Iima Documentation

Release 0.5

Bernhard Weitzhofer
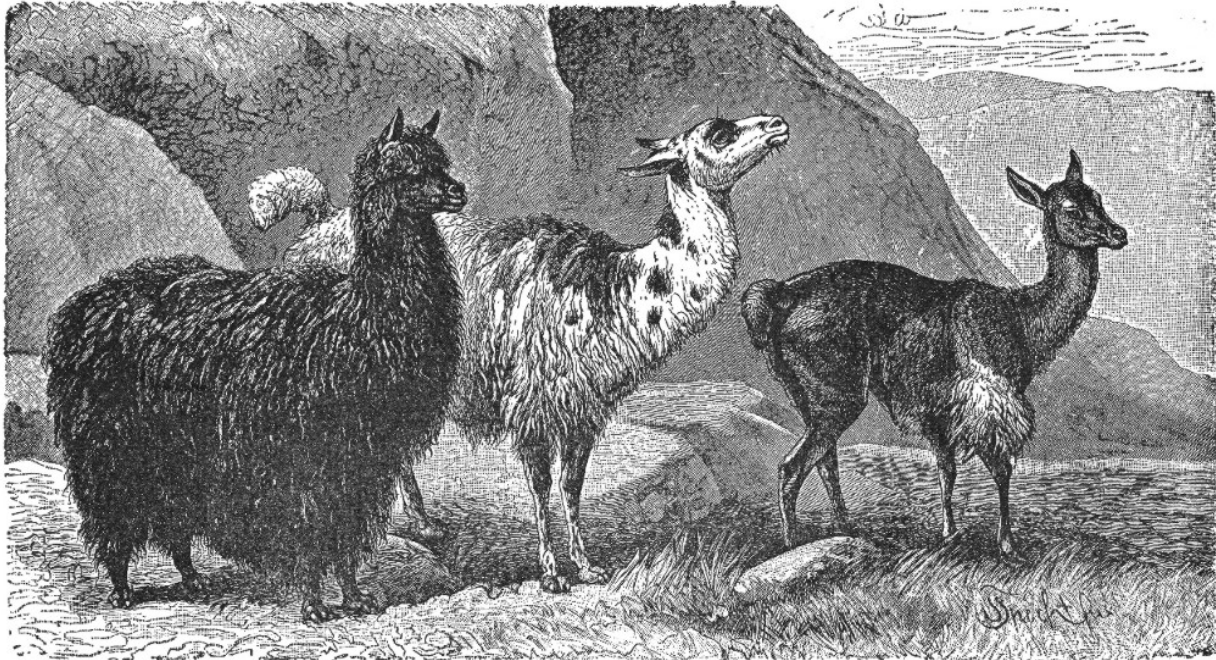
May 12, 2015

Contents

lima takes arbitrary Python objects and converts them into data structures native to Python. The result can easily be serialized into JSON, XML, and all sorts of other things. lima is Free Software, lightweight and fast.

ALPACA   LLAMA   VICUNA

lima at a Glance

```python
import datetime
import lima

# a model
class Book:
    def __init__(self, title, date_published):
        self.title = title
        self.date_published = date_published

# a marshalling schema
class BookSchema(lima.Schema):
    title = lima.fields.String()
    published = lima.fields.Date(attr='date_published')

book = Book('The Old Man and the Sea', datetime.date(1952, 9, 1))
schema = BookSchema()
schema.dump(book)
# {'published': '1952-09-01', 'title': 'The Old Man and the Sea'}
```

Key Features

Lightweight lima has only a few hundred SLOC. lima has no external dependencies.

Fast lima tries to be as fast as possible while still remaining pure Python 3.

Well documented lima has a comprehensive *tutorial* and more than one line of *docstring* per line of Python code.

Free lima is Free Software, licensed under the terms of the *MIT license*.

Documentation

3.1 Installation

The recommended way to install lima is via pip.

Just make sure you have at least Python 3.3 and a matching version of pip available and installing lima becomes a one-liner:

```
$ pip install lima
```

Most of the time it's also a good idea to do this in an isolated virtual environment.

Starting with version 3.4, Python handles all of this (creation of virtual environments, ensuring the availability of pip) out of the box:

```
$ python3 -m venv /path/to/my_venv
$ source /path/to/my_venv/bin/activate
(my_venv) $ pip install lima
```

If you should run into trouble, the Tutorial on Installing Distributions from the Python Packaging User Guide might be helpful.

3.2 First Steps

lima tries to be lean, consistent, and easy to learn. Assuming you already have *installed* lima, this section should help you getting started.

Note: Throughout this documentation, the terms *marshalling* and *serialization* will be used synonymously.

3.2.1 A simple Example

Let's say we want to expose our data to the world via a web API and we've chosen JSON as our preferred serialization format. We have defined a data model in the ORM of our choice. It might behave something like this:

```python
class Person:
    def __init__(self, first_name, last_name, date_of_birth):
        self.first_name = first_name
        self.last_name = last_name
        self.date_of_birth = date_of_birth
```

Our person objects look like this:

```python
import datetime
person = Person('Ernest', 'Hemingway', datetime.date(1899, 7, 21))
```

If we want to serialize such person objects, we can't just feed them to Python's `json.dumps()` function: per default it only knows how to deal with a very basic set of data types.

Here's where lima comes in: Defining an appropriate Schema, we can convert person objects into data structures accepted by `json.dumps()`.

```python
from lima import fields, Schema


class PersonSchema(Schema):
    first_name = fields.String()
    last_name = fields.String()
    date_of_birth = fields.Date()

schema = PersonSchema()
serialized = schema.dump(person)
# {'date_of_birth': '1899-07-21',
#  'first_name': 'Ernest',
#  'last_name': 'Hemingway'}
```

... and to conclude our example:

```python
import json
json.dumps(serialized)
# '{"last_name": "Hemingway", "date_of_birth": "1899-07-21", ...
```

3.2.2 First Steps Recap

- You now know how to do basic marshalling (Create a schema class with appropriate fields. Create a schema object. Pass the object(s) to marshal to the schema object's `dump()` method.

- You now know how to get JSON for arbitrary objects (pass the result of a schema object's `dump()` method to `json.dumps()`).

3.3 Working with Schemas

Schemas collect fields for object serialization.

3.3.1 Defining Schemas

We already know how to define schemas: subclass `lima.Schema` (the shortcut for `lima.schema.Schema`) and add fields as class attributes.

