
TinyDB Documentation

Release 3.15.2

Markus Siemens

Jan 02, 2020

Contents

1	User's Guide	3
1.1	Introduction	3
1.2	Getting Started	4
1.3	Advanced Usage	5
2	Extending TinyDB	17
2.1	How to Extend TinyDB	17
2.2	Extensions	21
3	API Reference	23
3.1	API Documentation	23
4	Additional Notes	33
4.1	Contribution Guidelines	33
4.2	Changelog	34
4.3	Upgrading to Newer Releases	40
	Python Module Index	43
	Index	45

Welcome to TinyDB, your tiny, document oriented database optimized for your happiness :)

```
>>> from tinydb import TinyDB, Query
>>> db = TinyDB('path/to/db.json')
>>> User = Query()
>>> db.insert({'name': 'John', 'age': 22})
>>> db.search(User.name == 'John')
[{'name': 'John', 'age': 22}]
```


1.1 Introduction

Great that you've taken time to check out the TinyDB docs! Before we begin looking at TinyDB itself, let's take some time to see whether you should use TinyDB.

1.1.1 Why Use TinyDB?

- **tiny:** The current source code has 1200 lines of code (with about 40% documentation) and 1000 lines tests. For comparison: [Buzhug](#) has about 2500 lines of code (w/o tests), [CodernityDB](#) has about 7000 lines of code (w/o tests).
- **document oriented:** Like [MongoDB](#), you can store any document (represented as `dict`) in TinyDB.
- **optimized for your happiness:** TinyDB is designed to be simple and fun to use by providing a simple and clean API.
- **written in pure Python:** TinyDB neither needs an external server (as e.g. [PyMongo](#)) nor any dependencies from PyPI.
- **works on Python 2.6 + 2.7 and 3.3 – 3.7 and PyPy:** TinyDB works on all modern versions of Python and PyPy.
- **powerfully extensible:** You can easily extend TinyDB by writing new storages or modify the behaviour of storages with Middlewares.
- **100% test coverage:** No explanation needed.

In short: If you need a simple database with a clean API that just works without lots of configuration, TinyDB might be the right choice for you.

1.1.2 Why Not Use TinyDB?

- **You need advanced features like:**

- access from multiple processes or threads,
 - creating indexes for tables,
 - an HTTP server,
 - managing relationships between tables or similar,
 - [ACID guarantees](#).
- You are really concerned about **performance** and need a high speed database.

To put it plainly: If you need advanced features or high performance, TinyDB is the wrong database for you – consider using databases like [SQLite](#), [Buzhug](#), [CodernityDB](#) or [MongoDB](#).

1.2 Getting Started

1.2.1 Installing TinyDB

To install TinyDB from PyPI, run:

```
$ pip install tinydb
```

You can also grab the latest development version from [GitHub](#). After downloading and unpacking it, you can install it using:

```
$ python setup.py install
```

1.2.2 Basic Usage

Let's cover the basics before going more into detail. We'll start by setting up a TinyDB database:

```
>>> from tinydb import TinyDB, Query
>>> db = TinyDB('db.json')
```

You now have a TinyDB database that stores its data in `db.json`. What about inserting some data? TinyDB expects the data to be Python dicts:

```
>>> db.insert({'type': 'apple', 'count': 7})
>>> db.insert({'type': 'peach', 'count': 3})
```

Note: The `insert` method returns the inserted document's ID. Read more about it here: [Using Document IDs](#).

Now you can get all documents stored in the database by running:

```
>>> db.all()
[{'count': 7, 'type': 'apple'}, {'count': 3, 'type': 'peach'}]
```

You can also iter over stored documents:

```
>>> for item in db:
>>>     print(item)
{'count': 7, 'type': 'apple'}
{'count': 3, 'type': 'peach'}
```


Of course you'll also want to search for specific documents. Let's try:

```
>>> Fruit = Query()
>>> db.search(Fruit.type == 'peach')
[{'count': 3, 'type': 'peach'}]
>>> db.search(Fruit.count > 5)
[{'count': 7, 'type': 'apple'}]
```

Next we'll update the count field of the apples:

```
>>> db.update({'count': 10}, Fruit.type == 'apple')
>>> db.all()
[{'count': 10, 'type': 'apple'}, {'count': 3, 'type': 'peach'}]
```

In the same manner you can also remove documents:

```
>>> db.remove(Fruit.count < 5)
>>> db.all()
[{'count': 10, 'type': 'apple'}]
```

And of course you can throw away all data to start with an empty database:

```
>>> db.purge()
>>> db.all()
[]
```

Recap

Before we dive deeper, let's recapitulate the basics:

Inserting	
<code>db.insert(...)</code>	Insert a document
Getting data	
<code>db.all()</code>	Get all documents
<code>iter(db)</code>	Iter over all documents
<code>db.search(query)</code>	Get a list of documents matching the query
Updating	
<code>db.update(fields, query)</code>	Update all documents matching the query to contain fields
Removing	
<code>db.remove(query)</code>	Remove all documents matching the query
<code>db.purge()</code>	Purge all documents
Querying	
<code>Query()</code>	Create a new query object
<code>Query().field == 2</code>	Match any document that has a key field with value == 2 (also possible: != > >= < <=)

1.3 Advanced Usage

1.3.1 Remarks on Storage

Before we dive deeper into the usage of TinyDB, we should stop for a moment and discuss how TinyDB stores data.

To convert your data to a format that is writable to disk TinyDB uses the [Python JSON](#) module by default. It's great when only simple data types are involved but it cannot handle more complex data types like custom classes. On Python 2 it also converts strings to Unicode strings upon reading (described [here](#)).

If that causes problems, you can write *your own storage*, that uses a more powerful (but also slower) library like [pickle](#) or [PyYAML](#).

Hint: Opening multiple TinyDB instances on the same data (e.g. with the `JSONStorage`) may result in unexpected behavior due to query caching. See [query_caching](#) on how to disable the query cache.

1.3.2 Queries

With that out of the way, let's start with TinyDB's rich set of queries. There are two main ways to construct queries. The first one resembles the syntax of popular ORM tools:

```
>>> from tinydb import Query
>>> User = Query()
>>> db.search(User.name == 'John')
```

As you can see, we first create a new Query object and then use it to specify which fields to check. Searching for nested fields is just as easy:

```
>>> db.search(User.birthday.year == 1990)
```

Not all fields can be accessed this way if the field name is not a valid Python identifier. In this case, you can switch to array indexing notation:

```
>>> # This would be invalid Python syntax:
>>> db.search(User.country-code == 'foo')
>>> # Use this instead:
>>> db.search(User['country-code'] == 'foo')
```

The second, traditional way of constructing queries is as follows:

```
>>> from tinydb import where
>>> db.search(where('field') == 'value')
```

Using `where('field')` is a shorthand for the following code:

```
>>> db.search(Query() ['field'] == 'value')
```

Accessing nested fields with this syntax can be achieved like this:

```
>>> db.search(where('birthday').year == 1900)
>>> db.search(where('birthday') ['year'] == 1900)
```

Advanced queries

In the [Getting Started](#) you've learned about the basic comparisons (`==`, `<`, `>`, `...`). In addition to these TinyDB supports the following queries:

```
>>> # Existence of a field:
>>> db.search(User.name.exists())
```

```
>>> # Regex:
>>> # Full item has to match the regex:
>>> db.search(User.name.matches('[aZ]*'))
>>> # Case insensitive search for 'John':
>>> import re
>>> db.search(User.name.matches('John', flags=re.IGNORECASE))
>>> # Any part of the item has to match the regex:
>>> db.search(User.name.search('b+'))
```

```
>>> # Custom test:
>>> test_func = lambda s: s == 'John'
>>> db.search(User.name.test(test_func))
```

```
>>> # Custom test with parameters:
>>> def test_func(val, m, n):
>>>     return m <= val <= n
>>> db.search(User.age.test(test_func, 0, 21))
>>> db.search(User.age.test(test_func, 21, 99))
```

