
miniff

Release 0.1

miniff authors

Sep 01, 2021

CONTENTS:

1	Introduction	1
1.1	Glossary	1
2	Tutorials	5
2.1	Toy function fitting	5
2.2	Computing descriptors	10
3	Miniff Design Documentation	15
3.1	Structure	15
3.2	Deployment	16
4	Indices and tables	17

INTRODUCTION

1.1 Glossary

- **atom, point:** a vector in real (typically, 3D) space defining the location and the metadata associated with this location. This gives us information about an individual small object in space.
- **structure, cell, box:** a set of *points* enclosed into a parallelogram and representing a snapshot of atoms in media (molecular, solid, liquid, etc.). A **box** is a way to describe how multiple *atoms* share the same space. This can simply be a list of coordinates, for example. Related classes: `miniff.kernel.Cell`, `miniff.kernel.NeighborWrapper`.
- **potential, interaction, spring:** a protocol (a function) taking atomic coordinates of two or three *atoms* (optionally, matching a set of conditions) and producing a single floating-point number. For example, **potentials** may describe how strongly atoms interact with each other.
- **atomic environment:** a single *point* picked in a *structure*. **Atomic environment** is a very abstract way of telling which *interactions* are important and how *atoms* are grouped by these interactions.
- **partial energy, atomic energy, potential energy:** a sum of all *interaction* values the chosen atom participates in. The **potential energy** value may be subject to double-counting issues when a single *potential* is shared between many atoms. Related classes: `miniff.potentials.LocalPotential`, `miniff.ml.NNPPotential`.
- **machine learning (ML) potential:** a variant of the *partial energy* where the sum is replaced by a more complex process involving machine learning techniques.
- **(total) energy:** a sum of all *atomic energies* defining the cumulative energy accumulated in the *structure* and originating from attractions and repulsions of individual *atoms*. Predicting **total energy** from **structures** is the primary purpose of this package.
- **(total) energy landscape:** *total energy* as a function of one or more parameters of a *structure*. It is simply a way to look at the *total energy* as a scalar function of many variables.
- **(total) energy minimum:** a *structure* and the corresponding *total energy* minimum of the *potential landscape*. Finding *potential energy minimum* is one of the primary purposes of the *total energy* description.
- **charge:** a scalar belonging to *atomic* metadata with the properties of *partial energy*. **Charges** are used to treat long-range *potentials* which cannot be described by *atomic environments*. **Charges** are not necessary physical (Coulomb) charges: they may also be electronegativities or any other scalar property of an *atom*.
- **force, stress:** negative gradients of the *total energy* with respect to coordinates describing *structures*. *Forces* indicate the direction in a multidimensional space where the *total energy* becomes smaller.

- **descriptors**: a special sort of *potential energy* intended to describe *atomic environments*. Descriptors are typically defined by *local environments*. Unlike *potentials*, **descriptor** functions usually have a simple smooth form.

TUTORIALS

2.1 Toy function fitting

`miniff` is built around the idea of using neural-network fits to describe inter-atomic interactions. This is done in three basic steps:

1. constructing the dataset;
2. fitting the data (= finding optimal map parameters);
3. using the fit.

This tutorial demonstrates the first two steps in the context of simple scalar function fitting. We will pick a function and will try to perform a least-squares fit. Then, we will repeat the procedure with the help of `miniff.ml` and `miniff.ml_util` modules to introduce the key concepts of dataset construction and fitting in `miniff`.