But there's more to schemas than this. First of all – schemas are composible:

```python
from lima import Schema, fields


class PersonSchema(Schema):
    first_name = fields.String()
    last_name = fields.String()
```

```
class AccountSchema(Schema):
    login = fields.String()
    password_hash = fields.String()


class UserSchema(PersonSchema, AccountSchema):
    pass


list(UserSchema.__fields__)
# ['first_name', 'last_name', 'login', 'password_hash']
```

Secondly, it's possible to *remove* fields from subclasses that are present in superclasses. This is done by setting a special class attribute __lima_args__ like so:

```
class UserProfileSchema(UserSchema):
    __lima_args__ = {'exclude': ['last_name', 'password_hash']}


list(UserProfileSchema.__fields__)
# ['first_name', 'login']
```

If there's only one field to exclude, you don't have to put its name inside a list - lima does that for you:

```
class NoLastNameSchema(UserSchema):
    __lima_args__ = {'exclude': 'last_name'}  # string instead of list


list(NoLastNameSchema.__fields__)
# ['first_name', 'login', 'password_hash']
```

If, on the other hand, there are lots of fields to exclude, you *could* provide __lima_args__['only'] (Note that "exclude" and "only" are mutually exclusive):

```
class JustNameSchema(UserSchema):
    __lima_args__ = {'only': ['first_name', 'last_name']}


list(JustNameSchema.__fields__)
# ['first_name', 'last_name']
```

> **Warning:** Having to provide "only" on Schema definition hints at bad design - why would you add a lot of fields just to remove all but one of them afterwards? Have a look at *Schema Objects* for the preferred way to selectively remove fields.

And finally, we can't just *exclude* fields, we can *include* them too. So here is a user schema with fields provided via __lima_args__:

```
class UserSchema(Schema):
    __lima_args__ = {
        'include': {
            'first_name': fields.String(),
            'last_name': fields.String(),
            'login': fields.String(),
            'password_hash': fields.String()
        }
    }


list(UserSchema.__fields__)
# ['password_hash', 'last_name', 'first_name', 'login']
```

> **Note:** It's possible to mix and match all those features to your heart's content. lima tries to fail early if something doesn't add up (remember, "exclude" and "only" are mutually exclusive).

Note: The inheritance and precedence rules for fields are intuitive, but should there ever arise the need for clarification, you can read about how a schema's fields are determined in the documentation of `lima.schema.SchemaMeta`.

3.3.2 Schema Objects

Up until now we only ever needed a single instance of a schema class to marshal the fields defined in this class. But schema objects can do more.

Providing the keyword-only argument `exclude`, we may exclude certain fields from being serialized.

Keyword-only arguments

Keyword-only arguments can be recognized by their position in a method/function signature: Every argument coming after the varargs argument like `*args` (or after a single `*`) is a keyword-only argument.

A function that is defined as `def foo(*, x, y):  pass` *must* be called like this: `foo(x=1, y=2)`; calling `foo(1, 2)` will raise a `TypeError`.

It is the author's opinion that enforcing keyword arguments in the right places makes the resulting code more readable.

For more information about keyword-only arguments, see PEP 3102

```python
import datetime
from lima import Schema, fields

# again, our model
class Person:
    def __init__(self, first_name, last_name, birthday):
        self.first_name = first_name
        self.last_name = last_name
        self.birthday = birthday

# again, our schema
class PersonSchema(Schema):
    first_name = fields.String()
    last_name = fields.String()
    date_of_birth = fields.Date(attr='birthday')

# again, our person
person = Person('Ernest', 'Hemingway', datetime.date(1899, 7, 21))

# as before, for reference
person_schema = PersonSchema()
person_schema.dump(person)
# {'date_of_birth': '1899-07-21',
#  'first_name': 'Ernest',
#  'last_name': 'Hemingway'}

birthday_schema = PersonSchema(exclude=['first_name', 'last_name'])
birthday_schema.dump(person)
# {'date_of_birth': '1899-07-21'}
```

The same thing can be achieved via the `only` keyword-only argument:

```
birthday_schema = PersonSchema(only='date_of_birth')
birthday_schema.dump(person)
# {'date_of_birth': '1899-07-21'}
```

You may have already guessed: both `exclude` and `only` take lists of field names as well as simple strings for a single field name – just like `__lima_args__['exclude']` and `__lima_args__['only']`.

For some use cases, `exclude` and `only` save the need to define lots of almost similar schema classes.

You *could* also include fields on schema object creation time:

```
getter = lambda o: '{}, {}'.format(o.last_name, o.first_name)

schema = PersonSchema(include={'sort_name': fields.String(get=getter)})

schema.dump(person)
# {'date_of_birth': '1899-07-21',
#  'first_name': 'Ernest',
#  'last_name': 'Hemingway',
#  'sort_name': 'Hemingway, Ernest'}
```

> **Warning:** Having to provide `include` on Schema object creation hints at bad design - why not just include the fields in the Schema itself?

3.3.3 Field Order

Lima marshals objects to dictionaries. Field order doesn't matter. Unless you want it to:

```
person_schema = PersonSchema(ordered=True)
person_schema.dump(person)
# OrderedDict([
#     ('first_name', 'Ernest'),
#     ('last_name', 'Hemingway'),
#     ('date_of_birth', '1899-07-21')])
# ])
```

Just provide the keyword-only argument `ordered=True` to a schema's constructor, and the resulting instance will dump ordered dictionaries.

The order of the resulting key-value-pairs reflects the order in which the fields were defined at schema definition time.

If you use `__lima_args__['include']`, make sure to provide an instance of `collections.OrderedDict` if you care about the order of those fields as well.

Fields specified via `__lima_args__['include']` are inserted at the position of the `__lima_args__` class attribute in the Schema class. Here is a more complex example:

```
from collections import OrderedDict

class FooSchema(Schema):
    one = fields.String()
    two = fields.String()

class BarSchema(FooSchema):
    three = fields.String()
    __lima_args__ = {
        'include': OrderedDict([
            ('four', fields.String()),
```

```
                ('five', fields.String())
            ])
        }
    six = fields.String()

bar_schema = BarSchema(ordered=True)
```

`bar_schema` will dump ordered dictionaries with keys ordered from `one` to `six`.

Note: For the exact rules on how a complex schema's fields are going to be ordered, see `lima.schema.SchemaMeta` or have a look at the source code.

3.3.4 Marshalling Collections

Consider this:

```
persons = [
    Person('Ernest', 'Hemingway', datetime.date(1899, 7, 21)),
    Person('Virginia', 'Woolf', datetime.date(1882, 1, 25)),
    Person('Stefan', 'Zweig', datetime.date(1881, 11, 28)),
]
```

Instead of looping over this collection ourselves, we can ask the schema object to do this for us by specifying `many=True` to the schema's constructor):

```
many_persons_schema = PersonSchema(only='last_name', many=True)
many_persons_schema.dump(persons)
# [{'last_name': 'Hemingway'},
#  {'last_name': 'Woolf'},
#  {'last_name': 'Zweig'}]
```

3.3.5 Schema Recap

- You now know how to compose bigger schemas from smaller ones (inheritance of schema classes).