When a field contains a list, you also can use the `any` and `all` methods. There are two ways to use them: with lists of values and with nested queries. Let's start with the first one. Assuming we have a user object with a groups list like this:

```
>>> db.insert({'name': 'user1', 'groups': ['user']})
>>> db.insert({'name': 'user2', 'groups': ['admin', 'user']})
>>> db.insert({'name': 'user3', 'groups': ['sudo', 'user']})
```

Now we can use the following queries:

```
>>> # User's groups include at least one value from ['admin', 'sudo']
>>> db.search(User.groups.any(['admin', 'sudo']))
[{'name': 'user2', 'groups': ['admin', 'user']},
 {'name': 'user3', 'groups': ['sudo', 'user']}]
>>>
>>> # User's groups include all values from ['admin', 'user']
>>> db.search(User.groups.all(['admin', 'user']))
[{'name': 'user2', 'groups': ['admin', 'user']}]
```

In some cases you may want to have more complex `any/all` queries. This is where nested queries come in as helpful. Let's set up a table like this:

```
>>> Group = Query()
>>> Permission = Query()
>>> groups = db.table('groups')
>>> groups.insert({
>>>     'name': 'user',
>>>     'permissions': [{'type': 'read'}]})
>>> groups.insert({
>>>     'name': 'sudo',
>>>     'permissions': [{'type': 'read'}, {'type': 'sudo'}]})
>>> groups.insert({
>>>     'name': 'admin',
>>>     'permissions': [{'type': 'read'}, {'type': 'write'}, {'type': 'sudo'}]})
```

Now let's search this table using nested `any/all` queries:

```
>>> # Group has a permission with type 'read'
>>> groups.search(Group.permissions.any(Permission.type == 'read'))
[{'name': 'user', 'permissions': [{'type': 'read'}]},
 {'name': 'sudo', 'permissions': [{'type': 'read'}, {'type': 'sudo'}]},
 {'name': 'admin', 'permissions':
   [{'type': 'read'}, {'type': 'write'}, {'type': 'sudo'}]}]
>>> # Group has ONLY permission 'read'
>>> groups.search(Group.permissions.all(Permission.type == 'read'))
[{'name': 'user', 'permissions': [{'type': 'read'}]}]
```

As you can see, `any` tests if there is *at least one* document matching the query while `all` ensures *all* documents match the query.

The opposite operation, checking if a single item is contained in a list, is also possible using `one_of`:

```
>>> db.search(User.name.one_of(['jane', 'john']))
```

Query modifiers

TinyDB also allows you to use logical operations to modify and combine queries:

```
>>> # Negate a query:
>>> db.search(~ (User.name == 'John'))
```

```
>>> # Logical AND:
>>> db.search((User.name == 'John') & (User.age <= 30))
```

```
>>> # Logical OR:
>>> db.search((User.name == 'John') | (User.name == 'Bob'))
```

Note: When using `&` or `|`, make sure you wrap the conditions on both sides with parentheses or Python will mess up the comparison.

Also, when using negation (`~`) you'll have to wrap the query you want to negate in parentheses.

The reason for these requirements is that Python's binary operators that are used for query modifiers have a higher operator precedence than comparison operators. Simply put, `~ User.name == 'John'` is parsed by Python as `(~User.name) == 'John'` instead of `~(User.name == 'John')`. See also the [Python docs on operator precedence](#) for details.

Recap

Let's review the query operations we've learned:

Queries	
<code>Query().field.exists()</code>	Match any document where a field called <code>field</code> exists
<code>Query().field.matches(regex)</code>	Match any document with the whole field matching the regular expression
<code>Query().field.search(regex)</code>	Match any document with a substring of the field matching the regular expression
<code>Query().field.test(func, *args)</code>	Matches any document for which the function returns <code>True</code>
<code>Query().field.all(query list)</code>	If given a query, matches all documents where all documents in the list <code>field</code> match the query. If given a list, matches all documents where all documents in the list <code>field</code> are a member of the given list
<code>Query().field.any(query list)</code>	If given a query, matches all documents where at least one document in the list <code>field</code> match the query. If given a list, matches all documents where at least one documents in the list <code>field</code> are a member of the given list
<code>Query().field.one_of(list)</code>	Match if the field is contained in the list
Logical operations on queries	
<code>~ (query)</code>	Match documents that don't match the query
<code>(query1) & (query2)</code>	Match documents that match both queries
<code>(query1) (query2)</code>	Match documents that match at least one of the queries

1.3.3 Handling Data

Next, let's look at some more ways to insert, update and retrieve data from your database.

Inserting data

As already described you can insert a document using `db.insert(...)`. In case you want to insert multiple documents, you can use `db.insert_multiple(...)`:

```
>>> db.insert_multiple([
    {'name': 'John', 'age': 22},
    {'name': 'John', 'age': 37}])
>>> db.insert_multiple({'int': 1, 'value': i} for i in range(2))
```

Updating data

Sometimes you want to update all documents in your database. In this case, you can leave out the `query` argument:

```
>>> db.update({'foo': 'bar'})
```

When passing a dict to `db.update(fields, query)`, it only allows you to update a document by adding or overwriting its values. But sometimes you may need to e.g. remove one field or increment its value. In that case you can pass a function instead of `fields`:

```
>>> from tinydb.operations import delete
>>> db.update(delete('key1'), User.name == 'John')
```

This will remove the key `key1` from all matching documents. TinyDB comes with these operations:

- `delete(key)`: delete a key from the document
- `increment(key)`: increment the value of a key
- `decrement(key)`: decrement the value of a key
- `add(key, value)`: add value to the value of a key (also works for strings)
- `subtract(key, value)`: subtract value from the value of a key
- `set(key, value)`: set key to value

Of course you also can write your own operations:

```
>>> def your_operation(your_arguments):
...     def transform(doc):
...         # do something with the document
...         # ...
...     return transform
...
>>> db.update(your_operation(arguments), query)
```

1.3.4 Data access and modification

Upserting data

In some cases you'll need a mix of both `update` and `insert`: `upsert`. This operation is provided a document and a query. If it finds any documents matching the query, they will be updated with the data from the provided document. On the other hand, if no matching document is found, it inserts the provided document into the table:

```
>>> db.upsert({'name': 'John', 'logged-in': True}, User.name == 'John')
```

This will update all users with the name `John` to have `logged-in` set to `True`. If no matching user is found, a new document is inserted with both the name set and the `logged-in` flag.

Retrieving data

There are several ways to retrieve data from your database. For instance you can get the number of stored documents:

```
>>> len(db)
3
```

Hint: This will return the number of documents in the default table (see the notes on the [default table](#)).

Then of course you can use `db.search(...)` as described in the *Getting Started* section. But sometimes you want to get only one matching document. Instead of using

```
>>> try:
...     result = db.search(User.name == 'John')[0]
... except IndexError:
...     pass
```

you can use `db.get(...)`:

```
>>> db.get(User.name == 'John')
{'name': 'John', 'age': 22}
>>> db.get(User.name == 'Bobby')
None
```

Caution: If multiple documents match the query, probably a random one of them will be returned!

