken somewhere.
    response = web_poet.HttpResponse(...)

    # It uses Python's dict interface.
    page_params = web_poet.PageParams({"arbitrary_value": 1234, "cool": True})

page = SomePage(response=response, page_params=page_params)
```

However, similar with *HttpResponse*, developers using *PageParams* shouldn't care about how they are being passed into Page Objects. This will depend on the framework that would use **web-poet**.

Let's checkout some examples on how to use it inside a Page Object.

6.1 Controlling item values

```
import attrs
import web_poet

@attrs.define
class ProductPage(web_poet.ItemWebPage):
    page_params: web_poet.PageParams
```

(continues on next page)

(continued from previous page)

```

default_tax_rate = 0.10

def to_item(self):
    item = {
        "url": self.url,
        "name": self.css("#main h3.name ::text").get(),
        "price": self.css("#main .price ::text").get(),
    }
    self.calculate_price_with_tax(item)
    return item

@staticmethod
def calculate_price_with_tax(item):
    tax_rate = self.page_params.get("tax_rate") or self.default_tax_rate
    item["price_with_tax"] = item["price"] * (1 + tax_rate)

```

From the example above, we were able to provide an optional information regarding the **tax rate** of the product. This could be useful when trying to support the different tax rates for each state or territory. However, since we're treating the **tax_rate** as optional information, notice that we also have a the `default_tax_rate` as a backup value just in case it's not available.

6.2 Controlling Page Object behavior

Let's try an example wherein *PageParams* is able to control how *Additional Requests* are being used. Specifically, we are going to use *PageParams* to control the number of paginations being made.

```

from typing import List

import attrs
import web_poet

@attrs.define
class ProductPage(web_poet.ItemWebPage):
    http_client: web_poet.HttpClient
    page_params: web_poet.PageParams

    default_max_pages = 5

    async def to_item(self):
        return {"product_urls": await self.get_product_urls()}

    async def get_product_urls(self) -> List[str]:
        # Simulates scrolling to the bottom of the page to load the next
        # set of items in an "Infinite Scrolling" category list page.
        max_pages = self.page_params.get("max_pages") or self.default_max_pages
        requests = [
            self.create_next_page_request(page_num)
            for page_num in range(2, max_pages + 1)
        ]

```

(continues on next page)

(continued from previous page)

```
]
responses = await http_client.batch_execute(*requests)
return [
    url
    for response in responses
    for product_urls in self.parse_product_urls(response)
    for url in product_urls
]

@staticmethod
def create_next_page_request(page_num):
    next_page_url = f"https://example.com/category/products?page={page_num}"
    return web_poet.Request(url=next_page_url)

@staticmethod
def parse_product_urls(response: web_poet.HttpResponse):
    return response.css("#main .products a.link ::attr(href)").getall()
```

From the example above, we can see how *PageParams* is able to arbitrarily limit the pagination behavior by passing an optional **max_pages** info. Take note that a `default_max_pages` value is also present in the Page Object in case the *PageParams* instance did not provide it.

7.1 Page Inputs

class `web_poet.page_inputs.browser.BrowserHtml`

Bases: `SelectableMixin`, `str`

HTML returned by a web browser, i.e. snapshot of the DOM tree in HTML format.

css(*query*) → `SelectorList`

A shortcut to `.selector.css()`.

property selector: `Selector`

Cached instance of `parsel.selector.Selector`.

xpath(*query*, ***kwargs*) → `SelectorList`

A shortcut to `.selector.xpath()`.

class `web_poet.page_inputs.http.RequestUrl(*args, **kwargs)`

Bases: `RequestUrl`

class `web_poet.page_inputs.http.ResponseUrl(*args, **kwargs)`

Bases: `ResponseUrl`

class `web_poet.page_inputs.http.HttpRequestBody`

Bases: `bytes`

A container for holding the raw HTTP request body in bytes format.

class `web_poet.page_inputs.http.HttpResponseBody`

Bases: `bytes`

A container for holding the raw HTTP response body in bytes format.

declared_encoding() → `Optional[str]`

Return the encoding specified in meta tags in the html body, or `None` if no suitable encoding was found

json()

Deserialize a JSON document to a Python object.

class `web_poet.page_inputs.http.HttpRequestHeaders`

Bases: `_HttpHeaders`

A container for holding the HTTP request headers.

It's able to accept instantiation via an Iterable of Tuples:

```
>>> pairs = [("Content-Encoding", "gzip"), ("content-length", "648")]
>>> HttpRequestHeaders(pairs)
<HttpRequestHeaders('Content-Encoding': 'gzip', 'content-length': '648')>
```

It's also accepts a mapping of key-value pairs as well:

```
>>> pairs = {"Content-Encoding": "gzip", "content-length": "648"}
>>> headers = HttpRequestHeaders(pairs)
>>> headers
<HttpRequestHeaders('Content-Encoding': 'gzip', 'content-length': '648')>
```

Note that this also supports case insensitive header-key lookups:

```
>>> headers.get("content-encoding")
'gzip'
>>> headers.get("Content-Length")
'648'
```

These are just a few of the functionalities it inherits from `multidict.CIMultiDict`. For more info on its other features, read the API spec of `multidict.CIMultiDict`.

`copy()`

Return a copy of itself.

classmethod `from_name_value_pairs`(arg: *List[Dict]*) → T_headers

An alternative constructor for instantiation using a `List[Dict]` where the 'key' is the header name while the 'value' is the header value.

```
>>> pairs = [
...     {"name": "Content-Encoding", "value": "gzip"},
...     {"name": "content-length", "value": "648"}
... ]
>>> headers = _HttpHeaders.from_name_value_pairs(pairs)
>>> headers
<_HttpHeaders('Content-Encoding': 'gzip', 'content-length': '648')>
```

class `web_poet.page_inputs.http.HttpResponseHeaders`

Bases: `_HttpHeaders`

A container for holding the HTTP response headers.

It's able to accept instantiation via an Iterable of Tuples:

```
>>> pairs = [("Content-Encoding", "gzip"), ("content-length", "648")]
>>> HttpResponseHeaders(pairs)
<HttpResponseHeaders('Content-Encoding': 'gzip', 'content-length': '648')>
```

It's also accepts a mapping of key-value pairs as well:

```
>>> pairs = {"Content-Encoding": "gzip", "content-length": "648"}
>>> headers = HttpResponseHeaders(pairs)
>>> headers
<HttpResponseHeaders('Content-Encoding': 'gzip', 'content-length': '648')>
```

Note that this also supports case insensitive header-key lookups:

```
>>> headers.get("content-encoding")
'gzip'
>>> headers.get("Content-Length")
'648'
```

These are just a few of the functionalities it inherits from `multidict.CIMultiDict`. For more info on its other features, read the API spec of `multidict.CIMultiDict`.

classmethod `from_bytes_dict`(*arg: Dict[AnyStr, Union[AnyStr, List, Tuple[AnyStr, ...]]], encoding: str = 'utf-8'*) → `T_headers`

An alternative constructor for instantiation where the header-value pairs could be in raw bytes form.

This supports multiple header values in the form of `List[bytes]` and `Tuple[bytes]` alongside a plain bytes value. A value in `str` also works and wouldn't break the decoding process at all.

By default, it converts the bytes value using “utf-8”. However, this can easily be overridden using the `encoding` parameter.

```
>>> raw_values = {
...     b"Content-Encoding": [b"gzip", b"br"],
...     b"Content-Type": [b"text/html"],
...     b"content-length": b"648",
... }
>>> headers = HttpResponseHeaders.from_bytes_dict(raw_values)
>>> headers
<HttpResponseHeaders('Content-Encoding': 'gzip', 'Content-Encoding': 'br',
↳ 'Content-Type': 'text/html', 'content-length': '648')>
```

declared_encoding() → `Optional[str]`

Return encoding detected from the Content-Type header, or `None` if encoding is not found

copy()

Return a copy of itself.

classmethod `from_name_value_pairs`(*arg: List[Dict]*) → `T_headers`

An alternative constructor for instantiation using a `List[Dict]` where the ‘key’ is the header name while the ‘value’ is the header value.

