

---

# **stochqn Documentation**

**David Cortes**

**Apr 15, 2021**



---

## Contents:

---

<b>1</b>	<b>Indices and tables</b>	<b>17</b>
	<b>Python Module Index</b>	<b>19</b>
	<b>Index</b>	<b>21</b>



```
class SQN (x0, grad_fun, obj_fun=None, hess_vec_fun=None, pred_fun=None, batches_per_epoch=25,
            step_size=0.001, decr_step_size='auto', shuffle_data=True, random_state=1, nepochs=25,
            valset_frac=None, tol=0.1, callback_epoch=None, callback_iter=None, kwargs_cb={},
            verbose=True, mem_size=10, bfgs_upd_freq=20, min_curvature=0.0001, y_reg=None,
            use_grad_diff=False, check_nan=True, nthreads=-1, use_float=False)
SQN optimizer
```

Optimizes an empirical (convex) loss function over batches of sample data.

### Parameters

- **x0** (*array (m, )*) – Initial values of the variables to optimize (referred hereafter as ‘x’).
- **grad\_fun** (*function(x, X, y, sample\_weight, \*\*kwargs) -> array(m, )*) – Function that calculates the empirical gradient at values ‘x’ on data ‘X’ and ‘y’. Note: output must be one-dimensional and with the same number of entries as ‘x’, otherwise the Python session might segfault. (The extra keyword arguments are passed in the ‘fit’ method, not here)
- **obj\_fun** (*function(x, X, y, sample\_weight, \*\*kwargs) -> float*) – Function that calculates the empirical objective value at values ‘x’ on data ‘X’ and ‘y’. Only used when using a validation set (‘valset\_frac’ not None, or ‘valset’ passed to fit). Ignored when fitting the data in user-provided batches. (The extra keyword arguments are passed in the ‘fit’ method, not here)
- **hess\_vec\_fun** (*function(x, vec, X, y, sample\_weight, \*\*kwargs) -> array(m, )*) – Function that calculates the product of a vector the empirical Hessian at values ‘x’ on data ‘X’ and ‘y’. Ignored when using ‘use\_grad\_diff=True’. Note: output must be one-dimensional and with the same number of entries as ‘x’, otherwise the Python session might segfault. These products are calculated on a larger batch than the gradients (given by batch\_size \* bfgs\_upd\_freq). (The extra keyword arguments are passed in the ‘fit’ method, not here)
- **pred\_fun** (*None or function(xopt, X)*) – Prediction function taking as input the optimal ‘x’ values as obtained by the optimization procedure, and new observation ‘X’ on which to make predictions. If passed, will have an additional method oLBFGS.predict(X, \*args) that calls this function with current values of ‘x’.
- **batches\_per\_epoch** (*int*) – Number of batches per epoch (each batch will have the same number of observations except for the last one which might be smaller).
- **step\_size** (*float*) – Initial step size to use. (Can be modified after object is already initialized)
- **decr\_step\_size** (*str “auto”, None, or function(initial\_step\_size, epoch) -> float*) – Function that determines the step size during each epoch, taking as input the initial step size and the epoch number (starting at zero). If “auto”, will use 1/sqrt(iteration). If None, will use constant step size. For ‘partial\_fit’, it will take as input the number of iterations of the algorithm rather than epoch, so it’s very recommended to provide a custom function when passing data in user-provided batches. Can be modified after the object has been initialized (oLBFGS.decr\_step\_size)
- **shuffle\_data** (*bool*) – Whether to shuffle the data at the beginning of each epoch.
- **random\_state** (*int*) – Random seed to use for shuffling data and selecting validation set. The algorithm is deterministic so it’s not used for anything else.
- **nepochs** (*int*) – Number of epochs for which to run the optimization procedure. Might terminate earlier if using a validation set for monitoring.
- **valset\_frac** (*float(0, 1) or None*) – Percent of the data to use as validation set for early stopping. Can also pass a user-provided validation set to ‘fit’, in which case it will be ignored. If passing None, will run for the number of epochs passed in ‘nepochs’.