- You know how to exclude certain fields from schemas (`__lima_args__['exclude']`).

- You know three different ways to add fields to schemas (class attributes, `__lima_args__['include']` and inheriting from other schemas).

- You can fine-tune what gets dumped by a schema object (`only` and `exclude` keyword-only arguments)

- You can dump ordered dictionaries (`ordered=True`) and you can serialize collections of objects (`many=True`).

3.4 A closer Look at Fields

Fields are the basic building blocks of a Schema. Even though lima fields follow only the most basic protocol, they are rather powerful.

3.4.1 How a Field gets its Data

The `PersonSchema` from the last chapter contains three field objects named *first_name, last_name* and *date_of_birth*. These get their data from a person object's attributes of the same name. But what if those attributes were named differently?

Data from arbitrary Object Attributes

Let's say our model doesn't have an attribute `date_of_birth` but an attribute `birthday` instead.

To get the data for our `date_of_birth` field from the model's `birthday` attribute, we have to tell the field by supplying the attribute name via the `attr` argument:

```python
import datetime
from lima import Schema, fields


class Person:
    def __init__(self, first_name, last_name, birthday):
        self.first_name = first_name
        self.last_name = last_name
        self.birthday = birthday

person = Person('Ernest', 'Hemingway', datetime.date(1899, 7, 21))


class PersonSchema(Schema):
    first_name = fields.String()
    last_name = fields.String()
    date_of_birth = fields.Date(attr='birthday')

schema = PersonSchema()
schema.dump(person)
# {'date_of_birth': '1899-07-21',
#  'first_name': 'Ernest',
#  'last_name': 'Hemingway'}
```

Data from an Object's Items

If an object has it's data stored in items instead of attributes, we can tell the field about this by supplying the `key` argument instead of the `attr` argument:

```python
import datetime
from lima import Schema, fields

person_dict = {
    'first_name': 'Ernest',
    'last_name': 'Hemingway',
    'birthday': datetime.date(1899, 7, 21),
}


class PersonDictSchema(Schema):
    last_name = fields.String(key='last_name')
    date_of_birth = fields.Date(key='birthday')

schema = PersonDictSchema()
schema.dump(person_dict)
```

```
# {'date_of_birth': '1899-07-21',
#  'last_name': 'Hemingway'}
```

Note: It's currently not possible to provide `None` as key. use a *getter* (see below) if you need to do this.

Data derived by different Means

What if we can't get the information we need from a single attribute or key? Here *getters* come in handy.

A getter in this context is a callable that takes an object (in our case: a person object) and returns the value we're interested in. We tell a field about the getter via the `get` parameter:

```python
def sort_name_getter(obj):
    return '{}, {}'.format(obj.last_name, obj.first_name)


class PersonSchema(Schema):
    first_name = fields.String()
    last_name = fields.String()
    sort_name = fields.String(get=sort_name_getter)
    date_of_birth = fields.Date(attr='birthday')


schema = PersonSchema()
schema.dump(person)
# {'date_of_birth': '1899-07-21',
#  'first_name': 'Ernest',
#  'last_name': 'Hemingway'
#  'sort_name': 'Hemingway, Ernest'}
```

Note: For getters, lambda expressions come in handy. `sort_name` could just as well have been defined like this:

```python
sort_name = fields.String(
    get=lambda obj: '{}, {}'.format(obj.last_name, obj.first_name)
)
```

Constant Field Values

Sometimes a field's data is always the same. For example, if a schema provides a field for type information, this field will most likely always have the same value.

To reflect this, we could provide a getter that always returns the same value (here, for example, the string `'https:/schema.org/Person'`). But lima provides a better way to achieve the same result: Just provide the `val` parameter to a field's constructor:

```python
class TypedPersonSchema(Schema):
    _type = fields.String(val='https://schema.org/Person')
    givenName = fields.String(attr='first_name')
    familyName = fields.String(attr='last_name')
    birthDate = fields.Date(attr='birthday')


schema = TypedPersonSchema()
schema.dump(person)
# {'_type': 'https://schema.org/Person',
#  'birthDate': '1899-07-21',
#  'familyName': 'Hemingway',
#  'givenName': 'Ernest'}
```

Note: It's currently not possible to provide `None` as a constant value using `val` - use a getter if you need to do this.

On Field Parameters

`attr`, `get` and `val` are mutually exclusive. See `lima.fields.Field` for more information on this topic.

3.4.2 How a Field presents its Data

If a field has a static method (or instance method) `pack()`, this method is used to present a field's data. (Otherwise the field's data is just passed through on marshalling. Some of the more basic built-in fields behave that way.)

So by implementing a `pack()` static method (or instance method), we can support marshalling of any data type we want:

```python
from collections import namedtuple
from lima import fields, Schema

# a new data type
GeoPoint = namedtuple('GeoPoint', ['lat', 'long'])

# a field class for the new date type
class GeoPointField(fields.Field):
    @staticmethod
    def pack(val):
        ns = 'N' if val.lat > 0 else 'S'
        ew = 'E' if val.long > 0 else 'W'
        return '{}° {}, {}° {}'.format(val.lat, ns, val.long, ew)

# a model using the new data type
class Treasure:
    def __init__(self, name, location):
        self.name = name
        self.location = location

# a schema for that model
class TreasureSchema(Schema):
    name = fields.String()
    location = GeoPointField()

treasure = Treasure('The Amber Room', GeoPoint(lat=59.7161, long=30.3956))
schema = TreasureSchema()
schema.dump(treasure)
# {'location': '59.7161° N, 30.3956° E', 'name': 'The Amber Room'}
```

Or we can change how already supported data types are marshalled:

```python
class FancyDate(fields.Date):
    @staticmethod
    def pack(val):
        return val.strftime('%A, the %d. of %B %Y')

class FancyPersonSchema(Schema):
    first_name = fields.String()
    last_name = fields.String()
```

```
    date_of_birth = FancyDate(attr='birthday')

schema = FancyPersonSchema()
schema.dump(person)
# {'date_of_birth': 'Friday, the 21. of July 1899',
#  'first_name': 'Ernest',
#  'last_name': 'Hemingway'}
```

> **Warning:** Make sure the result of your `pack()` methods is JSON serializable (or at least in a format accepted by the serializer of your target format).
>
> Also, don't try to override an existing instance method with a static method. Have a look at the source if in doubt (in lima itself, currently only `lima.fields.Embed` and `lima.fields.Reference` implement `pack()` as instance methods.