```
>>> pairs = [
...     {"name": "Content-Encoding", "value": "gzip"},
...     {"name": "content-length", "value": "648"}
... ]
>>> headers = _HttpHeaders.from_name_value_pairs(pairs)
>>> headers
<_HttpHeaders('Content-Encoding': 'gzip', 'content-length': '648')>
```

class `web_poet.page_inputs.http.HttpRequest`(*url: Union[str, _Url], *, method: str = 'GET', headers=NOTHING, body=NOTHING*)

Bases: `object`

Represents a generic HTTP request used by other functionalities in **web-poet** like `HttpClient`.

url: `RequestUrl`

method: `str`

headers: [HttpRequestHeaders](#)

body: [HttpRequestBody](#)

urljoin(*url: Union[str, RequestUrl, ResponseUrl]*) → RequestUrl

Return *url* as an absolute URL.

If *url* is relative, it is made absolute relative to *url*.

```
class web_poet.page_inputs.http.HttpResponse(url: Union[str, _Url], body, *, status: Optional[int] =
None, headers=NOTHING, encoding: Optional[str] =
None)
```

Bases: [SelectableMixin](#)

A container for the contents of a response, downloaded directly using an HTTP client.

url should be a URL of the response (after all redirects), not a URL of the request, if possible.

body contains the raw HTTP response body.

The following are optional since it would depend on the source of the `HttpResponse` if these are available or not. For example, the responses could simply come off from a local HTML file which doesn't contain headers and status.

status should represent the int status code of the HTTP response.

headers should contain the HTTP response headers.

encoding encoding of the response. If `None` (default), encoding is auto-detected from headers and body content.

url: [ResponseUrl](#)

body: [HttpResponseBody](#)

status: [Optional\[int\]](#)

headers: [HttpResponseHeaders](#)

property text: [str](#)

Content of the HTTP body, converted to unicode using the detected encoding of the response, according to the web browser rules (respecting Content-Type header, etc.)

property encoding

Encoding of the response

json()

Deserialize a JSON document to a Python object.

urljoin(*url: Union[str, RequestUrl, ResponseUrl]*) → RequestUrl

Return *url* as an absolute URL.

If *url* is relative, it is made absolute relative to the base URL of *self*.

css(*query*) → [SelectorList](#)

A shortcut to `.selector.css()`.

property selector: [Selector](#)

Cached instance of `parsel.selector.Selector`.

`xpath(query, **kwargs) → SelectorList`

A shortcut to `.selector.xpath()`.

class `web_poet.page_inputs.page_params.PageParams`

Bases: `dict`

Container class that could contain any arbitrary data to be passed into a Page Object.

Note that this is simply a subclass of Python’s `dict`.

class `web_poet.page_inputs.client.HttpClient(request_downloader: Optional[Callable] = None)`

Async HTTP client to be used in Page Objects.

See *Additional Requests* for the usage information.

`HttpClient` doesn’t make HTTP requests on itself. It uses either the request function assigned to the `web_poet.request_downloader_var` `contextvar`, or a function passed via `request_downloader` argument of the `__init__()` method.

Either way, this function should be an `async def` function which receives an `HttpRequest` instance, and either returns a `HttpResponse` instance, or raises a subclass of `HttpError`. You can read more in the *Providing the Downloader* documentation.

async request(`url: Union[str, _Url]`, *, `method: str = 'GET'`, `headers: Optional[Union[Dict[str, str], HttpRequestHeaders]] = None`, `body: Optional[Union[bytes, HttpRequestBody]] = None`, `allow_status: Optional[List[Union[str, int]]] = None`) → `HttpResponse`

This is a shortcut for creating an `HttpRequest` instance and executing that request.

`HttpRequestError` is raised for *connection errors, connection and read timeouts*, etc.

An `HttpResponse` instance is returned for successful responses in the 100–3xx status code range.

Otherwise, an exception of type `HttpResponseError` is raised.

Raising `HttpResponseError` can be suppressed for certain status codes using the `allow_status` param - it is a list of status code values for which `HttpResponse` should be returned instead of raising `HttpResponseError`.

There is a special “*” `allow_status` value which allows any status code.

There is no need to include 100–3xx status codes in `allow_status`, because `HttpResponseError` is not raised for them.

async get(`url: Union[str, _Url]`, *, `headers: Optional[Union[Dict[str, str], HttpRequestHeaders]] = None`, `allow_status: Optional[List[Union[str, int]]] = None`) → `HttpResponse`

Similar to `request()` but performing a GET request.

async post(`url: Union[str, _Url]`, *, `headers: Optional[Union[Dict[str, str], HttpRequestHeaders]] = None`, `body: Optional[Union[bytes, HttpRequestBody]] = None`, `allow_status: Optional[List[Union[str, int]]] = None`) → `HttpResponse`

Similar to `request()` but performing a POST request.

async execute(`request: HttpRequest`, *, `allow_status: Optional[List[Union[str, int]]] = None`) → `HttpResponse`

Execute the specified `HttpRequest` instance using the request implementation configured in the `HttpClient` instance.

`HttpRequestError` is raised for *connection errors, connection and read timeouts*, etc.

`HttpResponse` instance is returned for successful responses in the 100–3xx status code range.

Otherwise, an exception of type `HttpResponseError` is raised.

Raising `HttpResponseError` can be suppressed for certain status codes using the `allow_status` param - it is a list of status code values for which `HttpResponse` should be returned instead of raising `HttpResponseError`.

There is a special "*" `allow_status` value which allows any status code.

There is no need to include 100-3xx status codes in `allow_status`, because `HttpResponseError` is not raised for them.