Often you don't want to search for documents but only know whether they are stored in the database. In this case `db.contains(...)` is your friend:

```
>>> db.contains(User.name == 'John')
```

In a similar manner you can look up the number of documents matching a query:

```
>>> db.count(User.name == 'John')
2
```

Replacing data

Another occasionally useful operation is to replace a list of documents. If you have a list of documents with IDs (see *document_ids*), you can pass them to `db.write_back(list)`:

```
>>> docs = db.search(User.name == 'John')
[{'name': 'John', 'age': 12}, {'name': 'John', 'age': 44}]
>>> for doc in docs:
...     doc['name'] = 'Jane'
>>> db.write_back(docs) # Will update the documents we retrieved
>>> docs = db.search(User.name == 'John')
[]
>>> docs = db.search(User.name == 'Jane')
[{'name': 'Jane', 'age': 12}, {'name': 'Jane', 'age': 44}]
```

Alternatively you can pass a list of documents along with a list of document IDs to achieve the same goal. In this case, the length of the document list and the ID list has to be equal.

Recap

Let's summarize the ways to handle data:

Inserting data	
<code>db.insert_multiple(...)</code>	Insert multiple documents
Updating data	
<code>db.update(operation, ...)</code>	Update all matching documents with a special operation
<code>db.write_back(docs)</code>	Replace all documents with the updated versions
Retrieving data	
<code>len(db)</code>	Get the number of documents in the database
<code>db.get(query)</code>	Get one document matching the query
<code>db.contains(query)</code>	Check if the database contains a matching document
<code>db.count(query)</code>	Get the number of matching documents

Note: This was a new feature in v3.6.0

1.3.5 Using Document IDs

Internally TinyDB associates an ID with every document you insert. It's returned after inserting a document:

```
>>> db.insert({'name': 'John', 'age': 22})
3
>>> db.insert_multiple([{'...'}, {'...'}, {'...'}])
[4, 5, 6]
```

In addition you can get the ID of already inserted documents using `document.doc_id`. This works both with `get` and `all`:

```
>>> e1 = db.get(User.name == 'John')
>>> e1.doc_id
3
>>> e1 = db.all()[0]
>>> e1.doc_id
12
```

Different TinyDB methods also work with IDs, namely: `update`, `remove`, `contains` and `get`. The first two also return a list of affected IDs.

```
>>> db.update({'value': 2}, doc_ids=[1, 2])
>>> db.contains(doc_ids=[1])
True
>>> db.remove(doc_ids=[1, 2])
>>> db.get(doc_id=3)
{'...'}
```

Using `doc_id` instead of `Query()` again is slightly faster in operation.

Recap

Let's sum up the way TinyDB supports working with IDs:

1.3.6 Tables

TinyDB supports working with multiple tables. They behave just the same as the `TinyDB` class. To create and use a table, use `db.table(name)`.

```
>>> table = db.table('table_name')
>>> table.insert({'value': True})
>>> table.all()
[{'value': True}]
>>> for row in table:
>>>     print(row)
{'value': True}
```

To remove a table from a database, use:

```
>>> db.purge_table('table_name')
```

If on the other hand you want to remove all tables, use the counterpart:

```
>>> db.purge_tables()
```

Finally, you can get a list with the names of all tables in your database:

```
>>> db.tables()
['_default', 'table_name']
```

Default Table

TinyDB uses a table named `_default` as the default table. All operations on the database object (like `db.insert(...)`) operate on this table. The name of this table can be modified by either passing `default_table` to the `TinyDB` constructor or by setting the `DEFAULT_TABLE` class variable to modify the default table name for all instances:

```
>>> #1: for a single instance only
>>> TinyDB(storage=SomeStorage, default_table='my-default')
>>> #2: for all instances
>>> TinyDB.DEFAULT_TABLE = 'my-default'
```

You also can modify the keyword arguments that are passed to the default table by setting `TinyDB.DEFAULT_TABLE_KWARGS`. For example, you can disable the query cache for the default table by setting like this:

```
>>> TinyDB.DEFAULT_TABLE_KWARGS = {'cache_size': 0}
```

Query Caching

TinyDB caches query result for performance. You can optimize the query cache size by passing the `cache_size` to the `table(...)` function:

```
>>> table = db.table('table_name', cache_size=30)
```

Hint: You can set `cache_size` to `None` to make the cache unlimited in size. Also, you can set `cache_size` to 0 to disable it.

1.3.7 Storage & Middleware

Storage Types

TinyDB comes with two storage types: JSON and in-memory. By default TinyDB stores its data in JSON files so you have to specify the path where to store it:

```
>>> from tinydb import TinyDB, where
>>> db = TinyDB('path/to/db.json')
```

To use the in-memory storage, use:

```
>>> from tinydb.storages import MemoryStorage
>>> db = TinyDB(storage=MemoryStorage)
```

Hint: All arguments except for the `storage` argument are forwarded to the underlying storage. For the JSON storage you can use this to pass additional keyword arguments to Python's `json.dump(...)` method. For example, you can set it to create prettified JSON files like this:

```
>>> db = TinyDB('db.json', sort_keys=True, indent=4, separators=(',', ': '))
```

To modify the default storage for all TinyDB instances, set the `DEFAULT_STORAGE` class variable:

```
>>> TinyDB.DEFAULT_STORAGE = MemoryStorage
```

In case you need to access the storage instance directly, you can use the `storage` property of your TinyDB instance. This may be useful to call method directly on the storage or middleware:

```
>>> db = TinyDB(storage=CachingMiddleware(MemoryStorage))
<tinydb.middlewares.CachingMiddleware at 0x10991def0>
>>> db.storage.flush()
```

Middleware

Middleware wraps around existing storage allowing you to customize their behaviour.

```
>>> from tinydb.storages import JSONStorage
>>> from tinydb.middlewares import CachingMiddleware
>>> db = TinyDB('/path/to/db.json', storage=CachingMiddleware(JSONStorage))
```

Hint: You can nest middleware:

```
>>> db = TinyDB('/path/to/db.json',
                storage=FirstMiddleware(SecondMiddleware(JSONStorage)))
```

CachingMiddleware

The `CachingMiddleware` improves speed by reducing disk I/O. It caches all read operations and writes data to disk after a configured number of write operations.

To make sure that all data is safely written when closing the table, use one of these ways:

```
# Using a context manager:
with database as db:
    # Your operations
```

```
# Using the close function
db.close()
```

1.3.8 What's next

Congratulations, you've made through the user guide! Now go and build something awesome or dive deeper into TinyDB with these resources:

- Want to learn how to customize TinyDB (storages, middlewares) and what extensions exist? Check out [How to Extend TinyDB](#) and [Extensions](#).
- Want to study the API in detail? Read [API Documentation](#).
- Interested in contributing to the TinyDB development guide? Go on to the [Contribution Guidelines](#).

2.1 How to Extend TinyDB

There are three main ways to extend TinyDB and modify its behaviour:

1. custom storage,
2. custom middleware, and
3. custom table classes.

Let's look at them in this order.

