

L4-Python-and-Texts

March 25, 2018

1 Regular Expression (正则表达式)

Regular expressions (called REs, or regexes, or regex patterns) are essentially a tiny, highly specialized programming language embedded inside Python and made available through the `re` module.

The regular expression language is relatively small and restricted, so not all possible string processing tasks can be done using regular expressions. There are also tasks that can be done with regular expressions, but the expressions turn out to be very complicated. In these cases, you may be better off writing Python code to do the processing; while Python code will be slower than an elaborate regular expression, it will also probably be more understandable.

1.1 Regular Expression Syntax (正则表达式语法)

Regular expressions can contain both special and ordinary characters. Most ordinary characters, like 'A', 'a', or '0', are the simplest regular expressions; they simply match themselves. You can concatenate ordinary characters, so `last` matches the string 'last'.

Some characters, like '|' or '(', are special. **Special characters** either stand for classes of ordinary characters, or affect how the regular expressions around them are interpreted.

- `[]`: Used to indicate a set of characters. In a set:
 - Characters can be listed individually, e.g. `[amk]` will match 'a', 'm', or 'k'.
 - Ranges of characters can be indicated by giving two characters and separating them by a '-', for example `[a-z]` will match any lowercase ASCII letter, `[0-5][0-9]` will match all the two-digits numbers from 00 to 59, and `[0-9A-Fa-f]` will match any hexadecimal digit. If - is escaped (e.g. `[a-z]`) or if it's placed as the first or last character (e.g. `[a-]`), it will match a literal '-'
- `(...)`: Matches whatever regular expression is inside the parentheses, and indicates the start and end of a group; the contents of a group can be retrieved after a match has been performed, and can be matched later in the string with the `\number` special sequence, described below. To match the literals '(' or ')', use `(or)`, or enclose them inside a character class: `[()]`.
- `|`: `A|B`, where A and B can be arbitrary REs, creates a regular expression that will match either A or B. An arbitrary number of REs can be separated by the '|' in this way. This can be used inside groups (see below) as well. As the target string is scanned, REs separated by '|' are tried from left to right. When one pattern completely matches, that branch is accepted. This means that once A matches, B will not be tested further, even if it would produce a

longer overall match. In other words, the '|' operator is never greedy. To match a literal '|', use '|', or enclose it inside a character class, as in [|].

The special sequences consist of " and a character from the list below. If the ordinary character is not on the list, then the resulting RE will match the second character. For example, \$ matches the character '\$'

- \d: Matches any decimal digit; this is equivalent to the class [0-9].
- \D: Matches any non-digit character; this is equivalent to the class [^0-9].
- \s: Matches any whitespace character; this is equivalent to the class [\t\n\r\f\v].
- \S: Matches any non-whitespace character; this is equivalent to the class [^\t\n\r\f\v].
- \w: Matches any alphanumeric character; this is equivalent to the class [a-zA-Z0-9_].
- \W: Matches any non-alphanumeric character; this is equivalent to the class [^a-zA-Z0-9_].

These sequences can be included inside a character class. For example, [\s,.] is a character class that will match any whitespace character, or ',' or '.'.

The final metacharacter in this section is .. It matches anything except a newline character, and there's an alternate mode (re.DOTALL) where it will match even a newline. '.' is often used where you want to match "any character".

For a complete list of sequences and expanded class definitions for Unicode string patterns, see the last part of [Regular Expression Syntax in the Standard Library](#) reference.

1.2 Python re module (Python 正则表达式模块)

This module provides regular expression matching operations similar to those found in Perl.

1.2.1 Compiling Regular Expressions (编译正则表达式)

The module defines several functions, constants, and an exception. Some of the functions are simplified versions of the full featured methods for compiled regular expressions. Most non-trivial applications always use the compiled form. Compile a regular expression pattern into a regular expression object, which can be used for matching using its `match()` and `search()` methods

```
re.compile(pattern, flags=0)
```

Regular expressions are compiled into pattern objects, which have methods for various operations such as searching for pattern matches or performing string substitutions.

```
In [3]: import re
        p = re.compile('oh my god')
        p
```

```
Out[3]: re.compile(r'oh my god', re.UNICODE)
```

```
In [5]: print(p.search('oh my god, I love Python'))
```

```
<_sre.SRE_Match object; span=(0, 9), match='oh my god'>
```

1.2.2 Backslash character (")

Regular expressions use the backslash character (") to indicate special forms or to allow special characters to be used without invoking their special meaning. This collides with Python's usage of the same character for the same purpose in string literals; for example, to match a literal backslash, one might have to write '\\\' as the pattern string, because the regular expression must be \, and each backslash must be expressed as \ inside a regular Python string literal.

```
In [6]: print('\\\\')
```

```
\\
```

The solution is to use Python's raw string notation for regular expression patterns; backslashes are not handled in any special way in a string literal prefixed with 'r'. So r"\n" is a two-character string containing " and 'n', while \n is a one-character string containing a newline. Usually patterns will be expressed in Python code using this raw string notation.

```
In [2]: print('\n') # print a new line
```

```
In [3]: print(r'\n') # print '\n' string
```

```
\n
```