```
async batch_execute(*requests: HttpRequest, return_exceptions: bool = False, allow_status:
    Optional[List[Union[str, int]]] = None) → List[Union[HttpResponse, Exception]]
```

Similar to `execute()` but accepts a collection of `HttpRequest` instances that would be batch executed.

The order of the `HttpResponses` would correspond to the order of `HttpRequest` passed.

If any of the `HttpRequest` raises an exception upon execution, the exception is raised.

To prevent this, the actual exception can be returned alongside any successful `HttpResponse`. This enables salvaging any usable responses despite any possible failures. This can be done by setting `True` to the `return_exceptions` parameter.

Like `execute()`, `HttpResponseError` will be raised for responses with status codes in the 400-5xx range. The `allow_status` parameter could be used the same way here to prevent these exceptions from being raised.

You can omit `allow_status="*"` if you're passing `return_exceptions=True`. However, it would be returning `HttpResponseError` instead of `HttpResponse`.

Lastly, a `HttpRequestError` may be raised on cases like *connection errors, connection and read timeouts*, etc.

7.2 Pages

`class web_poet.pages.Injectable`

Base Page Object class, which all Page Objects should inherit from (probably through `Injectable` subclasses).

Frameworks which are using web-poet Page Objects should use `is_injectable()` function to detect if an object is an `Injectable`, and if an object is injectable, allow building it automatically through dependency injection, using <https://github.com/scrapinghub/andi> library.

Instead of inheriting you can also use `Injectable.register(MyWebPage)`. `Injectable.register` can also be used as a decorator.

```
web_poet.pages.is_injectable(cls: Any) → bool
```

Return `True` if `cls` is a class which inherits from `Injectable`.

`class web_poet.pages.ItemPage`

Bases: `Injectable`, `ABC`

Base Page Object with a required `to_item()` method. Make sure you're creating Page Objects with `to_item` methods if their main goal is to extract a single data record from a web page.

`abstract to_item()`

Extract an item from a web page

```
class web_poet.pages.WebPage(response: HttpResponse)
```

Bases: `Injectable`, `ResponseShortcutsMixin`

Base Page Object which requires `HttpResponse` and provides XPath / CSS shortcuts.

Use this class as a base class for Page Objects which work on HTML downloaded using an HTTP client directly.

response: *HttpResponse*

property base_url: *str*

Return the base url of the given response

css(*query*) → *SelectorList*

A shortcut to `.selector.css()`.

property html

Shortcut to HTML Response's content.

property selector: *Selector*

Cached instance of `parse.selector.Selector`.

property url

Shortcut to HTML Response's URL, as a string.

urljoin(*url: str*) → *str*

Convert url to absolute, taking in account url and baseurl of the response

xpath(*query, **kwargs*) → *SelectorList*

A shortcut to `.selector.xpath()`.

class `web_poet.pages.ItemWebPage`(*response: HttpResponse*)

Bases: *WebPage, ItemPage*

WebPage that requires the `to_item()` method to be implemented.

7.3 Mixins

class `web_poet.mixins.ResponseShortcutsMixin`

Common shortcut methods for working with HTML responses. This mixin could be used with Page Object base classes.

It requires "response" attribute to be present.

property url

Shortcut to HTML Response's URL, as a string.

property html

Shortcut to HTML Response's content.

property base_url: *str*

Return the base url of the given response

urljoin(*url: str*) → *str*

Convert url to absolute, taking in account url and baseurl of the response

7.4 Requests