2.1.1 Write Custom Storage

First, we have support for custom storage. By default TinyDB comes with an in-memory storage mechanism and a JSON file storage mechanism. But of course you can add your own. Let's look how you could add a [YAML](#) storage using [PyYAML](#):

```
import yaml

def represent_doc(dumper, data):
    # Represent `Document` objects as their dict's string representation
    # which PyYAML understands
    return dumper.represent_data(dict(data))

yaml.add_representer(Document, represent_doc)

class YAMLStorage(Storage):
    def __init__(self, filename): # (1)
        self.filename = filename

    def read(self):
        with open(self.filename) as handle:
```

(continues on next page)

(continued from previous page)

```

        try:
            data = yaml.safe_load(handle.read()) # (2)
            return data
        except yaml.YAMLError:
            return None # (3)

    def write(self, data):
        with open(self.filename, 'w') as handle:
            yaml.dump(data, handle)

    def close(self): # (4)
        pass

```

There are some things we should look closer at:

1. The constructor will receive all arguments passed to TinyDB when creating the database instance (except storage which TinyDB itself consumes). In other words calling `TinyDB('something', storage=YAMLStorage)` will pass 'something' as an argument to `YAMLStorage`.
2. We use `yaml.safe_load` as recommended by the [PyYAML documentation](#) when processing data from a potentially untrusted source.
3. If the storage is uninitialized, TinyDB expects the storage to return `None` so it can do any internal initialization that is necessary.
4. If your storage needs any cleanup (like closing file handles) before an instance is destroyed, you can put it in the `close()` method. To run these, you'll either have to run `db.close()` on your TinyDB instance or use it as a context manager, like this:

```

with TinyDB('db.yml', storage=YAMLStorage) as db:
    # ...

```

Finally, using the YAML storage is very straight-forward:

```

db = TinyDB('db.yml', storage=YAMLStorage)
# ...

```

2.1.2 Write Custom Middleware

Sometimes you don't want to write a new storage module but rather modify the behaviour of an existing one. As an example we'll build middleware that filters out any empty items.

Because middleware acts as a wrapper around a storage, they needs a `read()` and a `write(data)` method. In addition, they can access the underlying storage via `self.storage`. Before we start implementing we should look at the structure of the data that the middleware receives. Here's what the data that goes through the middleware looks like:

```

{
    '_default': {
        1: {'key': 'value'},
        2: {'key': 'value'},
        # other items
    },
    # other tables
}

```

Thus, we'll need two nested loops:

1. Process every table
2. Process every item

Now let's implement that:

```
class RemoveEmptyItemsMiddleware(Middleware):
    def __init__(self, storage_cls=TinyDB.DEFAULT_STORAGE):
        # Any middleware has to call the super constructor
        # with storage_cls
        super(CustomMiddleware, self).__init__(storage_cls)

    def read(self):
        data = self.storage.read()

        for table_name in data:
            table = data[table_name]

            for doc_id in table:
                item = table[doc_id]

                if item == {}:
                    del table[doc_id]

        return data

    def write(self, data):
        for table_name in data:
            table = data[table_name]

            for doc_id in table:
                item = table[doc_id]

                if item == {}:
                    del table[doc_id]

        self.storage.write(data)

    def close(self):
        self.storage.close()
```

Two remarks:

1. You have to use the `super(...)` call as shown in the example. To run your own initialization, add it below the `super(...)` call.
2. This is an example for middleware, not an example for clean code. Don't use it as shown here without at least refactoring the loops into a separate method.

To wrap storage with this new middleware, we use it like this:

```
db = TinyDB(storage=RemoveEmptyItemsMiddleware(SomeStorageClass))
```

Here `SomeStorageClass` should be replaced with the storage you want to use. If you leave it empty, the default storage will be used (which is the `JSONStorage`).

2.1.3 Creating a Custom Table Classes

Custom storage and middleware are useful if you want to modify the way TinyDB stores its data. But there are cases where you want to modify how TinyDB itself behaves. For that use case TinyDB supports custom table classes. Internally TinyDB creates a `Table` instance for every table that is used. You can overwrite which class is used by setting `TinyDB.table_class` before creating a TinyDB instance. This class has to support the *Table API*. The best way to accomplish that is to subclass it:

```
from tinydb import TinyDB
from tinydb.database import Table

class YourTableClass(Table):
    pass # Modify original methods as needed

TinyDB.table_class = YourTableClass
```

For an more advanced example, see the source of the `tinydb-smartcache` extension.

2.1.4 Creating a Custom Storage Proxy Classes

Warning: This extension requires knowledge of TinyDB internals. Use it if you understand how TinyDB works in detail.

Another way to modify TinyDB's behavior is to create a custom storage proxy class. Internally, TinyDB uses a proxy class to allow tables to access a storage object. The proxy makes sure the table only accesses its own table data and doesn't accidentally modify other table's data.

In this class you can modify how a table can read and write from a storage instance. Also, the proxy class has a method called `_new_document` which creates a new document object. If you want to replace it with a different document class, you can do it right here.

```
from tinydb import TinyDB
from tinydb.database import Table, StorageProxy, Document
from tinydb.storages import MemoryStorage

class YourStorageProxy(StorageProxy):
    def _new_document(self, key, val):
        # Modify document object creation
        doc_id = int(key)
        return Document(val, doc_id)

    def read(self):
        return {} # Modify reading

    def write(self, data):
        pass # Modify writing

TinyDB.storage_proxy_class = YourStorageProxy
# Or:
TinyDB(storage=..., storage_proxy_class=YourStorageProxy)
```


2.2 Extensions

Here are some extensions that might be useful to you:

2.2.1 `tinyindex`

Repo: <https://github.com/eugene-eeo/tinyindex>

Status: *experimental*

Description: Document indexing for TinyDB. Basically ensures deterministic (as long as there aren't any changes to the table) yielding of documents.

2.2.2 `tinymongo`

Repo: <https://github.com/schapman1974/tinymongo>

Status: *experimental*

Description: A simple wrapper that allows to use TinyDB as a flat file drop-in replacement for MongoDB.

2.2.3 `TinyMP`

Repo: <https://github.com/alshapton/TinyMP>

Status: *stable*

Description: A MessagePack-based storage extension to tinydb using <http://msgpack.org>

2.2.4 `tinyrecord`

Repo: <https://github.com/eugene-eeo/tinyrecord>

Status: *stable*

Description: Tinyrecord is a library which implements experimental atomic transaction support for the TinyDB NoSQL database. It uses a record-first then execute architecture which allows us to minimize the time that we are within a thread lock.

2.2.5 `tinydb-appengine`

Repo: <https://github.com/imalento/tinydb-appengine>

Status: *stable*

Description: `tinydb-appengine` provides TinyDB storage for App Engine. You can use JSON readonly.

2.2.6 `tinydb-serialization`

Repo: <https://github.com/msiemens/tinydb-serialization>

Status: *stable*

Description: `tinydb-serialization` provides serialization for objects that TinyDB otherwise couldn't handle.

2.2.7 `tinydb-smartcache`

Repo: <https://github.com/msiemens/tinydb-smartcache>

Status: *stable*

Description: `tinydb-smartcache` provides a smart query cache for TinyDB. It updates the query cache when inserting/removing/updating documents so the cache doesn't get invalidated. It's useful if you perform lots of queries while the data changes only little.

2.2.8 `aiotinydb`

Repo: <https://github.com/ASMfreaK/aiotinydb>

Status: *stable*

Description: asyncio compatibility shim for TinyDB. Enables usage of TinyDB in asyncio-aware contexts without slow synchronous IO.

3.1 API Documentation

3.1.1 `tinydb.database`

class `tinydb.database.TinyDB(*args, **kwargs)`

The main class of TinyDB.

Gives access to the database, provides methods to insert/search/remove and getting tables.

DEFAULT_STORAGE

alias of `tinydb.storages.JSONStorage`

__getattr__(*name*)

Forward all unknown attribute calls to the underlying standard table.

__init__(**args, **kwargs*)

Create a new instance of TinyDB.

All arguments and keyword arguments will be passed to the underlying storage class (default: `JSONStorage`).