- **tol** (*float*) – If the objective function calculated on the validation set decrease by less than ‘tol’ upon completion of an epoch, will terminate the optimization procedure. Ignored when not using a validation set.
- **callback\_epoch** (*None or function\*(x, \*\*kwargs)*) – Callback function to call at the end of each epoch
- **callback\_iter** (*None or function\*(x, \*\*kwargs)*) – Callback function to call at the end of each iteration
- **kwargs\_cb** (*tuple*) – Additional arguments to pass to ‘callback’ and ‘stop\_crit’. (Can be modified after object is already initialized)
- **verbose** (*bool*) – Whether to print messages when there is some problem during an iteration (e.g. correction pair not meeting minum curvature).
- **mem\_size** (*int*) – Number of correction pairs to store for approximation of Hessian-vector products.
- **bfgs\_upd\_freq** (*int*) – Number of iterations (batches) after which to generate a BFGS correction pair.
- **min\_curvature** (*float or None*) – Minimum value of  $s^*y / s^*s$  in order to accept a correction pair.
- **y\_reg** (*float or None*) – regularizer for ‘y’ vector (gets added  $y_{reg} * s$ )
- **use\_grad\_diff** (*bool*) – Whether to create the correction pairs using differences between gradients instead of Hessian-vector products. These gradients are calculated on a larger batch than the regular ones (given by  $batch\_size * bfgs\_upd\_freq$ ).
- **check\_nan** (*bool*) – Whether to check for variables becoming NaN after each iteration, and reverting the step if they do (will also reset BFGS memory).
- **nthreads** (*int*) – Number of parallel threads to use. If set to -1, will determine the number of available threads and use all of them. Note however that not all the computations can be parallelized.
- **use\_float** (*bool*) – Whether to use C ‘float’ type (np.float32). If ‘False’ (the default), will use ‘double’ type (np.float64). The variables, gradient, and hessian-vector must be of this same dtype.

## References

**fit** (*X, y, sample\_weight=None, additional\_kwargs={}, valset=None*)

Fit model to sample data

### Parameters

- **X** (*array(n\_samples, m)*) – Sample data to which to fit the model.
- **y** (*array(n\_samples, )*) – Labels or target values for the sample data.
- **sample\_weight** (*None or array(n\_samples, )*) – Observations weights for the sample data.
- **additional\_kwargs** (*dict*) – Additional keyword arguments to pass to the objective, gradient, and Hessian-vector functions.
- **valset** (*tuple(3)*) – User-provided validation set containing (X\_val, y\_val, sample\_weight\_val). At the end of each epoch, will calculate objective function on this set, and if the decrease from the objective function in the previous epoch is below tolerance, will terminate procedure earlier. If ‘valset\_frac’ was provided and a validation set is passed,

‘valset\_frac’ will be ignored. Must provide objective function in order to use a validation set.

**Returns self** – This object.

**Return type** obj

**get\_x()**

Get a copy of current values of the variables

**Returns x** – Current variable values.

**Return type** array(n, )

**niter**

**partial\_fit** (X, y, *sample\_weight=None*, *additional\_kwargs={}*)

Update model with user-provided batches of data

---

**Note:** In SQN and adaQN, the data passed to all calls in partial fit will be stored in a limited-memory container which will be used to calculate Hessian-vector products or large-batch gradients. The size of this container is determined by the inputs ‘batch\_size’ and ‘bfgs\_upd\_freq’ passed in the constructor call.

---



---

**Note:** The step size in partial fit is determined by the number of optimizer iterations rather than the number of epochs, thus for a given amount of data, the default step size will be much smaller than when calling ‘fit’. Recommended to provide a custom step size function (‘decr\_step\_size’ in the initialization), as otherwise the step size sequence will be too small.