2.1.1 The toy problem

Neural-network fitting is useful for interpolating black-box functions which otherwise require intensive computations. For pedagogical reasons, let us instead try to fit a simple analytic function defined for positive r and resembling some features of realistic interatomic potentials:

$$f(r) = (r^2 - r + 1) e^{-r}.$$

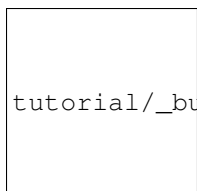
Let's first plot the function:

```
import numpy as np
from matplotlib import pyplot

def f(r):
    return (r ** 2 - r + 1) * np.exp(-r)

r = np.linspace(0.1, 10, 100)
f_r = f(r)
pyplot.plot(r, f_r)
```

```
[<matplotlib.lines.Line2D at 0x7f37f8683e50>]
```



tutorial/_build/jupyter_execute/tutorial/fitting_1d_function_0_1.png

The function has two extrema at $r = 1, 2$ and decays rapidly: these are the features which we will attempt to reproduce.

When fitting, we only have access to discretized sampling values f_{-r} , but not to the functional form $f(r)$. In principle, we may still use $f(r)$ as a black box to produce more sampling points but we consider this operation to be computationally expensive, thus, to only be used *before* the fitting.

2.1.2 1. Constructing the dataset

The fitting data (or the dataset) includes two pieces of information: blackbox outputs and descriptors. They can be also seen as right-hand and left-hand sides of the map $g(r) \rightarrow f(r)$ respectively, where $f(r)$ are previously computed blackbox function values f_{-r} .

What are descriptors?

The purpose of descriptors $g(r)$ is to describe values of r where $f(r)$ was computed in a physically reasonable way. Can we use values of $g(r)=r$ directly? Yes, but $g(r)$ play the role of pre-conditioners: we input some intuition about how $f(r)$ would most reasonably look like while leaving particular details to be determined during the fitting process.

Additionally, we allow the blackbox function to have an arbitrary number of inputs such as

$$f(\vec{r}) = \sum_i (r_i^2 - r_i + 1) e^{-r_i}.$$

In this case we may want to construct the dataset from vectors of a non-constant dimension $r = ([0], [0.1, 0.2], [2.2, 3, 2], \dots)$. To employ neural network machinery, however, we usually need a constant number of inputs. This can be achieved by using a descriptor of the form

$$g(\vec{r}) = \sum_i r_i$$

in place of r . Importantly, the above descriptor also respect the permutation symmetry of $f(r)$: $f(x, y) = f(y, x)$, $g(x, y) = g(y, x)$. Thus, the follow-up fitting process does not need to deduce this symmetry in a hard way numerically.

How to choose descriptors?

- $g(r)$ should include as much intuition as possible (symmetries, asymptotic behavior, ...);
- $g(r)$ should be simple enough functions fast to compute.

Here we choose the following family of descriptors which obviously satisfies both recommendations:

$$g_i(r) = r^i e^{-r}.$$

We also know the exact form of the map to be fitted:


$$f(r) = g_2(r) - g_1(r) + g_0(r).$$

So, to complete the dataset we compute g_{-0} , g_{-1} , g_{-2} at previously chosen sampling points.

```
def g(r, i):
    return r ** i * np.exp(-r)

g_r = np.array(tuple(g(r, i) for i in range(3)))
pyplot.plot(r, f_r, label="target function")
for i, _g in enumerate(g_r):
    pyplot.plot(r, _g, ls="--", label=f"g_{i}")
pyplot.legend()
```

```
<matplotlib.legend.Legend at 0x7f37f8ad2410>
```



tutorial/_build/jupyter_execute/tutorial/fitting_1d_function_1_1.png

2.1.3 2. Fitting the data

For this example we know there exists a linear combination of descriptors perfectly reproducing the black-box function

$$f(r) = \sum_i a_i \cdot g_i(r).$$

Thus, the least-squares fit will find map parameters a_i .

```
gf_map, res, rank, s = np.linalg.lstsq(g_r.T, f_r)
print(gf_map)

from numpy import testing
testing.assert_allclose(gf_map, (1, -1, 1))
```

```
[ 1. -1.  1.]
```

According to the last formula in the previous section, the answer is $[1, -1, 1]$.

2.1.4 miniff setting

The above solution can be also obtained from miniff with the same workflow.