Parameters

- **storage** – The class of the storage to use. Will be initialized with `args` and `kwargs`.
- **default_table** – The name of the default table to populate.

__iter__()

Iter over all documents from default table.

__len__()

Get the total number of documents in the default table.

```
>>> db = TinyDB('db.json')
>>> len(db)
0
```

`__repr__()` $\Leftrightarrow repr(x)$

`close()`

Close the database.

`purge_table(name)`

Purge a specific table from the database. **CANNOT BE REVERSED!**

Parameters `name` (*str*) – The name of the table.

`purge_tables()`

Purge all tables from the database. **CANNOT BE REVERSED!**

`storage`

Access the storage used for this TinyDB instance.

Returns This instance's storage

`storage_proxy_class`

alias of `StorageProxy`

`table(name='_default', **options)`

Get access to a specific table.

Creates a new table, if it hasn't been created before, otherwise it returns the cached `Table` object.

Parameters

- **name** (*str*) – The name of the table.
- **cache_size** – How many query results to cache.
- **table_class** – Which table class to use.

`table_class`

alias of `Table`

`tables()`

Get the names of all tables in the database.

Returns a set of table names

Return type `set[str]`

class `tinydb.database.Table(storage, name, cache_size=10)`

Represents a single TinyDB Table.

`__init__(storage, name, cache_size=10)`

Get access to a table.

Parameters

- **storage** (`StorageProxy`) – Access to the storage
- **name** – The table name
- **cache_size** – Maximum size of query cache.

`__iter__()`

Iter over all documents stored in the table.

Returns an iterator over all documents.

Return type `listiterator[Element]`

`__len__()`

Get the total number of documents in the table.

`__repr__()` $\Leftrightarrow repr(x)$

all()

Get all documents stored in the table.

Returns a list with all documents.

Return type list[Element]

clear_cache()

Clear the query cache.

A simple helper that clears the internal query cache.

contains (*cond=None, doc_ids=None, eids=None*)

Check whether the database contains a document matching a condition or an ID.

If *eids* is set, it checks if the db contains a document with one of the specified.

Parameters

- **cond** (Query) – the condition use
- **doc_ids** – the document IDs to look for

count (*cond*)

Count the documents matching a condition.

Parameters **cond** (Query) – the condition use

get (*cond=None, doc_id=None, eid=None*)

Get exactly one document specified by a query or an ID.

Returns `None` if the document doesn't exist

Parameters

- **cond** (Query) – the condition to check against
- **doc_id** – the document's ID

Returns the document or `None`

Return type Element | None

insert (*document*)

Insert a new document into the table.

Parameters **document** – the document to insert

Returns the inserted document's ID

insert_multiple (*documents*)

Insert multiple documents into the table.

Parameters **documents** – a list of documents to insert

Returns a list containing the inserted documents' IDs

name

Get the table name.

process_elements (*func, cond=None, doc_ids=None, eids=None*)

Helper function for processing all documents specified by condition or IDs.

A repeating pattern in TinyDB is to run some code on all documents that match a condition or are specified by their ID. This is implemented in this function. The function passed as *func* has to be a callable. Its

first argument will be the data currently in the database. Its second argument is the document ID of the currently processed document.

See: `update()`, `remove()`

Parameters

- **func** – the function to execute on every included document. first argument: all data second argument: the current eid
- **cond** – query that matches documents to use, or
- **doc_ids** – list of document IDs to use
- **eids** – list of document IDs to use (deprecated)

Returns the document IDs that were affected during processing

purge()

Purge the table by removing all documents.

remove (*cond=None, doc_ids=None, eids=None*)

Remove all matching documents.

Parameters

- **cond** (*query*) – the condition to check against
- **doc_ids** (*list*) – a list of document IDs

Returns a list containing the removed document's ID

search (*cond*)

Search for all documents matching a 'where' cond.

Parameters **cond** (*Query*) – the condition to check against

Returns list of matching documents

Return type list[Element]

update (*fields, cond=None, doc_ids=None, eids=None*)

Update all matching documents to have a given set of fields.

Parameters

- **fields** (*dict | dict -> None*) – the fields that the matching documents will have or a method that will update the documents
- **cond** (*query*) – which documents to update
- **doc_ids** (*list*) – a list of document IDs

Returns a list containing the updated document's ID

upsert (*document, cond*)

Update a document, if it exist - insert it otherwise.

Note: this will update *all* documents matching the query.

Parameters

- **document** – the document to insert or the fields to update
- **cond** – which document to look for

Returns a list containing the updated document's ID

write_back (*documents*, *doc_ids=None*, *eids=None*)

Write back documents by doc_id

Parameters

- **documents** – a list of document to write back
- **doc_ids** – a list of document IDs which need to be written back

Returns a list of document IDs that have been written

class `tinydb.database.Document` (*value*, *doc_id*, ***kwargs*)

Represents a document stored in the database.

This is a transparent proxy for database records. It exists to provide a way to access a record's id via `el.doc_id`.

doc_id

The document's id

__init__ (*value*, *doc_id*, ***kwargs*)

`x.__init__(...)` initializes `x`; see `help(type(x))` for signature

`tinydb.database.Element`

alias of `tinydb.database.Document`

3.1.2 tinydb.queries

class `tinydb.queries.Query`

TinyDB Queries.

Allows to build queries for TinyDB databases. There are two main ways of using queries:

1) ORM-like usage:

```
>>> User = Query()
>>> db.search(User.name == 'John Doe')
>>> db.search(User['logged-in'] == True)
```

2) Classical usage:

```
>>> db.search(where('value') == True)
```

Note that `where(...)` is a shorthand for `Query(...)` allowing for a more fluent syntax.

Besides the methods documented here you can combine queries using the binary AND and OR operators:

```
>>> # Binary AND:
>>> db.search((where('field1').exists()) & (where('field2') == 5))
>>> # Binary OR:
>>> db.search((where('field1').exists()) | (where('field2') == 5))
```

Queries are executed by calling the resulting object. They expect to get the document to test as the first argument and return `True` or `False` depending on whether the documents matches the query or not.

__eq__ (*rhs*)

Test a dict value for equality.

```
>>> Query().f1 == 42
```

Parameters *rhs* – The value to compare against

`__ge__ (rhs)`

Test a dict value for being greater than or equal to another value.

```
>>> Query().f1 >= 42
```

Parameters *rhs* – The value to compare against

`__gt__ (rhs)`

Test a dict value for being greater than another value.

```
>>> Query().f1 > 42
```

Parameters *rhs* – The value to compare against

`__hash__ ()` $\leq \Rightarrow hash(x)$

`__init__ ()`

`x.__init__(...)` initializes `x`; see `help(type(x))` for signature

`__le__ (rhs)`

Test a dict value for being lower than or equal to another value.

```
>>> where('f1') <= 42
```

Parameters *rhs* – The value to compare against

`__lt__ (rhs)`

Test a dict value for being lower than another value.

```
>>> Query().f1 < 42
```

Parameters *rhs* – The value to compare against

`__ne__ (rhs)`

Test a dict value for inequality.

```
>>> Query().f1 != 42
```

Parameters *rhs* – The value to compare against

`__repr__ ()` $\leq \Rightarrow repr(x)$

`all (cond)`

Check if a condition is met by all documents in a list, where a condition can also be a sequence (e.g. list).

```
>>> Query().f1.all(Query().f2 == 1)
```

Matches:

```
{ 'f1': [{ 'f2': 1 }, { 'f2': 1 }] }
```



```
>>> Query().f1.all([1, 2, 3])
```