---

#### Parameters

- **X** (*array(n\_samples, m)*) – Sample data to with which to update the model.
- **y** (*array(n\_samples, )*) – Labels or target values for the sample data.
- **sample\_weight** (*None or array(n\_samples, )*) – Observations weights for the sample data.
- **additional\_kwargs** (*dict*) – Additional keyword arguments to pass to the objective, gradient, and Hessian-vector functions.

**Returns self** – This object.

**Return type** obj

**predict** (X, *additional\_kwargs={}*)

Make predictions on new data

---

**Note:** Using this method requires passing ‘pred\_fun’ in the initialization.

---

#### Parameters

- **X** (*array(n\_samples, m)*) – New data to pass to user-provided predict function.
- **additional\_kwargs** (*dict*) – Additional keyword arguments to pass to user-provided predict function.

```
class SQN_free (mem_size=10, bfgs_upd_freq=20, min_curvature=0.0001, y_reg=None,  

               use_grad_diff=False, check_nan=True, nthreads=-1, use_float=False)  

    SQN optimizer (free mode)
```

Optimizes an empirical (convex) loss function over batches of sample data. Compared to class ‘SQN’, this version lets the user do all the calculations from the outside, only interacting with the object by means of a function that returns a request type and is fed the required calculation through methods ‘update\_gradient’ and ‘update\_hess\_vec’.

Order in which requests are made:

```
===== loop ===== * calc_grad  

... (repeat calc_grad)
```

**if ‘use\_grad\_diff’:**

- calc\_grad\_big\_batch

**else:**

- calc\_hess\_vec

### Parameters

- **mem\_size** (*int*) – Number of correction pairs to store for approximation of Hessian-vector products.
- **bfgs\_upd\_freq** (*int*) – Number of iterations (batches) after which to generate a BFGS correction pair.
- **min\_curvature** (*float or None*) – Minimum value of  $s^*y / s^*s$  in order to accept a correction pair.
- **y\_reg** (*float or None*) – Regularizer for ‘y’ vector (gets added  $y\_reg * s$ ).
- **use\_grad\_diff** (*bool*) – Whether to create the correction pairs using differences between gradients instead of Hessian-vector products. These gradients are calculated on a larger batch than the regular ones (given by  $batch\_size * bfgs\_upd\_freq$ ).
- **check\_nan** (*bool*) – Whether to check for variables becoming NaN after each iteration, and reverting the step if they do (will also reset BFGS memory).
- **nthreads** (*int*) – Number of parallel threads to use. If set to -1, will determine the number of available threads and use all of them. Note however that not all the computations can be parallelized.
- **use\_float** (*bool*) – Whether to use C ‘float’ type (np.float32). If ‘False’ (the default), will use ‘double’ type (np.float64). The variables and gradient must be of this same dtype.

```
run_optimizer (x, step_size)
```

Continue optimization process after supplying the calculation requested from the last run

Continue the optimization process from where it was left since the last calculation was requested. Will internally do all the updates that are possible until the moment some calculation of function/gradient/hessian-vector is required.

---

**Note:** The first time this is run, no calculation needs to be supplied.

---

### Parameters



- **x** (*array(m, )*) – Current values of the variables. Will be modified in-place.
- **step\_size** (*float*) – Step size for the next update (note that variables are not updated during all runs).

### Returns

**request** – Dictionary with the calculation required to proceed and iteration information. Structure:

- **task** : str - one of “calc\_grad”, “calc\_grad\_same\_batch” (oLBFGS w. ‘min\_curvature’ or ‘check\_nan’), “calc\_hess\_vec” (SQN wo. ‘use\_grad\_diff’), “calc\_fun\_val\_batch” (adaQN w. ‘max\_incr’), “calc\_grad\_big\_batch” (SQN and adaQN w. ‘use\_grad\_diff’). \* requested\_on : array(m, ) or tuple(array(m, ), array(m, )), containing the values on which the request in “task” has to be evaluated. In the case of Hessian-vector products (SQN), the first vector is the values of ‘x’ and the second is the vector with which the product is required. \* info : dict(x\_changed\_in\_run : bool, iteration\_number : int, iteration\_info : str), iteration\_info can be one of “no\_problems\_encountered”, “search\_direction\_was\_nan”, “func\_increased”, “curvature\_too\_small”.