3.4.3 Data Validation

In short: *There is none.*

lima is opinionated in this regard. It assumes you have control over the data you want to serialize and have already validated it *before* putting it in your database.

But this doesn't mean it can't be done. You'll just have to do it yourself. The `pack()` method would be the place for this:

```python
import re

class ValidEmailField(fields.String):
    @staticmethod
    def pack(val):
        if not re.match(r'[^@]+@[^@]+\.[^@]+', val):
            raise ValueError('Not an email address: {!r}'.format(val))
        return val
```

> **Note:** If you need full-featured validation of your existing data at marshalling time, have a look at marshmallow.

3.4.4 Fields Recap

- You now know how it's determined where a field's data comes from. (from least to highest precedence: field name < attr < getter < constant field value.

- You know how a field presents its data (`pack()` method).

- You know how to support your own data types (subclass `lima.fields.Field`) and implement `pack()`

- And you know how to change the marshalling of already supported data types (subclass the appropriate field class and override `pack()`)

- Also, you're able to implement data validation should the need arise (implement/override `pack()`).

3.5 Linked Data

Let's model a relationship between a book and a book review:

```python
class Book:
    def __init__(self, isbn, author, title):
        self.isbn = isbn
        self.author = author
        self.title = title

# A review links to a book via the "book" attribute
class Review:
    def __init__(self, rating, text, book=None):
        self.rating = rating
        self.text = text
        self.book = book

book = Book('0-684-80122-1', 'Hemingway', 'The Old Man and the Sea')
review = Review(10, 'Has lots of sharks.')
review.book = book
```

To serialize this construct, we have to tell lima that a `Review` object links to a `Book` object via the `book` attribute (many ORMs represent related objects in a similar way).

3.5.1 Embedding linked Objects

We can use a field of type `lima.fields.Embed` to *embed* the serialized book into the serialization of the review. For this to work we have to tell the `Embed` field what to expect by providing the `schema` parameter:

```python
from lima import fields, Schema

class BookSchema(Schema):
    isbn = fields.String()
    author = fields.String()
    title = fields.String()

class ReviewSchema(Schema):
    book = fields.Embed(schema=BookSchema)
    rating = fields.Integer()
    text = fields.String()

review_schema = ReviewSchema()
review_schema.dump(review)
# {'book': {'author': 'Hemingway',
#           'isbn': '0-684-80122-1',
#           'title': 'The Old Man and the Sea'},
#  'rating': 10,
#  'text': 'Has lots of sharks.'}
```

Along with the mandatory keyword-only argument `schema`, `Embed` accepts the optional keyword-only-arguments we already know (`attr`, `get`, `val`). All other keyword arguments provided to `Embed` get passed through to the constructor of the associated schema. This allows us to do stuff like the following:

```python
class ReviewSchemaPartialBook(Schema):
    rating = fields.Integer()
    text = fields.String()
    partial_book = fields.Embed(attr='book',
                                schema=BookSchema,
                                exclude='isbn')

review_schema_partial_book = ReviewSchemaPartialBook()
```

```
review_schema_partial_book.dump(review)
# {'partial_book': {'author': 'Hemingway',
#                    'title': 'The Old Man and the Sea'},
#  'rating': 10,
#  'text': 'Has lots of sharks.'}
```

3.5.2 Referencing linked Objects

Embedding linked objects is not always what we want. If we just want to reference linked objects, we can use a field of type `lima.fields.Reference`. This field type yields the value of a single field of the linked object's serialization.

Referencing is similar to embedding save one key difference: In addition to the schema of the linked object we also provide the name of the field that acts as our reference to the linked object. We may, for example, reference a book via its ISBN like this:

```
class ReferencingReviewSchema(Schema):
    book = fields.Reference(schema=BookSchema, field='isbn')
    rating = fields.Integer()
    text = fields.String()


referencing_review_schema = ReferencingReviewSchema()
referencing_review_schema.dump(review)
# {'book': '0-684-80122-1',
#  'rating': 10,
#  'text': 'Has lots of sharks.'}
```

3.5.3 Hyperlinks

One application of `Reference` is linking to ressources via hyperlinks in RESTful Web services. Here is a quick sketch:

```
# your framework should provide something like this
def book_url(book):
    return 'https://my.service/books/{}'.format(book.isbn)


class BookSchema(Schema):
    url = fields.String(get=book_url)
    isbn = fields.String()
    author = fields.String()
    title = fields.String()


class ReviewSchema(Schema):
    book = fields.Reference(schema=BookSchema, field='url')
    rating = fields.Integer()
    text = fields.String()


review_schema = ReviewSchema()
review_schema.dump(review)
# {'book': 'https://my.service/books/0-684-80122-1',
#  'rating': 10,
#  'text': 'Has lots of sharks.'}
```

Note: If you want to do JSON-LD, you may want to have fields with names like `"@id"` or `"@context"`. Have a look at the section on *Field Name Mangling* for an easy way to accomplish this.

3.5.4 Two-way Relationships

Up until now, we've only dealt with one-way relationships (*From* a review *to* its book). If not only a review should link to its book, but a book should also link to it's most popular review, we can adapt our model like this:

```python
# books now link to their most popular review
class Book:
    def __init__(self, isbn, author, title, pop_review=None):
        self.isbn = isbn
        self.author = author
        self.title = title
        self.pop_review = pop_review

# unchanged: reviews still link to their books
class Review:
    def __init__(self, rating, text, book=None):
        self.rating = rating
        self.text = text
        self.book = book

book = Book('0-684-80122-1', 'Hemingway', 'The Old Man and the Sea')
review = Review(4, "Why doesn't he just kill ALL the sharks?")

book.pop_review = review
review.book = book
```

If we want to construct schemas for models like this, we will have to adress two problems:

1. **Definition order:** If we define our `BookSchema` first, its `pop_review` attribute will have to reference a `ReviewSchema` - but this doesn't exist yet, since we decided to define `BookSchema` first. If we decide to define `ReviewSchema` first instead, we run into the same problem with its `book` attribute.

2. **Recursion:** A review links to a book that links to a review that links to a book that links to a review that links to a book that links to a review that links to a book `RuntimeError: maximum recursion depth exceeded`

lima makes it easy to deal with those problems:

To overcome the problem of recursion, just exclude the attribute on the other side that links back.