```
web_poet.requests.request_downloader_var: ContextVar = <ContextVar
name='request_downloader'>
```

Frameworks that wants to support additional requests in `web-poet` should set the appropriate implementation of `request_downloader_var` for requesting data.

7.5 Exceptions

7.5.1 Core Exceptions

These exceptions are tied to how `web-poet` operates.

exception `web_poet.exceptions.core.RequestDownloaderVarError`

The `web_poet.request_downloader_var` had its contents accessed but there wasn't any value set during the time requests are executed.

See the documentation section about *setting up the contextvars* to learn more about this.

exception `web_poet.exceptions.core.Retry`

The page object found that the input data is partial or empty, and a request retry may provide better input.

See *Retries*.

7.5.2 HTTP Exceptions

These are exceptions pertaining to common issues faced when executing HTTP operations.

exception `web_poet.exceptions.http.HttpError` (*msg: Optional[str] = None, request: Optional[HttpRequest] = None*)

Bases: `OSError`

Indicates that an exception has occurred when handling an HTTP operation.

This is used as a **base class** for more specific errors and could be vague since it could denote problems either in the HTTP Request or Response.

For more specific errors, it would be better to use `HttpRequestError` and `HttpResponseError`.

Parameters

request (`HttpRequest`) – The `HttpRequest` instance that was used.

exception `web_poet.exceptions.http.HttpRequestError` (*msg: Optional[str] = None, request: Optional[HttpRequest] = None*)

Bases: `HttpError`

Indicates that an exception has occurred when the **HTTP Request** was being handled.

Parameters

request (`HttpRequest`) – The `HttpRequest` instance that was used.

exception `web_poet.exceptions.http.HttpResponseError` (*msg: Optional[str] = None, response: Optional[HttpResponse] = None, request: Optional[HttpRequest] = None*)

Bases: [HttpError](#)

Indicates that an exception has occurred when the **HTTP Response** was received.

For responses that are in the status code **100-3xx** range, this exception shouldn't be raised at all. However, for responses in the **400-5xx**, this will be raised by **web-poet**.

Note: Frameworks implementing **web-poet** should **NOT** raise this exception.

This exception is raised by web-poet itself, based on `allow_status` parameter found in the methods of [HttpClient](#).

Parameters

- **request** ([HttpRequest](#)) – The [HttpRequest](#) instance that was used.
- **response** ([HttpResponse](#)) – The [HttpResponse](#) instance that was received, if any. Note that this wouldn't exist if the problem occurred when executing the [HttpRequest](#).

7.6 Overrides

See the tutorial section on [Overrides](#) for more context about its use cases and some examples.

`web_poet.handle_urls(include: Union[str, Iterable[str]], *, overrides: Callable, exclude: Optional[Union[str, Iterable[str]]] = None, priority: int = 500, **kwargs)`

Class decorator that indicates that the decorated Page Object should be used instead of the overridden one for a particular set the URLs.

The Page Object that is **overridden** is declared using the `overrides` parameter.

The **override** mechanism only works on certain URLs that match the `include` and `exclude` parameters. See the documentation of the [url-matcher](#) package for more information about them.

Any extra parameters are stored as meta information that can be later used.

Parameters

- **include** – The URLs that should be handled by the decorated Page Object.
- **overrides** – The Page Object that should be *replaced*.
- **exclude** – The URLs over which the override should **not** happen.
- **priority** – The resolution priority in case of *conflicting* rules. A conflict happens when the `include`, `override`, and `exclude` parameters are the same. If so, the *highest priority* will be chosen.