**Return type** dict

### update\_gradient (*gradient*)

Pass requested gradient to optimizer

**Parameters gradient** (*array(m, )*) – Gradient calculated as requested, evaluated at values given in “requested\_on”, calculated either in a regular batch (task = “calc\_grad”), same batch as before (task = “calc\_grad\_same\_batch” - oLBFGS only), or a larger batch of data (task = “calc\_grad\_big\_batch”), perhaps including all the cases from the last such calculation (SQN and adaQN with ‘use\_grad\_diff=True’).

### update\_hess\_vec (*hess\_vec*)

Pass requested Hessian-vector product to optimizer (task = “calc\_hess\_vec”)

**Parameters hess\_vec** (*array(m, )*) – Product of the Hessian evaluated at “requested\_on”[0] with the vector “requested\_on”[1], calculated a larger batch of data than the gradient, perhaps including all the cases from the last such calculation.

```
class adaQN(x0, grad_fun, obj_fun=None, pred_fun=None, batches_per_epoch=25, step_size=0.1,
decr_step_size=None, shuffle_data=True, random_state=1, nepochs=25, valset_frac=None,
tol=0.1, callback_epoch=None, callback_iter=None, kwargs_cb={}, verbose=True,
mem_size=10, fisher_size=100, bfgs_upd_freq=20, max_incr=1.01, min_curvature=0.0001,
y_reg=None, scal_reg=0.0001, rmsprop_weight=None, use_grad_diff=False,
check_nan=True, nthreads=-1, use_float=False)
```

adaQN optimizer

Optimizes an empirical (possibly non-convex) loss function over batches of sample data.

### Parameters

- **x0** (*array(m, )*) – Initial values of the variables to optimize (referred hereafter as ‘x’).
- **grad\_fun** (*function(x, X, y, sample\_weight, \*\*kwargs) -> array(m, )*) – Function that calculates the empirical gradient at values ‘x’ on data ‘X’ and ‘y’. Note: output must be one-dimensional and with the same number of entries as ‘x’, otherwise the Python session might segfault. (The extra keyword arguments are passed in the ‘fit’ method, not here)
- **obj\_fun** (*function(x, X, y, sample\_weight, \*\*kwargs) -> float*) – Function that calculates the empirical objective value at values ‘x’ on data ‘X’ and ‘y’. Will be ignored if passing

‘max\_incr=None’ and no validation set (‘valset\_frac=None’, and no ‘valset’ passed to fit). (The extra keyword arguments are passed in the ‘fit’ method, not here)