Matches:

```
{'f1': [1, 2, 3, 4, 5]}
```

Parameters **cond** – Either a query that all documents have to match or a list which has to be contained in the tested document.

any (*cond*)

Check if a condition is met by any document in a list, where a condition can also be a sequence (e.g. list).

```
>>> Query().f1.any(Query().f2 == 1)
```

Matches:

```
{'f1': [{'f2': 1}, {'f2': 0}]}
```

```
>>> Query().f1.any([1, 2, 3])
```

Matches:

```
{'f1': [1, 2]}
{'f1': [3, 4, 5]}
```

Parameters **cond** – Either a query that at least one document has to match or a list of which at least one document has to be contained in the tested document.

exists ()

Test for a dict where a provided key exists.

```
>>> Query().f1.exists()
```

matches (*regex, flags=0*)

Run a regex test against a dict value (whole string has to match).

```
>>> Query().f1.matches(r'^\w+$')
```

Parameters

- **regex** – The regular expression to use for matching
- **flags** – regex flags to pass to `re.match`

one_of (*items*)

Check if the value is contained in a list or generator.

```
>>> Query().f1.one_of(['value 1', 'value 2'])
```

Parameters **items** – The list of items to check with

search (*regex, flags=0*)

Run a regex test against a dict value (only substring string has to match).

```
>>> Query().fl.search(r'^\w+$')
```

Parameters

- **regex** – The regular expression to use for matching
- **flags** – regex flags to pass to `re.match`

test (*func*, **args*)

Run a user-defined test function against a dict value.

```
>>> def test_func(val):  
...     return val == 42  
...  
>>> Query().fl.test(test_func)
```

Parameters

- **func** – The function to call, passing the dict as the first argument
- **args** – Additional arguments to pass to the test function

3.1.3 tinydb.storage

Contains the *base class* for storages and implementations.

class `tinydb.storages.Storage`

The abstract base class for all Storages.

A Storage (de)serializes the current state of the database and stores it in some place (memory, file on disk, ...).

read()

Read the last stored state.

write(*data*)

Write the current state of the database to the storage.

close()

Optional: Close open file handles, etc.

class `tinydb.storages.JSONStorage` (*path*, *create_dirs=False*, *encoding=None*, ***kwargs*)

Store the data in a JSON file.

__init__ (*path*, *create_dirs=False*, *encoding=None*, ***kwargs*)

Create a new instance.

Also creates the storage file, if it doesn't exist.

Parameters *path* (*str*) – Where to store the JSON data.

close()

Optional: Close open file handles, etc.

read()

Read the last stored state.

Any kind of deserialization should go here. Return `None` here to indicate that the storage is empty.

Return type dict

write (*data*)

Write the current state of the database to the storage.

Any kind of serialization should go here.

Parameters *data* (*dict*) – The current state of the database.

class `tinydb.storages.MemoryStorage`

Store the data as JSON in memory.

__init__ ()

Create a new instance.

read ()

Read the last stored state.

Any kind of deserialization should go here. Return `None` here to indicate that the storage is empty.

Return type `dict`

write (*data*)

Write the current state of the database to the storage.

Any kind of serialization should go here.

Parameters *data* (*dict*) – The current state of the database.

3.1.4 `tinydb.middlewares`

Contains the *base class* for middlewares and implementations.

class `tinydb.middlewares.Middleware`

The base class for all Middlewares.

Middlewares hook into the read/write process of TinyDB allowing you to extend the behaviour by adding caching, logging, ...

If `read()` or `write()` are not overloaded, they will be forwarded directly to the storage instance.

storage

Type *Storage*

Access to the underlying storage instance.

read ()

Read the last stored state.

write (*data*)

Write the current state of the database to the storage.

close ()

Optional: Close open file handles, etc.

class `tinydb.middlewares.CachingMiddleware` (*storage_cls=<class 'tinydb.storages.JSONStorage'>*)

Add some caching to TinyDB.

This Middleware aims to improve the performance of TinyDB by writing only the last DB state every `WRITE_CACHE_SIZE` time and reading always from cache.

__init__ (*storage_cls=<class 'tinydb.storages.JSONStorage'>*)

`x.__init__(...)` initializes x; see `help(type(x))` for signature

flush()

Flush all unwritten data to disk.

4.1 Contribution Guidelines

Whether reporting bugs, discussing improvements and new ideas or writing extensions: Contributions to TinyDB are welcome! Here's how to get started:

1. Check for open issues or open a fresh issue to start a discussion around a feature idea or a bug
2. Fork [the repository](#) on Github, create a new branch off the *master* branch and start making your changes (known as [GitHub Flow](#))
3. Write a test which shows that the bug was fixed or that the feature works as expected
4. Send a pull request and bug the maintainer until it gets merged and published :)

4.1.1 Philosophy of TinyDB

TinyDB aims to be simple and fun to use. Therefore two key values are simplicity and elegance of interfaces and code. These values will contradict each other from time to time. In these cases, try using as little magic as possible. In any case don't forget documenting code that isn't clear at first glance.

4.1.2 Code Conventions

In general the TinyDB source should always follow [PEP 8](#). Exceptions are allowed in well justified and documented cases. However we make a small exception concerning docstrings:

When using multiline docstrings, keep the opening and closing triple quotes on their own lines and add an empty line after it.

```
def some_function():  
    """  
    Documentation ...  
    """
```

(continues on next page)

```
# implementation ...
```

4.1.3 Version Numbers

TinyDB follows the [SemVer versioning guidelines](#). This implies that backwards incompatible changes in the API will increment the major version. So think twice before making such changes.

4.2 Changelog

4.2.1 Version Numbering

TinyDB follows the SemVer versioning guidelines. For more information, see semver.org

4.2.2 unreleased

- *nothing yet*

4.2.3 v3.15.2 (2019-12-18)

- Fix: handle `Query().search` and `Query().matches` with non-string values correctly (see [a forum post by foprel](#)).

4.2.4 v3.15.1 (2019-10-26)

- Internal change: fix missing values handling for `LRUCache`

4.2.5 v3.15.0 (2019-10-12)

- Feature: allow setting the parameters of TinyDB's default table (see [issue 278](#))

4.2.6 v3.14.2 (2019-09-13)

- Internal change: support correct iteration for `LRUCache` objects

4.2.7 v3.14.1 (2019-07-03)

- Internal change: fix `Query` class to permit subclass creation (see [pull request 270](#))

4.2.8 v3.14.0 (2019-06-18)

- Change: support for `ujson` is now deprecated (see [issue 258](#))

4.2.9 v3.13.0 (2019-03-16)

- Feature: direct access to a TinyDB instance's storage (see [issue 258](#))

4.2.10 v3.12.2 (2018-12-12)

- Internal change: convert documents to dicts during insertion (see [pull request 256](#))
- Internal change: use tuple literals instead of tuple class/constructor (see [pull request 247](#))
- Infra: ensure YAML tests are run (see [pull request 252](#))

4.2.11 v3.12.1 (2018-11-09)

- Fix: Don't break when searching the same query multiple times (see [pull request 249](#))
- Internal change: allow `collections.abc.Mutable` as valid document types (see [pull request 245](#))

4.2.12 v3.12.0 (2018-11-06)

- Feature: Add encoding option to `JSONStorage` (see [pull request 238](#))
- Internal change: allow `collections.abc.Mutable` as valid document types (see [pull request 245](#))

4.2.13 v3.11.1 (2018-09-13)

- Bugfix: Make path queries (`db.search(where('key'))`) work again (see [issue 232](#))
- Improvement: Add custom `repr` representations for main classes (see [pull request 229](#))