1.3 Matching Characters (字符匹配)

1.4 re.match() and re.search()

Python offers two different primitive operations based on regular expressions: `re.match()` checks for a match only at the beginning of the string, while `re.search()` checks for a match anywhere in the string (this is what Perl does by default).

```
In [8]: import re
        out1 = re.match('c', "china, I love coding")
        out2 = re.search('c', "I love coding")

        print(out1)
        print(out2)
```

```
<_sre.SRE_Match object; span=(0, 1), match='c'>
<_sre.SRE_Match object; span=(7, 8), match='c'>
```

Regular expressions beginning with '^' can be used with `search()` to **restrict the match at the beginning of the string**:

```
In [10]: print(re.match("c", "abcdef")) # No match
         print(re.search("^c", "abcdef")) # No match
         print(re.search("^a", "abcdef")) # Match
```

None

None

```
<_sre.SRE_Match object; span=(0, 1), match='a'>
```

In MULTILINE mode `match()` only matches at the beginning of the string, whereas using `search()` with a regular expression beginning with `^` will match at the beginning of each line.

```
In [12]: print('A\nB\nX')
         print(re.match('X', 'A\nB\nX', re.MULTILINE)) # No match
         print(re.search('^X', 'A\nB\nX', re.MULTILINE)) # Match
```

A

B

X

None

```
<_sre.SRE_Match object; span=(4, 5), match='X'>
```

1.5 `match.group([group1, ...])`

Returns one or more subgroups of the match. If there is a single argument, the result is a single string; if there are multiple arguments, the result is a tuple with one item per argument. Without arguments, `group1` defaults to zero (the whole match is returned). If a `groupN` argument is zero, the corresponding return value is the entire matching string; if it is in the inclusive range `[1..99]`, it is the string matching the corresponding parenthesized group. If a group number is negative or larger than the number of groups defined in the pattern, an `IndexError` exception is raised. If a group is contained in a part of the pattern that did not match, the corresponding result is `None`. If a group is contained in a part of the pattern that matched multiple times, the last match is returned.

```
In [13]: m = re.match(r"(\w+) (\w+)", "Isaac Newton, physicist")
         print(m.group(0)) # The entire match

         print(m.group(1)) # The first parenthesized subgroup.

         print(m.group(2)) # The second parenthesized subgroup.

         print(m.group(1, 2)) # Multiple arguments give us a tuple.
```

Isaac Newton

Isaac

Newton

```
('Isaac', 'Newton')
```

1.6 `match.start([group])` and `match.end([group])`

Return the indices of the start and end of the substring matched by group; group defaults to zero (meaning the whole matched substring). Return -1 if group exists but did not contribute to the match. For a match object `m`, and a group `g` that did contribute to the match, the substring matched by group `g` (equivalent to `m.group(g)`) is

```
In [14]: email = "tony@tiremove_thisger.net"
        m = re.search("remove_this", email)
        print(m)
        email[:m.start()] + email[m.end():]

<_sre.SRE_Match object; span=(7, 18), match='remove_this'>
```

```
Out[14]: 'tony@tiger.net'
```

1.7 Splitting Strings (字符串拆分)

The `split()` method of a pattern splits a string apart wherever the RE matches, returning a list of the pieces. It's similar to the `split()` method of strings but provides much more generality in the delimiters that you can split by; string `split()` only supports splitting by whitespace or by a fixed string.

```
In [15]: re.split('\W+', 'Words, words, words.')
Out[15]: ['Words', 'words', 'words', '']

In [16]: re.split('(\W+)', 'Words, words, words.')
Out[16]: ['Words', ', ', ', ', 'words', ', ', ', ', 'words', '. ', '']

In [17]: re.split('\W+', 'Words, words, words.', 1)
Out[17]: ['Words', 'words, words.']

In [18]: re.split('[a-f]+', '0a3B9', flags=re.IGNORECASE)
Out[18]: ['0', '3', '9']
```

1.8 Substitution (字符替换)

```
re.sub(pattern, repl, string, count=0, flags=0)
```

Return the string obtained by replacing the leftmost non-overlapping occurrences of `pattern` in `string` by the replacement `repl`. If the `pattern` isn't found, `string` is returned unchanged. `repl` can be a string or a function; if it is a string, any backslash escapes in it are processed. That is, `\n` is converted to a single newline character, `\r` is converted to a carriage return, and so forth. Unknown escapes such as `\j` are left alone. Backreferences, such as `\6`, are replaced with the substring matched by group 6 in the pattern. For example:

```
In [19]: re.sub(r'def\s+([a-zA-Z_][a-zA-Z_0-9]*)\s*(\s*):',
               r'static PyObject*\numpy_\1(void)\n{',
               'def myfunc():')
Out[19]: 'static PyObject*\numpy_myfunc(void)\n{'
```