1. Constructing the dataset (miniff)

miniff.ml contains handy pipelines for dataset construction and pre-processing. PerCellDataset and PerPointDataset are the two key containers where descriptors and target function values are stored. Some data pieces (such as descriptors, but also atomic charges and more) can be assigned to individual points (atoms) and are stored in PerPointDataset. Other data pieces (such as total interatomic energies) are cumulative values that are stored in PerCellDataset. Due to the nature of force-field fitting, there can be only one PerCellDataset but one or more "PerPointDataset"s.

In this case we store f_r into the total interaction energy slot of PerCellDataset to benefit from the energy fitting pipeline.

```
import torch
from miniff import ml

rhs = ml.PerCellDataset(
    energy=torch.tensor(f_r, dtype=torch.float32)[: , None],
)
```

All dataset pieces have to be torch tensors of a fixed shape. In this case “energy” is required to be a column array with the second dimension equal to 1: otherwise a `ValueError` will be raised.

Descriptors are stored in `PerPointDataset` as a standard.

```
lhs = ml.PerPointDataset(
    features=torch.tensor(g_r.T, dtype=torch.float32)[: , None, :],
    mask=torch.ones(len(r), 1),
)
```

Here, we additionally specify the mask which weights individual data points. Again, the dimensions of input tensors are well-defined. As a final step, we collect the two pieces of fitting data into the `Dataset` object.

```
dataset = ml.Dataset(rhs, lhs)
```

At this point all tensor shapes are checked and we are ready to set up the fitting.

2. Fitting the data (miniff)

First, we need to define the functional form of the map. In this case the map is linear which is supported by miniff out of the box through `ml.SequentialSoleEnergyNN` torch module.

$$f = W \cdot g,$$

where W stands for weights.

```
nn = ml.SequentialSoleEnergyNN(3, n_layers=1, bias=False)
```

Here, 3 stands for descriptor count, `n_layers` specifies the count of linear layers in the neural network (1 means the linear fit of the above form), `bias=False` tells that no constant is added to the linear form.

The `SequentialSoleEnergyNN` object contains all information needed to optimize and to reproduce the map in the context of machine learning. In this case the object wraps 3 previously mentioned parameters `a_i` to be optimized.

To obtain the best fit we use energy fit optimisation pipeline defined in `miniff.ml_util`. First, we define the closure to be energy difference closure. In general, closure defines the loss function to optimize but the corresponding `ml_util` object also includes default optimization settings.

```
from miniff import ml_util

tolerance = dict(
    tolerance_change=1e-14,
    tolerance_grad=1e-14,
)

closure = ml_util.simple_energy_closure([nn], dataset,
    optimizer_kwargs=tolerance)
```

Let’s check that the initial (random) guess is non-optimal.

```
print("Initial loss:", closure.loss().loss_value.item())
```

```
Initial loss: 0.09587489813566208
```

Finally, let's run the optimization and check the loss function value becomes small.

```
closure.optimizer_init()
closure.optimizer_step() # LBFGS algorithm by default
loss = closure.last_loss.loss_value.item()
print("Final loss:", loss)
assert loss < 1e-8
```

```
Final loss: 5.424805570031221e-16
```

Parameters are now stored in the neural network module.

```
weights = nn[0].weight.detach().squeeze().numpy()
print(weights)
testing.assert_allclose(weights.squeeze(), [1, -1, 1], atol=1e-3)
```

```
[ 1.          -0.9999999  1.          ]
```

Checking the result

Finally, let's check whether the fit is reasonable by plotting it.

```
def nn_fit(r, descriptor_fns, nn):
    # Compute descriptors
    descriptors = np.array(tuple(g(r) for g in descriptor_fns))
    # Compute the map
    result = nn(torch.tensor(descriptors.T, dtype=torch.float32))
    return result.detach().numpy().squeeze()

from functools import partial

f_r_test = nn_fit(
    r,
    tuple(partial(g, i=i) for i in range(3)),
    nn,
)
pyplot.plot(r, f_r)
pyplot.scatter(r, f_r_test)
```

```
<matplotlib.collections.PathCollection at 0x7f37a8223910>
```



The function `nn_fit` demonstrates how to propagate the map forward: one needs to (1) compute descriptors as a function of input `r` and (2) to feed them to the neural network.