To overcome the problem of definition order, lima supports lazy evaluation of schemas. Just pass the *qualified name* (or the *fully module-qualified name*) of a schema class to Embed instead of the class itself:

```python
class BookSchema(Schema):
    isbn = fields.String()
    author = fields.String()
    title = fields.String()
    pop_review = fields.Embed(schema='ReviewSchema', exclude='book')

class ReviewSchema(Schema):
    book = fields.Embed(schema=BookSchema, exclude='pop_review')
    rating = fields.Integer()
    text = fields.String()
```

Now embedding works both ways:

```python
book_schema = BookSchema()
book_schema.dump(book)
```

```
# {'author': 'Hemingway',
#  'isbn': '0-684-80122-1',
#  'pop_review': {'rating': 4,
#                 'text': "Why doesn't he just kill ALL the sharks?"},
#  'title': The Old Man and the Sea'}

review_schema = ReviewSchema()
review_schema.dump(review)
# {'book': {'author': 'Hemingway',
#           'isbn': '0-684-80122-1',
#           'title': 'The Old Man and the Sea'},
#  'rating': 4,
#  'text': "Why doesn't he just kill ALL the sharks?"}
```

On class names

For referring to classes via their name, the lima documentation only ever talks about two different kinds of class names: the *qualified name* (*qualname* for short) and the *fully module-qualified name*:

The qualified name This is the value of the class's `__qualname__` attribute. Most of the time, it's the same as the class's `__name__` attribute (except if you define classes within classes or functions ...). If you define `class Foo: pass` at the top level of your module, the class's qualified name is simply `Foo`. Qualified names were introduced with Python 3.3 via PEP 3155

The fully module-qualified name This is the qualified name of the class prefixed with the full name of the module the class is defined in. If you define `class Qux: pass` within a class `Baz` (resulting in the qualified name `Baz.Qux`) at the top level of your `foo.bar` module, the class's fully module-qualified name is `foo.bar.Baz.Qux`.

> **Warning:** If you define schemas in local namespaces (at function execution time), their names become meaningless outside of their local context. For example:
>
> ```
> def make_schema():
>     class FooSchema(Schema):
>         foo = fields.String()
>     return FooSchema
>
> schemas = [make_schema() for i in range(1000)]
> ```
>
> Which of those one thousend schemas would we refer to, would we try to link to a `FooSchema` by name? To avoid ambiguity, lima will refuse to link to schemas defined in local namespaces.

By the way, there's nothing stopping us from using the idioms we just learned for models that link to themselves - everything works as you'd expect:

```
class MarriedPerson:
    def __init__(self, first_name, last_name, spouse=None):
        self.first_name = first_name
        self.last_name = last_name
        self.spouse = spouse

class MarriedPersonSchema(Schema):
    first_name = fields.String()
    last_name = fields.String()
    spouse = fields.Embed(schema='MarriedPersonSchema', exclude='spouse')
```

3.5.5 One-to-many and many-to-many Relationships

Until now, we've only dealt with one-to-one relations. What about one-to-many and many-to-many relations? Those link to collections of objects.

We know the necessary building blocks already: Providing additional keyword arguments to `Embed` (or `Reference` respectively) passes them through to the specified schema's constructor. And providing `many=True` to a schema's construtor will have the schema marshalling collections - so if our model looks like this:

```python
# books now have a list of reviews
class Book:
    def __init__(self, isbn, author, title):
        self.isbn = isbn
        self.author = author
        self.title = title
        self.reviews = []


class Review:
    def __init__(self, rating, text, book=None):
        self.rating = rating
        self.text = text
        self.book = book

book = Book('0-684-80122-1', 'Hemingway', 'The Old Man and the Sea')
book.reviews = [
    Review(10, 'Has lots of sharks.', book),
    Review(4, "Why doesn't he just kill ALL the sharks?", book),
    Review(8, 'Better than the movie!', book),
]
```

... we wourld define our schemas like this:

```python
class BookSchema(Schema):
    isbn = fields.String()
    author = fields.String()
    title = fields.String()
    reviews = fields.Embed(schema='ReviewSchema',
                           many=True,
                           exclude='book')


class ReviewSchema(Schema):
    book = fields.Embed(schema=BookSchema, exclude='reviews')
    rating = fields.Integer()
    text = fields.String()
```

... which enables us to serialize a book object with many reviews:

```python
book_schema = BookSchema()
book_schema.dump(book)
# {'author': 'Hemingway',
#  'isbn': '0-684-80122-1',
#  'reviews': [
#      {'rating': 10, 'text': 'Has lots of sharks.'},
#      {'rating': 4, 'text': "Why doesn't he just kill ALL the sharks?"},
#      {'rating': 8, 'text': 'Better than the movie!'}],
#  'title': The Old Man and the Sea'
```

3.5.6 Linked Data Recap

- You now know how to marshal embedded linked objects (via a field of type `lima.fields.Embed`)

- You now know how to marshal references to linked objects (via a field of type `lima.fields.References`)

- You know about lazy evaluation of linked schemas and how to specify those via qualified and fully module-qualified names.

- You know how to implement two-way relationships between objects (pass `exclude` or `only` to the linked schema through `lima.fields.Embed`)

- You know how to marshal linked collections of objects (pass `many=True` to the linked schema through `lima.fields.Embed`)

3.6 Advanced Topics

3.6.1 Automated Schema Definition

Validating ORM agnosticism for a moment, let's see how we could utilize `__lima_args__['include']` to create our Schema automatically.

We start with this SQLAlchemy model (skip this section if you don't want to install SQLAlchemy):

```python
import sqlalchemy as sa
from sqlalchemy.ext.declarative import declarative_base


Base = declarative_base()


class Account(Base):
    __tablename__ = 'accounts'
    id = sa.Column(sa.Integer, primary_key=True)
    login = sa.Column(sa.String)
    password_hash = sa.Column(sa.String)
```

`lima.fields` defines a mapping `lima.fields.TYPE_MAPPING` of some Python types to field classes. We can utilize this as follows:

```python
from lima import fields


def fields_for_model(model):
    result = {}
    for name, col in model.__mapper__.columns.items():
        field_class = fields.TYPE_MAPPING[col.type.python_type]
        result[name] = field_class()
    return result
```

Defining lima schemas becomes a piece of cake now:

```python
from lima import Schema


class AccountSchema(Schema):
    __lima_args__ = {'include': fields_for_model(Account)}

dict(AccountSchema.__fields__)
# {'id': <lima.fields.Integer at 0x...>,
#  'login': <lima.fields.String at 0x...>,
#  'password_hash': <lima.fields.String at 0x...>}
```

... and of course you still can manually add, exclude or inherit anything you like.

> **Warning:** Neither `lima.fields.TYPE_MAPPING` nor the available field classes are as exhaustive as they should be. Expect above code to fail on slightly exotic column types. There is still work to be done.

3.6.2 Field Name Mangling

Fields specified via `__lima_args__['include']` can have arbitrary names. Fields provided via class attributes have a drawback: class attribute names have to be valid Python identifiers.

lima implements a simple name mangling mechanism to allow the specification of some common non-Python-identifier field names (like JSON-LD's `"@id"`) as class attributes.