```
class web_poet.overrides.OverrideRule(for_patterns: ~url_matcher.matcher.Patterns, use:
    ~typing.Callable, instead_of: ~typing.Callable, meta:
    ~typing.Dict[str, ~typing.Any] = <factory>)
```

A single override rule that specifies when a Page Object should be used in lieu of another.

This is instantiated when using the `web_poet.handle_urls()` decorator. It's also being returned as a `List[OverrideRule]` when calling the `web_poet.default_registry.get_overrides()` method.

You can access any of its attributes:

- `for_patterns` - contains the list of URL patterns associated with this rule. You can read the API documentation of the `url-matcher` package for more information about the patterns.
- `use` - The Page Object that will be **used**.
- `instead_of` - The Page Object that will be **replaced**.
- `meta` - Any other information you may want to store. This doesn't do anything for now but may be useful for future API updates.

Tip: The `OverrideRule` is also hashable. This makes it easy to store unique rules and identify any duplicates.

class `web_poet.overrides.PageObjectRegistry`

This contains the mapping rules that associates the Page Objects available for a given URL matching rule.

Note that it's simply a dict subclass with added functionalities on storing, retrieving, and searching for the `OverrideRule` instances. The **value** represents the `OverrideRule` instance from which the Page Object in the **key** is allowed to be used. Since it's essentially a dict, you can use any dict operations with it.

`web-poet` already provides a default Registry named `default_registry` for convenience. It can be directly accessed via:

```
from web_poet import handle_urls, default_registry, ItemWebPage

@handle_urls("example.com", overrides=ProductPageObject)
class ExampleComProductPage(ItemWebPage):
    ...

override_rules = default_registry.get_overrides()
```

Notice that the `@handle_urls` that we're using is a part of the `default_registry`. This provides a shorter and quicker way to interact with the built-in default `PageObjectRegistry` instead of writing the longer `@default_registry.handle_urls`.

Note: It is encouraged to simply use and import the already existing registry via `from web_poet import default_registry` instead of creating your own `PageObjectRegistry` instance. Using multiple registries would be unwieldy in most cases.

However, it might be applicable in certain scenarios like storing custom rules to separate it from the `default_registry`. This *example* from the tutorial section may provide some context.

classmethod `from_override_rules(rules: List[OverrideRule])` → `PageObjectRegistryTV`

An alternative constructor for creating a `PageObjectRegistry` instance by accepting a list of `OverrideRule`.

This is useful in cases wherein you need to store some selected rules from multiple external packages.

get_overrides() → `List[OverrideRule]`

Returns all of the `OverrideRule` that were declared using the `@handle_urls` annotation.

Warning: Remember to consider calling `consume_modules()` beforehand to recursively import all submodules which contains the `@handle_urls` annotations from external Page Objects.

search_overrides(*kwargs) → List[OverrideRule]

Returns any *OverrideRule* that has any of its attributes match the rules inside the registry.

Sample usage:

```
rules = registry.search_overrides(use=ProductPO, instead_of=GenericPO)
print(len(rules)) # 1
```

web_poet.overrides.consume_modules(*modules: str) → None

This recursively imports all packages/modules so that the `@handle_urls` annotation are properly discovered and imported.

Let's take a look at an example:

```
# FILE: my_page_obj_project/load_rules.py

from web_poet import default_registry, consume_modules

consume_modules("other_external_pkg.po", "another_pkg.lib")
rules = default_registry.get_overrides()
```

For this case, the *OverrideRule* are coming from:

- `my_page_obj_project` (since it's the same module as the file above)
- `other_external_pkg.po`
- `another_pkg.lib`
- any other modules that was imported in the same process inside the packages/modules above.

If the `default_registry` had other `@handle_urls` annotations outside of the packages/modules listed above, then the corresponding *OverrideRule* won't be returned. Unless, they were recursively imported in some way similar to `consume_modules()`.

CONTRIBUTING

web-poet is an open-source project. Your contribution is very welcome!

8.1 Issue Tracker

If you have a bug report, a new feature proposal or simply would like to make a question, please check our issue tracker on Github: <https://github.com/scrapinghub/web-poet/issues>

8.2 Source code

Our source code is hosted on Github: <https://github.com/scrapinghub/web-poet>

Before opening a pull request, it might be worth checking current and previous issues. Some code changes might also require some discussion before being accepted so it might be worth opening a new issue before implementing huge or breaking changes.

8.3 Testing

We use `tox` to run tests with different Python versions:

```
tox
```

The command above also runs type checks; we use `mypy`.

CHANGELOG

9.1 0.4.0 (2022-07-26)

- New `HttpResponse.urljoin` method, which take page's base url in account.
- New `HttpRequest.urljoin` method.
- standardized `web_poet.exceptions.Retry` exception, which allows to initiate a retry from the Page Object, e.g. based on page content.
- Documentation improvements.

9.2 0.3.0 (2022-06-14)

- Backwards Incompatible Change:
 - `web_poet.requests.request_backend_var` is renamed to `web_poet.requests.request_downloader_var`.
- Documentation and CI improvements.

9.3 0.2.0 (2022-06-10)

- Backward Incompatible Change:
 - `ResponseData` is replaced with `HttpResponse`.
`HttpResponse` exposes methods useful for web scraping (such as xpath and css selectors, json loading), and handles web page encoding detection. There are also new types like `HttpResponseBody` and `HttpResponseHeaders`.
- Added support for performing additional requests using `web_poet.HttpClient`.
- Introduced `web_poet.BrowserHtml` dependency
- Introduced `web_poet.PageParams` to pass arbitrary information inside a Page Object.
- Added `web_poet.handle_urls` decorator, which allows to declare which websites should be handled by the page objects. Lower-level `PageObjectRegistry` class is also available.
- removed support for Python 3.6
- added support for Python 3.10

9.4 0.1.1 (2021-06-02)

- `base_url` and `urljoin` shortcuts

9.5 0.1.0 (2020-07-18)

- Documentation
- `WebPage`, `ItemPage`, `ItemWebPage`, `Injectable` and `ResponseData` are available as top-level imports (e.g. `web_poet.ItemPage`)

9.6 0.0.1 (2020-04-27)

Initial release.

LICENSE

Copyright (c) Zyte Group Ltd All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. Neither the name of Zyte nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

PYTHON MODULE INDEX

W

`web_poet.exceptions.core`, 60
`web_poet.exceptions.http`, 60
`web_poet.mixins`, 59
`web_poet.overrides`, 61
`web_poet.page_inputs.browser`, 53
`web_poet.page_inputs.client`, 57
`web_poet.page_inputs.http`, 53
`web_poet.page_inputs.page_params`, 57
`web_poet.pages`, 58
`web_poet.requests`, 60

B

`base_url` (*web_poet.mixins.ResponseShortcutsMixin* property), 59
`base_url` (*web_poet.pages.WebPage* property), 59
`batch_execute()` (*web_poet.page_inputs.client.HttpClient* method), 58
`body` (*web_poet.page_inputs.http.HttpRequest* attribute), 56
`body` (*web_poet.page_inputs.http.HttpResponse* attribute), 56
`BrowserHtml` (class in *web_poet.page_inputs.browser*), 53

C

`consume_modules()` (in module *web_poet.overrides*), 63
`copy()` (*web_poet.page_inputs.http.HttpRequestHeaders* method), 54
`copy()` (*web_poet.page_inputs.http.HttpResponseHeaders* method), 55
`css()` (*web_poet.page_inputs.browser.BrowserHtml* method), 53
`css()` (*web_poet.page_inputs.http.HttpResponse* method), 56
`css()` (*web_poet.pages.WebPage* method), 59

D

`declared_encoding()` (*web_poet.page_inputs.http.HttpResponseBody* method), 53
`declared_encoding()` (*web_poet.page_inputs.http.HttpResponseHeaders* method), 55

E

`encoding` (*web_poet.page_inputs.http.HttpResponse* property), 56
`execute()` (*web_poet.page_inputs.client.HttpClient* method), 57

F

`from_bytes_dict()` (*web_poet.page_inputs.http.HttpResponseHeaders*

class method), 55
`from_name_value_pairs()` (*web_poet.page_inputs.http.HttpRequestHeaders* class method), 54
`from_name_value_pairs()` (*web_poet.page_inputs.http.HttpResponseHeaders* class method), 55
`from_override_rules()` (*web_poet.overrides.PageObjectRegistry* class method), 62

G

`get()` (*web_poet.page_inputs.client.HttpClient* method), 57
`get_overrides()` (*web_poet.overrides.PageObjectRegistry* method), 62