2.1.5 Advanced

It does not make sense to use anything but linear map for this particular problem with a known exact solution. However, if we chose to, for example, reduce the descriptors to `g_0` only or to chose different descriptors then the linear map becomes a crude approximation. In this case one typically switches to a more general neural-network setting with non-linear activation layers helping to map various curvature features in the dataset. One can use `torch` machinery for this or just otherwise increase `n_layers` when initializing the map. With `n_layers=2`, for example, one sigmoid activation layer is inserted between two linear layers.

The use of `ml_util` functions is entirely optional. They aim to provide default optimization setting to start with: loss functions (2-norm), optimizers (LBFGS), descriptor sets (Behler recipe). You are free to modify and replace these components at the point when you feel that the default choice is no longer optimal.

2.2 Computing descriptors

Computing descriptors is an important part of both constructing datasets and using neural-network fits.

2.2.1 The problem

Consider a unit box in 3 dimensions. $n = 100$ atoms are present in the box; their cartesian coordinates are written in a $[100 \times 3]$ matrix `r` such that $0 \leq r \leq 1$.

```
import numpy as np

n = 100
a = 0.3
r = np.random.rand(100, 3)
```

Our task is to count atomic neighbors inside a shell of size $a = 0.3$ for each atom. This is what a typical descriptor looks like: it counts neighbors in shells around each atom.

To make a continuous function out of neighbor count let's introduce a kernel function of neighbor distance r smoothly decaying from a unit value at $r = 0$ down to zero at cutoff value $r = a$. Let's take a half-period of cosine function for this example.

$$g(r) = \frac{1 + \cos \frac{\pi r}{a}}{2} \quad r < a$$

```
def g(r):
    return (1 + np.cos(np.pi * r / a)) / 2 * (r < a)

from matplotlib import pyplot
x = np.linspace(0, a)
pyplot.plot(x, g(x))
```

```
[<matplotlib.lines.Line2D at 0x7f3ffbe4ded0>]
```

tutorial/_build/jupyter_execute/tutorial/computing_descriptors_1_1.png

Instead of computing neighbors directly let's compute the cumulative value of the kernel function for all neighbors.

$$n_i = \sum_j g(r_{ij})$$

2.2.2 Computing naively

We start with computing all pair distances between atoms.

```
mdist = np.linalg.norm(r[:, None, :] - r[None, :, :], axis=-1)
```


Then, we reduce pair distances using the kernel function.

```
result = g(mdist).sum(axis=0) - 1
```

Finally, we can plot the distribution of neighbors across all points.

```
from matplotlib import pyplot
pyplot.hist(result)
```

```
(array([10., 12., 8., 20., 11., 11., 15., 7., 3., 3.]),
 array([0.10492551, 0.47059291, 0.8362603 , 1.2019277 , 1.5675951 ,
        1.93326249, 2.29892989, 2.66459729, 3.03026468, 3.39593208,
        3.76159948])),
 <BarContainer object of 10 artists>)
```



tutorial/_build/jupyter_execute/tutorial/computing_descriptors_4_1.png

This approach is straightforward but not efficient. First, matrix `mdist` includes information about very distant atoms we simply discard. Second, we can only perform pair reductions this way: counting atomic triples within the same sphere of radius `a` will require a tensor with three dimensions raising memory requirements by a factor of `n = 100`. Instead, `miniff` uses a more efficient approach with a quasi-linear complexity scaling.

2.2.3 Computing neighbors with miniff

`miniff` provides an interface to computing neighbors efficiently. It uses `KDTree` from `scipy` to compute neighbor lists and `cython` functions to iterate over the list and to reduce it to the desired quantities. As such, there are two key steps.