The following table shows how name prefixes will be replaced by lima when specifying fields as class attributes (note that every one of those prefixes ends with a double underscore):

name prefix	replacement
`'at__'`	`'@'`
`'dash__'`	`'-'`
`'dot__'`	`'.'`
`'hash__'`	`'#'`
`'plus__'`	`'+'`
`'nil__'`	`''` (the emtpy String)

This enables us to do the following:

```python
class FancyFieldNamesSchema(Schema):
    at__foo = fields.String(attr='foo')
    hash__bar = fields.String(attr='bar')
    nil__class = fields.String(attr='cls')   # Python Keyword

list(FancyFieldNamesSchema.__fields__)
# ['@foo', '#bar', 'class']
```

> **Note:** When using field names that aren't Python identifiers, lima obviously can't look for attributes with those same names, so make sure to specify explicitly how the data for these fields should be determined (see *How a Field gets its Data*).

3.6.3 Advanced Topics Recap

- You are now able to create schemas automatically (`__lima_args__['include']` with some model-specific code).
- You can specify a field named `'@context'` as a schema class attribute (using field name mangling: `'at__context'`).

3.7 The lima API

Please note that the lima API uses a relatively uncommon feature of Python 3: *Keyword-only arguments.*

Keyword-only arguments

Keyword-only arguments can be recognized by their position in a method/function signature: Every argument coming after the varargs argument like `*args` (or after a single `*`) is a keyword-only argument.

A function that is defined as `def foo(*, x, y):  pass` *must* be called like this: `foo(x=1, y=2)`; calling `foo(1, 2)` will raise a `TypeError`.

It is the author's opinion that enforcing keyword arguments in the right places makes the resulting code more readable.

For more information about keyword-only arguments, see PEP 3102

3.7.1 lima.abc

Abstract base classes for fields and schemas.

class `lima.abc.`**`FieldABC`**
> Abstract base class for fields.
>
> Being an instance of `FieldABC` marks a class as a field for internal type checks. You can use this class to implement your own type checks as well.
>
> ---
> **Note:** To create new fields, it's a better Idea to subclass `lima.fields.Field` directly instead of implementing FieldABC on your own.
>
> ---

class `lima.abc.`**`SchemaABC`**
> Abstract base class for schemas.
>
> Being an instance of `SchemaABC` marks a class as a schema for internal type checks. You can use this class to implement your own type checks as well.
>
> ---
> **Note:** To create new schemas, it's a *way* better Idea to subclass `lima.schema.Schema` directly instead of implementing SchemaABC on your own.
>
> ---

3.7.2 lima.exc

The lima exception hierarchy.

Note: Currently this module only holds Exceptions related to `lima.registry`, but this might change in the future.

---

exception `lima.exc.`**`AmbiguousClassNameError`**
> Raised when asking for a class with an ambiguous name.
>
> Usually this is the case if two or more classes of the same name were registered from within different modules, and afterwards a registry is asked for one of those classes without specifying the module in the class name.

exception `lima.exc.`**`ClassNotFoundError`**
> Raised when a class was not found by a registry.

exception `lima.exc.`**`RegisterLocalClassError`**
> Raised when trying to register a class defined in a local namespace.

exception `lima.exc.`**`RegistryError`**
> The base class for all registry-related exceptions.

3.7.3 lima.fields

Field classes and related code.

`lima.fields.`**`TYPE_MAPPING =dict(...)`**
> dict() -> new empty dictionary dict(mapping) -> new dictionary initialized from a mapping object's
>
>> (key, value) pairs
>
> **dict(iterable) -> new dictionary initialized as if via:** d = { } for k, v in iterable:
>
>> d[k] = v
>
> **dict(**kwargs) -> new dictionary initialized with the name=value pairs** in the keyword argument list. For
>> example: dict(one=1, two=2)

class `lima.fields.`**`Boolean`**(**, attr=None, key=None, get=None, val=None*)
> A boolean field.
>
> currently this class has no additional functionality compared to `Field`. Nevertheless it should be used over
> `Field` when referencing boolean values as an indicator for a field's type and to keep code future-proof.

class `lima.fields.`**`Date`**(**, attr=None, key=None, get=None, val=None*)
> A date field.
>
> **static** **`pack`**(*val*)
>> Return a string representation of `val`.
>>
>>> **Parameters** **val** – The :class: *datetime.date* object to convert.
>>>
>>> **Returns** The ISO 8601-representation of `val` (`YYYY-MM-DD`).

class `lima.fields.`**`DateTime`**(**, attr=None, key=None, get=None, val=None*)
> A DateTime field.
>
> **static** **`pack`**(*val*)
>> Return a string representation of `val`.
>>
>>> **Parameters** **val** – The :class: *datetime.datetime* object to convert.
>>>
>>> **Returns** The ISO 8601-representation of `val` (`YYYY-MM-DD%HH:MM:SS.mmmmmm+HH:MM`
>>> for `datetime.datetime` objects with Timezone information and microsecond preci-
>>> sion).

class `lima.fields.`**`Decimal`**(**, attr=None, key=None, get=None, val=None*)
> A decimal field.
>
> Decimal values get serialized as strings, this way, no precision is lost.

class `lima.fields.`**`Embed`**(**, schema, attr=None, key=None, get=None, val=None, **kwargs*)
> A Field to embed linked objects.
>
> **Parameters**
>
> - **schema** – The schema of the linked object. This can be specified via a schema *object,* a
>   schema *class* or the qualified *name* of a schema class (for when the named schema has not
>   been defined at the time of instantiation. If two or more schema classes with the same name
>   exist in different modules, the schema class name has to be fully module-qualified (see the
>   *entry on class names* for clarification of these concepts). Schemas defined within a local
>   namespace can not be referenced by name.
>
> - **attr** – See :class: *Field*.
>
> - **key** – See :class: *Field*.

- **get** – See :class: *Field*.

- **val** – See :class: *Field*.

- **kwargs** – Optional keyword arguments to pass to the :class: *Schema*'s constructor when the time has come to instance it. Must be empty if schema is a `lima.schema.Schema` object.

Examples:

```
# refer to PersonSchema class
author = Embed(schema=PersonSchema)

# refer to PersonSchema class with additional params
artists = Embed(schema=PersonSchema, exclude='email', many=True)

# refer to PersonSchema object
author = Embed(schema=PersonSchema())

# refer to PersonSchema object with additional params
# (note that Embed() itself gets no kwargs)
artists = Embed(schema=PersonSchema(exclude='email', many=true))

# refer to PersonSchema per name
author = Embed(schema='PersonSchema')

# refer to PersonSchema per name with additional params
author = Embed(schema='PersonSchema', exclude='email', many=True)

# refer to PersonSchema per module-qualified name
# (in case of ambiguity)
author = Embed(schema='project.persons.PersonSchema')

# specify attr name as well
user = Embed(attr='login_user', schema=PersonSchema)
```

pack(*val*)

Return the marshalled representation of val.