4.2.14 v3.11.0 (2018-08-20)

- **Drop official support for Python 3.3.** Python 3.3 has reached its official End Of Life as of September 29, 2017. It will probably continue to work, but will not be tested against ([issue 217](#))
- Feature: Allow extending TinyDB with a custom storage proxy class (see [pull request 224](#))
- Bugfix: Return list of document IDs for upsert when creating a new document (see [issue 223](#))

4.2.15 v3.10.0 (2018-07-21)

- Feature: Add support for regex flags (see [pull request 216](#))

4.2.16 v3.9.0 (2018-04-24)

- Feature: Allow setting a table class for single table only (see [issue 197](#))
- Internal change: call `fsync` after flushing `JSONStorage` (see [issue 208](#))

4.2.17 v3.8.1 (2018-03-26)

- Bugfix: Don't install tests as a package anymore (see [pull request #195](#))

4.2.18 v3.8.0 (2018-03-01)

- Feature: Allow disabling the query cache with `db.table(name, cache_size=0)` (see [pull request #187](#))
- Feature: Add `db.write_back(docs)` for replacing documents (see [pull request #184](#))

4.2.19 v3.7.0 (2017-11-11)

- Feature: `one_of` for checking if a value is contained in a list (see [issue 164](#))
- Feature: Upsert (insert if document doesn't exist, otherwise update; see <https://forum.m-siemens.de/d/30-primary-key-well-sort-of>)
- Internal change: don't read from storage twice during initialization (see <https://forum.m-siemens.de/d/28-reads-the-whole-data-file-twice>)

4.2.20 v3.6.0 (2017-10-05)

- Allow updating all documents using `db.update(fields)` (see [issue #157](#)).
- Rename elements to documents. Document IDs now available with `doc.doc_id`, using `doc.eid` is now deprecated (see [pull request #158](#))

4.2.21 v3.5.0 (2017-08-30)

- Expose the table name via `table.name` (see [issue #147](#)).
- Allow better subclassing of the `TinyDB` class (see [pull request #150](#)).

4.2.22 v3.4.1 (2017-08-23)

- Expose TinyDB version via `import tinyb; tinyb.__version__` (see [issue #148](#)).

4.2.23 v3.4.0 (2017-08-08)

- Add new update operations: `add(key, value)`, `subtract(key, value)`, and `set(key, value)` (see [pull request #145](#)).

4.2.24 v3.3.1 (2017-06-27)

- Use relative imports to allow vendoring TinyDB in other packages (see [pull request #142](#)).

4.2.25 v3.3.0 (2017-06-05)

- Allow iterating over a database or table yielding all documents (see [pull request #139](#)).

4.2.26 v3.2.3 (2017-04-22)

- Fix bug with accidental modifications to the query cache when modifying the list of search results (see [issue #132](#)).

4.2.27 v3.2.2 (2017-01-16)

- Fix the `Query` constructor to prevent wrong usage (see [issue #117](#)).

4.2.28 v3.2.1 (2016-06-29)

- Fix a bug with queries on documents that have a `path` key (see [pull request #107](#)).
- Don't write to the database file needlessly when opening the database (see [pull request #104](#)).

4.2.29 v3.2.0 (2016-04-25)

- Add a way to specify the default table name via `default_table` (see [pull request #98](#)).
- Add `db.purge_table(name)` to remove a single table (see [pull request #100](#)).
 - Along the way: celebrating 100 issues and pull requests! Thanks everyone for every single contribution!
- Extend API documentation (see [issue #96](#)).

4.2.30 v3.1.3 (2016-02-14)

- Fix a bug when using unhashable documents (lists, dicts) with `Query.any` or `Query.all` queries (see a [forum post by karibul](#)).

4.2.31 v3.1.2 (2016-01-30)

- Fix a bug when using unhashable documents (lists, dicts) with `Query.any` or `Query.all` queries (see a [forum post by karibul](#)).

4.2.32 v3.1.1 (2016-01-23)

- Inserting a dictionary with data that is not JSON serializable doesn't lead to corrupt files anymore (see [issue #89](#)).
- Fix a bug in the LRU cache that may lead to an invalid query cache (see [issue #87](#)).

4.2.33 v3.1.0 (2015-12-31)

- `db.update(...)` and `db.remove(...)` now return affected document IDs (see [issue #83](#)).
- Inserting an invalid document (i.e. not a `dict`) now raises an error instead of corrupting the database (see [issue #74](#)).

4.2.34 v3.0.0 (2015-11-13)

- Overhauled Query model:
 - `where('...').contains('...')` has been renamed to `where('...').search('...')`.
 - Support for ORM-like usage: `User = Query(); db.search(User.name == 'John')`.
 - `where('foo')` is an alias for `Query().foo`.
 - `where('foo').has('bar')` is replaced by either `where('foo').bar` or `Query().foo.bar`.
 - * In case the key is not a valid Python identifier, array notation can be used: `where('a.b.c')` is now `Query()['a.b.c']`.
 - Checking for the existence of a key has to be done explicitly: `where('foo').exists()`.
- Migrations from v1 to v2 have been removed.
- `SmartCacheTable` has been moved to [msiemens/tinydb-smartcache](#).
- Serialization has been moved to [msiemens/tinydb-serialization](#).
- Empty storages are now expected to return `None` instead of raising `ValueError`. (see [issue #67](#)).

4.2.35 v2.4.0 (2015-08-14)

- Allow custom parameters for custom test functions (see [issue #63](#) and [pull request #64](#)).

4.2.36 v2.3.2 (2015-05-20)

- Fix a forgotten debug output in the `SerializationMiddleware` (see [issue #55](#)).
- Fix an “ignored exception” warning when using the `CachingMiddleware` (see [pull request #54](#)).
- Fix a problem with symlinks when checking out TinyDB on OSX Yosemite (see [issue #52](#)).

4.2.37 v2.3.1 (2015-04-30)

- Hopefully fix a problem with using TinyDB as a dependency in a `setup.py` script (see [issue #51](#)).

4.2.38 v2.3.0 (2015-04-08)

- Added support for custom serialization. That way, you can teach TinyDB to store `datetime` objects in a JSON file :) (see [issue #48](#) and [pull request #50](#))
- Fixed a performance regression when searching became slower with every search (see [issue #49](#))
- Internal code has been cleaned up

4.2.39 v2.2.2 (2015-02-12)

- Fixed a data loss when using `CachingMiddleware` together with `JSONStorage` (see [issue #47](#))

4.2.40 v2.2.1 (2015-01-09)

- Fixed handling of IDs with the JSON backend that converted integers to strings (see [issue #45](#))

4.2.41 v2.2.0 (2014-11-10)

- Extended `any` and `all` queries to take lists as conditions (see [pull request #38](#))
- Fixed an `decode error` when installing TinyDB in a non-UTF-8 environment (see [pull request #37](#))
- Fixed some issues with `CachingMiddleware` in combination with `JSONStorage` (see [pull request #39](#))

4.2.42 v2.1.0 (2014-10-14)

- Added `where(...).contains(regex)` (see [issue #32](#))
- Fixed a bug that corrupted data after reopening a database (see [issue #34](#))

4.2.43 v2.0.1 (2014-09-22)

- Fixed handling of Unicode data in Python 2 (see [issue #28](#)).

4.2.44 v2.0.0 (2014-09-05)

Upgrade Notes

Warning: TinyDB changed the way data is stored. You may need to migrate your databases to the new scheme. Check out the [Upgrade Notes](#) for details.