1. Compute neighbor lists

The primary purpose of `miniff.kernel` module is to perform neighbor list construction with `NeighborWrapper` class. To do it, we first construct a unit cell object with atomic coordinates.

```
from miniff.kernel import Cell, NeighborWrapper

cell = Cell(np.eye(3), r, ['a'] * len(r))
```

The second step is to define the wrapper by providing coordinates and periodicity information.

```
nw = NeighborWrapper(cell, (1, 1, 1))
```

In this case `(1, 1, 1)` indicates a box with no periodic replica. Finally, neighbor computations are invoked separately by providing the neighbor cutoff value to method `compute_distances(cutoff)`.

```
nw.compute_distances(a)
print(len(nw.sparse_pair_distances.data))
```

```
742
```

As a result, `nw.sparse_pair_distances` is populated with neighbor distance data.

2. Reduce

In principle, neighbor counts can be deduced from the sparse matrix `nw.sparse_pair_distances` directly with previously defined method `g`. A more conventional workflow is to pick a reduction function from `miniff.potentials` module. Specifically, `miniff.potentials.sigmoid_descriptor_family` provides the following reduction function:

$$n_i = \sum_j \frac{1}{1 + e^{\frac{r_{ij} - r_0}{dr}}} \cdot \frac{1 + \cos \frac{\pi r_{ij}}{a}}{2},$$

where `r_ij` is the distance matrix, `r0`, `dr` and `a` are constant parameters. By setting `dr` to a small value and `r0` to `2a` we simplify the expression to the second term only

$$n_i \approx \sum_j \frac{1 + \cos \frac{\pi r_{ij}}{a}}{2},$$

resembling the kernel function defined previously. To instantiate the descriptor we simply call the factory object the specified parameters.

```
from miniff.potentials import sigmoid_descriptor_family

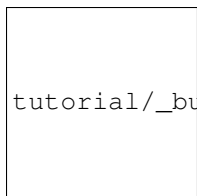
descriptor = sigmoid_descriptor_family(r0=10 * a, dr=a / 10, a=a)
```

Let's visualise the radial form of the descriptor.

```
from miniff.presentation import plot_potential_2

plot_potential_2(descriptor, energy_unit="1", length_unit="1")
```

```
<AxesSubplot:xlabel='r (1)', ylabel='Energy (1) '>
```

tutorial/_build/jupyter_execute/tutorial/computing_descriptors_9_1.png

To perform the reduction we simply feed the descriptor to NeighborWrapper.

```
result_miniff = nw.eval(descriptor.copy(tag="a-a"), "kernel")
```

Not that `descriptor.copy(tag="a-a")` makes a copy of the descriptor with the `descriptor.tag = "a-a"`. The tag is required by NeighborWrapper to distinguish between species and possible pairs. For example, having two kinds of species in the box "a" and "b" we may reduce several kinds of pairs such as "a-a" or "a-b".

Additionally, we specify the function we want to compute. The same descriptor may provide multiple functions at once: "kernel" stands for the descriptor value while "kernel_gradient" efficiently computes descriptor gradients.

Finally, let's plot and compare the result.

```
_, bins, _ = pyplot.hist(result)
pyplot.hist(result_miniff, bins, histtype="step")
```

```
(array([10., 12., 8., 20., 11., 11., 15., 7., 3., 3.]),
 array([0.10492551, 0.47059291, 0.8362603 , 1.2019277 , 1.5675951 ,
        1.93326249, 2.29892989, 2.66459729, 3.03026468, 3.39593208,
        3.76159948]),
 [matplotlib.patches.Polygon at 0x7f3f9caf9790])
```



tutorial/_build/jupyter_execute/tutorial/computing_descriptors_11_1.png

2.2.4 Advanced

`miniff.potentials` includes several most common reduction functions. But it is also possible to define yours using `miniff.potentials.general_pair_potential_family`. It accepts two methods: `f(r)` and `df_dr(r)`. As `f` is a python function called multiple times during the reduction, it will not benefit from cython optimizations and parallelism. Nevertheless, it is a handy prototyping tool.