> **Parameters val** – The linked object to embed.

> **Returns** The marshalled representation of `val` (or `None` if `val` is `None`).

Note that the return value is determined using an (internal) dump fields *function* of the associated schema object. This means that overriding the associated schema's `dump()` *method* has no effect on the result of this method.

class `lima.fields.`**Field**(***, *attr=None*, *key=None*, *get=None*, *val=None*)

Base class for fields.

> **Parameters**

- **attr** – The optional name of the corresponding attribute.

- **key** – The optional name of the corresponding key.

- **get** – An optional getter function accepting an object as its only parameter and returning the field value.

- **val** – An optional constant value for the field.

New in version 0.3: The `val` parameter.

`attr`, `key`, `get` and `val` are mutually exclusive.

When a `Field` object ends up with two or more of the attributes `attr`, `key`, `get` and `val` regardless (because one or more of them are implemented at the class level for example), `lima.schema.Schema.dump()` tries to get the field's value in the following order: `val` `get` `key` and finally `attr`.

If a `Field` object ends up with none of these attributes (not at the instance and not at the class level), `lima.schema.Schema.dump()` tries to get the field's value by looking for an attribute of the same name as the field has within the corresponding `lima.schema.Schema` instance.

class lima.fields.**Float**(***, *attr=None*, *key=None*, *get=None*, *val=None*)
A float field.

currently this class has no additional functionality compared to `Field`. Nevertheless it should be used over `Field` when referencing float values as an indicator for a field's type and to keep code future-proof.

class lima.fields.**Integer**(***, *attr=None*, *key=None*, *get=None*, *val=None*)
An integer field.

currently this class has no additional functionality compared to `Field`. Nevertheless it should be used over `Field` when referencing integer values as an indicator for a field's type and to keep code future-proof.

class lima.fields.**Reference**(***, *schema*, *field*, *attr=None*, *key=None*, *get=None*, *val=None*, ***kwargs*)
A Field to reference linked objects.

> **Parameters**
>
> - **schema** – A schema for the linked object (see :class: *Embed* for details on how to specify this schema). One field of this schema will act as reference to the linked object.
>
> - **field** – The name of the field to act as reference to the linked object.
>
> - **attr** – see :class: *Field*.
>
> - **key** – see :class: *Field*.
>
> - **get** – see :class: *Field*.
>
> - **val** – see :class: *Field*.
>
> - **kwargs** – see :class: *Embed*.

> **pack**(*val*)
> Return value of reference field of marshalled representation of val.
>
> > **Parameters** **val** – The nested object to get the reference to.
> >
> > **Returns** The value of the reference-field of the marshalled representation of val (see `field` argument of constructor) or `None` if `val` is `None`.
>
> Note that the return value is determined using an (internal) dump field *function* of the associated schema object. This means that overriding the associated schema's `dump()` *method* has no effect on the result of this method.

class lima.fields.**String**(***, *attr=None*, *key=None*, *get=None*, *val=None*)
A string field.

currently this class has no additional functionality compared to `Field`. Nevertheless it should be used over `Field` when referencing string values as an indicator for a field's type and to keep code future-proof.

3.7.4 lima.schema

Schema class and related code.

---

class `lima.schema.`**`Schema`**(***, *exclude=None*, *only=None*, *include=None*, *ordered=False*, *many=False*)
  Base class for Schemas.

  > **Parameters**
  >
  > - **exclude** – An optional sequence of field names to be removed from the fields of the new `Schema` instance. If only one field is to be removed, it's ok to supply a simple string instead of a list containing only one string for `exclude`. `exclude` may not be specified together with `only`.
  >
  > - **only** – An optional sequence of the names of the only fields that shall remain for the new `Schema` instance. If just one field is to remain, it's ok to supply a simple string instead of a list containing only one string for `only`. `only` may not be specified together with `exclude`.
  >
  > - **include** – An optional mapping of field names to fields to additionally include in the new Schema instance. Think twice before using this option - most of the time it's better to include fields at class level rather than at instance level.
  >
  > - **ordered** – An optional boolean indicating if the :meth: *Schema.dump* method should output `collections.OrderedDict` objects instead of simple `dict` objects. Defaults to `False`. This does not influence how nested fields are serialized.
  >
  > - **many** – An optional boolean indicating if the new Schema will be serializing single objects (`many=False`) or collections of objects (`many=True`) per default. This can later be overridden in the `dump()` Method.

  New in version 0.3: The `include` parameter.

  New in version 0.3: The `ordered` parameter.

  Upon creation, each Schema object gets an internal mapping of field names to fields. This mapping starts out as a copy of the class's `__fields__` attribute. (For an explanation on how this `__fields__` attribute is determined, see `SchemaMeta`.)

  Note that the fields themselves are not copied - changing the field of an instance would change this field for the other instances and classes referencing this field as well. In general it is *strongly* suggested to treat fields as immutable.

  The internal field mapping is then modified as follows:

  - If `include` was provided, fields specified therein are added (overriding any fields of the same name already present)

    If the order of your fields is important, make sure that `include` is of type `collections.OrderedDict` or similar.

  - If `exclude` was provided, fields specified therein are removed.

  - If `only` was provided, *all but* the fields specified therein are removed (unless `exclude` was provided as well, in which case a `ValueError` is raised.)

  Also upon creation, each Schema object gets an individually created dump function that aims to unroll most of the loops and to minimize the number of attribute lookups, resulting in a little speed gain on serialization.

  `Schema` classes defined outside of local namespaces can be referenced by name (used by `lima.fields.Nested`).

  dump(*obj*)
    Return a marshalled representation of obj.

    > **Parameters obj** – The object (or collection of objects, depending on the schema's `many` property) to marshall.

> **Returns** A representation of `obj` in the form of a JSON-serializable dict (or `collections.OrderedDict`, depending on the schema's `ordered` property), with each entry corresponding to one of the schema's fields. (Or a list of such dicts in case a collection of objects was marshalled)

Changed in version 0.4: Removed the `many` parameter of this method.

many
> Read-only property: does the dump method expect collections?

ordered
> Read-only property: does the dump method return ordered dicts?

class `lima.schema.`**`SchemaMeta`**
Metaclass of Schema.

Note: The metaclass `SchemaMeta` is used internally to simplify the configuration of new `Schema` classes. For users of the library there should be no need to use `SchemaMeta` directly.