H

`handle_urls()` (in module *web_poet*), 61
`headers` (*web_poet.page_inputs.http.HttpRequest* attribute), 55
`headers` (*web_poet.page_inputs.http.HttpResponse* attribute), 56
`html` (*web_poet.mixins.ResponseShortcutsMixin* property), 59
`html` (*web_poet.pages.WebPage* property), 59
`HttpClient` (class in *web_poet.page_inputs.client*), 57
`HttpError`, 60
`HttpRequest` (class in *web_poet.page_inputs.http*), 55
`HttpRequestBody` (class in *web_poet.page_inputs.http*), 53
`HttpRequestError`, 60
`HttpRequestHeaders` (class in *web_poet.page_inputs.http*), 53
`HttpResponse` (class in *web_poet.page_inputs.http*), 56
`HttpResponseBody` (class in *web_poet.page_inputs.http*), 53
`HttpResponseError`, 60
`HttpResponseHeaders` (class in *web_poet.page_inputs.http*), 54

I

`Injectable` (class in `web_poet.pages`), 58
`is_injectable()` (in module `web_poet.pages`), 58
`ItemPage` (class in `web_poet.pages`), 58
`ItemWebPage` (class in `web_poet.pages`), 59

J

`json()` (`web_poet.page_inputs.http.HttpResponse` method), 56
`json()` (`web_poet.page_inputs.http.HttpResponseBody` method), 53

M

`method` (`web_poet.page_inputs.http.HttpRequest` attribute), 55
`module`
 `web_poet.exceptions.core`, 60
 `web_poet.exceptions.http`, 60
 `web_poet.mixins`, 59
 `web_poet.overrides`, 61
 `web_poet.page_inputs.browser`, 53
 `web_poet.page_inputs.client`, 57
 `web_poet.page_inputs.http`, 53
 `web_poet.page_inputs.page_params`, 57
 `web_poet.pages`, 58
 `web_poet.requests`, 60

O

`OverrideRule` (class in `web_poet.overrides`), 61

P

`PageObjectRegistry` (class in `web_poet.overrides`), 62
`PageParams` (class in `web_poet.page_inputs.page_params`), 57
`post()` (`web_poet.page_inputs.client.HttpClient` method), 57

R

`request()` (`web_poet.page_inputs.client.HttpClient` method), 57
`request_downloader_var` (in module `web_poet.requests`), 60
`RequestDownloaderVarError`, 60
`RequestUrl` (class in `web_poet.page_inputs.http`), 53
`response` (`web_poet.pages.WebPage` attribute), 59
`ResponseShortcutsMixin` (class in `web_poet.mixins`), 59
`ResponseUrl` (class in `web_poet.page_inputs.http`), 53
`Retry`, 60

S

`search_overrides()` (`web_poet.overrides.PageObjectRegistry` method), 62

`selector` (`web_poet.page_inputs.browser.BrowserHtml` property), 53
`selector` (`web_poet.page_inputs.http.HttpResponse` property), 56
`selector` (`web_poet.pages.WebPage` property), 59
`status` (`web_poet.page_inputs.http.HttpResponse` attribute), 56

T

`text` (`web_poet.page_inputs.http.HttpResponse` property), 56
`to_item()` (`web_poet.pages.ItemPage` method), 58

U

`url` (`web_poet.mixins.ResponseShortcutsMixin` property), 59
`url` (`web_poet.page_inputs.http.HttpRequest` attribute), 55
`url` (`web_poet.page_inputs.http.HttpResponse` attribute), 56
`url` (`web_poet.pages.WebPage` property), 59
`urljoin()` (`web_poet.mixins.ResponseShortcutsMixin` method), 59
`urljoin()` (`web_poet.page_inputs.http.HttpRequest` method), 56
`urljoin()` (`web_poet.page_inputs.http.HttpResponse` method), 56
`urljoin()` (`web_poet.pages.WebPage` method), 59

W

`web_poet.exceptions.core`
 module, 60
`web_poet.exceptions.http`
 module, 60
`web_poet.mixins`
 module, 59
`web_poet.overrides`
 module, 61
`web_poet.page_inputs.browser`
 module, 53
`web_poet.page_inputs.client`
 module, 57
`web_poet.page_inputs.http`
 module, 53
`web_poet.page_inputs.page_params`
 module, 57
`web_poet.pages`
 module, 58
`web_poet.requests`
 module, 60
`WebPage` (class in `web_poet.pages`), 58

X

`xpath()` (*web_poet.page_inputs.browser.BrowserHtml* method), 53

`xpath()` (*web_poet.page_inputs.http.HttpResponse* method), 56

`xpath()` (*web_poet.pages.WebPage* method), 59