MINIFF DESIGN DOCUMENTATION

3.1 Structure

The project includes several modules performing energy and gradients computations, potential parameter optimization through machine learning, simple geometry optimization and presentation utilities.

- `potentials`: a core module implementing smooth classical interatomic potentials and descriptors. The module includes `LocalPotentialFamily`: a factory for constructing parameterized `LocalPotential` objects which, in turn, include all necessary data and interfaces to compute atomic energies and gradients. Several pre-built potential forms are provided through instantiating `LocalPotentialFamily`: for example, `lj_potential_family` (Lenard-Jones pair potential) or `behler5_descriptor_family` (Behler type 5 angular descriptor).
 - `kernel`: implements a `NeighborWrapper` class which computes neighbor information from atomic coordinate data and prepares contiguous data buffers which can be processed by `LocalPotential`'s. `NeighborWrapper` is implemented with 3D periodic boundary conditions in mind to model amorphous materials. However, it also supports molecular systems without any overhead. `NeighborWrapper` also implements a key interface `NeighborWrapper.eval` which, for a given structure and a list of potentials, computes total energy and gradients. `NeighborWrapper.relax` introduces a toy interface for structural relaxation using `scipy` minimizers.
 - `ml`: implements machine learning potentials.
 - `ml.Dataset` is the key container for atomic descriptors, energies and gradients. It ensures that all dataset pieces are `torch.Tensor`'s compatible with each other. `ml.Dataset` includes two large blocks of data, namely one `ml.PerCellDataset` block with target energies and energy gradients, and one or more `ml.PerPointDataset`'s with descriptor information.
 - `ml.Normalization` implements normalization of datasets in a physically reasonable way.
 - `ml.learn_cauldron` is a typical entry point for dataset creation which includes reasonable default values.
 - `ml.SequentialSoleEnergyNN` is Behler et al. suggestion for the neural network potential form.
 - `ml.forward_cauldron` is the core routine for machine-learning optimization. It combines dataset and neural-network models to produce the energy and gradients prediction.
 - `ml.NNPotential` is a neural-network potential subclassing `potentials.LocalPotential`.
- `miniff.ml` is built around `pytorch`.
- `ml_util`: includes reasonable recipes for optimizing neural networks from `ml`.
 - `ml_util.simple_energy_closure` provides defaults for running the optimization with *LBFGS*.

- `ml_util.*Workflow` are workflow classes for optimizing neural-network potentials. These classes accumulate and re-distribute many parameters related to the dataset organization, potential form and optimization process.
- `presentation`: various handy plotting routines to present potentials and visualize machine learning optimization process.

3.2 Deployment

miniff can be deployed on high-performance computing (HPC) clusters.

3.2.1 Parallelism

- *miniff* takes a full advantage of GPU parallelism in `pytorch`. Please note that it is often not enough to install pre-built bundles of `pytorch` as they support only a limited set of (very recent) GPU drivers. If your HPC hardware does not feature those you have several options:
 - It is best to ask your HPC support team for a suitable `pytorch` build specifically for the HPC machine. Such builds may be available through `module` script or other ways to manage the runtime environment on the cluster: please investigate such options first.
 - The second possibility is to use an older `pytorch` version which bundles kernels for older GPUs. *miniff* does its best to support a wide range of `pytorch` versions but you have to test the compatibility manually in your case.
 - The last possibility is to build `pytorch` manually. This is the most tedious approach, thus, not recommended for unexperienced users.
- OpenMP threading support is present in potential and gradient computations. This may be useful for computing energies and gradients in large atomic systems. The number of threads is controlled by usual means such as `OMP_NUM_THREADS` environment variable. For small atomic systems ~100 atoms up to 2-4 threads are beneficial: make sure your parallel cluster setup is reasonable.

INDICES AND TABLES

- `genindex`
- `modindex`
- `search`