When defining a new `Schema` (sub)class, `SchemaMeta` makes sure that the new class has a class attribute `__fields__` of type `collections.OrderedDict` containing the fields for the new Schema.

`__fields__` is determined like this:

- The `__fields__` of all base classes are copied (with base classes specified first having precedence).

  Note that the fields themselves are *not* copied - changing an inherited field would change this field for all base classes referencing this field as well. In general it is *strongly* suggested to treat fields as immutable.

- Fields (Class variables of type `lima.abc.FieldABC`) are moved out of the class namespace and into `__fields__`, overriding any fields of the same name therein.

- If present, the class attribute `__lima_args__` is removed from the class namespace and evaluated as follows:

  - Fields specified via `__lima_args__['include']` (an optional mapping of field names to fields) are inserted into `__fields__`. overriding any fields of the same name therein.

    If the order of your fields is important, make sure that `__lima_args__['include']` is of type `collections.OrderedDict` or similar.

    New fields in `__lima_args__['include']__` are inserted at the position where `__lima_args__` is specified in the class.

  - Fields named in an optional sequence `__lima_args__['exclude']` are removed from `__fields__`. If only one field is to be removed, it's ok to supply a simple string instead of a list containing only one string. `__lima_args__['exclude']` may not be specified together with `__lima_args__['only']`.

  - If in an optional sequence `__lima_args__['only']` is provided, *all but* the fields mentioned therein are removed from `__fields__`. If only one field is to remain, it's ok to supply a simple string instead of a list containing only one string. `__lima_args__['only']` may not be specified together with `__lima_args__['exclude']`.

    Think twice before using `__lima_args__['only']` - most of the time it's better to rethink your Schema than to remove a lot of fields that maybe shouldn't be there in the first place.

New in version 0.3: Support for `__lima_args__['only']`.

`SchemaMeta` also makes sure the new Schema class is registered with the lima class registry `lima.registry` (at least if the Schema isn't defined inside a local namespace, where we wouldn't find it later on).

> **classmethod __prepare__** (*metacls*, *name*, *bases*)
>   Return an OrderedDict as the class namespace.
>
>   This allows us to keep track of the order in which fields were defined for a schema.

3.8 Project Info

lima was started in 2014 by Bernhard Weitzhofer.

3.8.1 Acknowledgements

lima is heavily inspired by marshmallow, from which it lifts most of its concepts from.

Note: The key differences between lima and marshmallow are (from my, Bernhard's point of view):

- marshmallow supports Python 2 as well, lima is Python 3 only.

- marshmallow has more features, foremost among them deserialization and validation.

- *Skipping validation* and doing internal stuff differently, lima is (at the time of writing this) noticeably faster.

Although greatly inspired by marshmallow's API, the lima API differs from marshmallow's. lima is not a drop-in replacement for marshmallow and it does not intend to become one.

The lima sources include a copy of the Read the Docs Sphinx Theme.

The author believes to have benefited a lot from looking at the documentation and source code of other awesome projects, among them django, morepath, Pyramid (lima.util.reify was taken from there) and SQLAlchemy as well as the Python standard library itself. (Seriously, look in there!)

3.8.2 About the Image

The Vicuña is the smallest and lightest camelid in the world. In this 1914 illustration [1], it is depicted next to its bigger and heavier relatives, the Llama and the Alpaca.

Despite its delicate frame, the Vicuña is perfectly adapted to the harsh conditions in the high alpine regions of the Andes. It is a mainly wild animal long time believed to never have been domesticated. Reports of Vicuñas breathing fire are exaggerated.

3.9 Changelog

3.9.1 0.6 (unreleased)

Note: While unreleased, the changelog of lima 0.6 is itself subject to change.

---

[1] Beach, C. (Ed.). (1914). The New Student's Reference Work. Chicago: F. E. Compton and Company (via Wikisource).

3.9.2 0.5 (2015-05-11)

- Support getting field values from an object's items by providing the `key` argument to a Field constructor.

- Add a `fields.Decimal` field type that packs `decimal.Decimal` values into strings.

- Move Tests into directory `/test`.

- Remove deprecated field `fields.Nested`. Use `fields.Embed` instead.

3.9.3 0.4 (2015-01-15)

- **Breaking Change:** The `Schema.dump` method no longer supports the `many` argument. This makes `many` consistent with `ordered` and simplifies internals.

- Improve support for serializing linked data:

  - Add new field type `fields.Reference` for references to linked objects.

  - Add new name for `fields.Nested`: `fields.Embed`. Deprecate `fields.Nested` in favour of `fields.Embed`.

- Add read-only properties `many` and `ordered` for schema objects.

- Don't generate docs for internal modules any more - those did clutter up the documentation of the actual API (the docstrings remain though).

- Implement lazy evaluation and caching of some attributes (affects methods: `Schema.dump`, `Embed.pack` and `Reference.pack`). This means stuff is only evaluated if and when really needed, but it also means:

  - The very first time data is dumped/packed by a Schema/Embed/Reference object, there will be a tiny delay. Keep objects around to mitigate this effect.

  - Some errors might surface at a later time. lima mentions this when raising exceptions though.

- Allow quotes in field names.

- Small speed improvement when serializing collections.

- Remove deprecated name `fields.type_mapping`. Use `fields.TYPE_MAPPING` instead.

- Overall cleanup, improvements and bug fixes.

3.9.4 0.3.1 (2014-11-11)

- Fix inconsistency in changelog.

3.9.5 0.3 (2014-11-11)

- Support dumping of `OrderedDict` objects by providing `ordered=True` to a Schema constructor.

- Implement field name mangling: `at__foo` becomes `@foo` for fields specified as class attributes.

- Support constant field values by providing `val` to a Field constructor.

- Add new ways to specify a schema's fields:

  - Add support for `__lima_args__['only']` on schema definition

  - Add *include* parameter to Schema constructor

This makes specifying fields on schema definition (`__lima_args__` - options *include*, *exclude*, *only*) consistent with specifying fields on schema instantiation (schema constructor args *include*, *exclude*, *only*).

- Deprecate `fields.type_mapping` in favour of `fields.TYPE_MAPPING`.

- Improve the documentation.

- Overall cleanup, improvements and bug fixes.

3.9.6 0.2.2 (2014-10-27)

- Fix issue with package not uploading to PYPI

- Fix tiny issues with illustration

3.9.7 0.2.1 (2014-10-27)

- Fix issues with docs not building on readthedocs.org

3.9.8 0.2 (2014-10-27)

- Initial release

3.10 License

Copyright (c) 2014-2015, Bernhard Weitzhofer

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

l

Symbols

A

B

C

D

E

F

I

L

M

O

P

R

S

T