- The syntax `query in db` has been removed, use `db.contains` instead.
- The `ConcurrencyMiddleware` has been removed due to a insecure implementation (see [issue #18](#)). Consider [tinyrecord](#) instead.
- Better support for working with *Document IDs*.
- Added support for [nested comparisons](#).
- Added `all` and `any` [comparisons on lists](#).
- Added optional `<http://tinydb.readthedocs.io/en/v2.0.0/usage.html#smart-query-cache>‘_`.
- The query cache is now a *fixed size LRU cache*.

4.2.45 v1.4.0 (2014-07-22)

- Added `insert_multiple` function (see [issue #8](#)).

4.2.46 v1.3.0 (2014-07-02)

- Fixed [bug #7](#): IDs not unique.
- Extended the API: `db.count(where(...))` and `db.contains(where(...))`.
- The syntax query `in db` is now **deprecated** and replaced by `db.contains`.

4.2.47 v1.2.0 (2014-06-19)

- Added update method (see [issue #6](#)).

4.2.48 v1.1.1 (2014-06-14)

- Merged [PR #5](#): Fix minor documentation typos and style issues.

4.2.49 v1.1.0 (2014-05-06)

- Improved the docs and fixed some typos.
- Refactored some internal code.
- Fixed a bug with multiple TinyDB? instances.

4.2.50 v1.0.1 (2014-04-26)

- Fixed a bug in `JSONStorage` that broke the database when removing entries.

4.2.51 v1.0.0 (2013-07-20)

- First official release – consider TinyDB stable now.

4.3 Upgrading to Newer Releases

4.3.1 Version 3.0

Breaking API Changes

- Querying (see [Issue #62](#)):
 - `where('...').contains('...')` has been renamed to `where('...').search('...')`.
 - `where('foo').has('bar')` is replaced by either `where('foo').bar` or `Query().foo.bar`.
 - * In case the key is not a valid Python identifier, array notation can be used: `where('a.b.c')` is now `Query()['a.b.c']`.
- Checking for the existence of a key has to be done explicitly: `where('foo').exists()`.
- `SmartCacheTable` has been moved to [msiemens/tinydb-smartcache](#).
- Serialization has been moved to [msiemens/tinydb-serialization](#).

- Empty storages are now expected to return `None` instead of raising `ValueError` (see [Issue #67](#)).

4.3.2 Version 2.0

Breaking API Changes

- The syntax `query in db` is not supported any more. Use `db.contains(...)` instead.
- The `ConcurrencyMiddleware` has been removed due to a insecure implementation (see [Issue #18](#)). Consider *tinyrecord* instead.

Apart from that the API remains compatible to v1.4 and prior.

For migration from v1 to v2, check out the [v2.0 documentation](#)

t

`tinydb.middlewares`, [31](#)

`tinydb.storages`, [30](#)

Symbols

`__eq__()` (*tinydb.queries.Query method*), 27
`__ge__()` (*tinydb.queries.Query method*), 28
`__getattr__()` (*tinydb.database.TinyDB method*), 23
`__gt__()` (*tinydb.queries.Query method*), 28
`__hash__()` (*tinydb.queries.Query method*), 28
`__init__()` (*tinydb.database.Document method*), 27
`__init__()` (*tinydb.database.Table method*), 24
`__init__()` (*tinydb.database.TinyDB method*), 23
`__init__()` (*tinydb.middlewares.CachingMiddleware method*), 31
`__init__()` (*tinydb.queries.Query method*), 28
`__init__()` (*tinydb.storages.JSONStorage method*), 30
`__init__()` (*tinydb.storages.MemoryStorage method*), 31
`__iter__()` (*tinydb.database.Table method*), 24
`__iter__()` (*tinydb.database.TinyDB method*), 23
`__le__()` (*tinydb.queries.Query method*), 28
`__len__()` (*tinydb.database.Table method*), 24
`__len__()` (*tinydb.database.TinyDB method*), 23
`__lt__()` (*tinydb.queries.Query method*), 28
`__ne__()` (*tinydb.queries.Query method*), 28
`__repr__()` (*tinydb.database.Table method*), 24
`__repr__()` (*tinydb.database.TinyDB method*), 23
`__repr__()` (*tinydb.queries.Query method*), 28

A

`all()` (*tinydb.database.Table method*), 25
`all()` (*tinydb.queries.Query method*), 28
`any()` (*tinydb.queries.Query method*), 29

C

`CachingMiddleware` (*class in tinydb.middlewares*), 31
`clear_cache()` (*tinydb.database.Table method*), 25
`close()` (*tinydb.database.TinyDB method*), 24
`close()` (*tinydb.middlewares.Middleware method*), 31
`close()` (*tinydb.storages.JSONStorage method*), 30

`close()` (*tinydb.storages.Storage method*), 30
`contains()` (*tinydb.database.Table method*), 25
`count()` (*tinydb.database.Table method*), 25

D

`DEFAULT_STORAGE` (*tinydb.database.TinyDB attribute*), 23
`doc_id` (*tinydb.database.Document attribute*), 27
`Document` (*class in tinydb.database*), 27

E

`Element` (*in module tinydb.database*), 27
`exists()` (*tinydb.queries.Query method*), 29

F

`flush()` (*tinydb.middlewares.CachingMiddleware method*), 31

G

`get()` (*tinydb.database.Table method*), 25

I

`insert()` (*tinydb.database.Table method*), 25
`insert_multiple()` (*tinydb.database.Table method*), 25

J

`JSONStorage` (*class in tinydb.storages*), 30

M

`matches()` (*tinydb.queries.Query method*), 29
`MemoryStorage` (*class in tinydb.storages*), 31
`Middleware` (*class in tinydb.middlewares*), 31

N

`name` (*tinydb.database.Table attribute*), 25

O

`one_of()` (*tinydb.queries.Query method*), 29

P

`process_elements()` (*tinydb.database.Table method*), 25
`purge()` (*tinydb.database.Table method*), 26
`purge_table()` (*tinydb.database.TinyDB method*), 24
`purge_tables()` (*tinydb.database.TinyDB method*), 24

Q

`Query` (*class in tinydb.queries*), 27

R

`read()` (*tinydb.middlewares.Middleware method*), 31
`read()` (*tinydb.storages.JSONStorage method*), 30
`read()` (*tinydb.storages.MemoryStorage method*), 31
`read()` (*tinydb.storages.Storage method*), 30
`remove()` (*tinydb.database.Table method*), 26

S

`search()` (*tinydb.database.Table method*), 26
`search()` (*tinydb.queries.Query method*), 29
`Storage` (*class in tinydb.storages*), 30
`storage` (*tinydb.database.TinyDB attribute*), 24
`storage` (*tinydb.middlewares.Middleware attribute*), 31
`storage_proxy_class` (*tinydb.database.TinyDB attribute*), 24

T

`Table` (*class in tinydb.database*), 24
`table()` (*tinydb.database.TinyDB method*), 24
`table_class` (*tinydb.database.TinyDB attribute*), 24
`tables()` (*tinydb.database.TinyDB method*), 24
`test()` (*tinydb.queries.Query method*), 30
`TinyDB` (*class in tinydb.database*), 23
`tinydb.middlewares` (*module*), 31
`tinydb.storages` (*module*), 30

U

`update()` (*tinydb.database.Table method*), 26
`upsert()` (*tinydb.database.Table method*), 26

W

`write()` (*tinydb.middlewares.Middleware method*), 31
`write()` (*tinydb.storages.JSONStorage method*), 30
`write()` (*tinydb.storages.MemoryStorage method*), 31
`write()` (*tinydb.storages.Storage method*), 30
`write_back()` (*tinydb.database.Table method*), 26