- **pred\_fun** (*None or function(xopt, X)*) – Prediction function taking as input the optimal ‘x’ values as obtained by the optimization procedure, and new observation ‘X’ on which to make predictions. If passed, will have an additional method oLBFGS.predict(X, \*args) that calls this function with current values of ‘x’.
- **batches\_per\_epoch** (*int*) – Number of batches per epoch (each batch will have the same number of observations except for the last one which might be smaller).
- **step\_size** (*float*) – Initial step size to use. (Can be modified after object is already initialized)
- **decr\_step\_size** (*str “auto”, None, or function(initial\_step\_size, epoch) -> float*) – Function that determines the step size during each epoch, taking as input the initial step size and the epoch number (starting at zero). If “auto”, will use 1/sqrt(iteration). If None, will use constant step size. For ‘partial\_fit’, it will take as input the number of iterations of the algorithm rather than epoch, so it’s very recommended to provide a custom function when passing data in user-provided batches. Can be modified after the object has been initialized (oLBFGS.decr\_step\_size)
- **shuffle\_data** (*bool*) – Whether to shuffle the data at the beginning of each epoch.
- **random\_state** (*int*) – Random seed to use for shuffling data and selecting validation set. The algorithm is deterministic so it’s not used for anything else.
- **nepochs** (*int*) – Number of epochs for which to run the optimization procedure. Might terminate earlier if using a validation set for monitoring.
- **valset\_frac** (*float(0, 1) or None*) – Percent of the data to use as validation set for early stopping. Can also pass a user-provided validation set to ‘fit’, in which case it will be ignored. If passing None, will run for the number of epochs passed in ‘nepochs’.
- **tol** (*float*) – If the objective function calculated on the validation set decrease by less than ‘tol’ upon completion of an epoch, will terminate the optimization procedure. Ignored when not using a validation set.
- **callback\_epoch** (*None or function\*(x, \*\*kwargs)*) – Callback function to call at the end of each epoch
- **callback\_iter** (*None or function\*(x, \*\*kwargs)*) – Callback function to call at the end of each iteration
- **kwargs\_cb** (*tuple*) – Additional arguments to pass to ‘callback’ and ‘stop\_crit’. (Can be modified after object is already initialized)
- **verbose** (*bool*) – Whether to print messages when there is some problem during an iteration (e.g. correction pair not meeting minum curvature).
- **mem\_size** (*int*) – Number of correction pairs to store for approximation of Hessian-vector products.
- **fisher\_size** (*int or None*) – Number of gradients to store for calculation of the empirical Fisher product with gradients. If passing ‘None’, will force ‘use\_grad\_diff’ to ‘True’.
- **bfgs\_upd\_freq** (*int*) – Number of iterations (batches) after which to generate a BFGS correction pair.
- **max\_incr** (*float or None*) – Maximum ratio of function values in the validation set under the average values of ‘x’ during current epoch vs. previous epoch. If the ratio is above this

threshold, the BFGS and Fisher memories will be reset, and 'x' values reverted to their previous average. If not using a validation set, will take a longer batch for function evaluations (same as used for gradients when using 'use\_grad\_diff=True').

- **min\_curvature** (*float or None*) – Minimum value of  $s^*y / s^*s$  in order to accept a correction pair.
- **y\_reg** (*float or None*) – regularizer for 'y' vector (gets added  $y\_reg * s$ )
- **scal\_reg** (*float*) – Regularization parameter to use in the denominator for AdaGrad and RMSProp scaling.
- **rmsprop\_weight** (*float(0,1) or None*) – If not 'None', will use RMSProp formula instead of AdaGrad for approximated inverse-Hessian initialization. (Recommended to use lower initial step size + passing 'decr\_step\_size')
- **use\_grad\_diff** (*bool*) – Whether to create the correction pairs using differences between gradients instead of Fisher matrix. These gradients are calculated on a larger batch than the regular ones (given by  $batch\_size * bfgs\_upd\_freq$ ). If 'True', fisher\_size will be set to None, and empirical Fisher matrix will not be used.
- **check\_nan** (*bool*) – Whether to check for variables becoming NaN after each iteration, and reverting the step if they do (will also reset BFGS memory).
- **nthreads** (*int*) – Number of parallel threads to use. If set to -1, will determine the number of available threads and use all of them. Note however that not all the computations can be parallelized.
- **use\_float** (*bool*) – Whether to use C 'float' type (np.float32). If 'False' (the default), will use 'double' type (np.float64). The variables and gradient must be of this same dtype.

## References

**fit** (*X, y, sample\_weight=None, additional\_kwargs={}, valset=None*)

Fit model to sample data

### Parameters

- **X** (*array(n\_samples, m)*) – Sample data to which to fit the model.
- **y** (*array(n\_samples, )*) – Labels or target values for the sample data.
- **sample\_weight** (*None or array(n\_samples, )*) – Observations weights for the sample data.
- **additional\_kwargs** (*dict*) – Additional keyword arguments to pass to the objective, gradient, and Hessian-vector functions.
- **valset** (*tuple(3)*) – User-provided validation set containing (X\_val, y\_val, sample\_weight\_val). At the end of each epoch, will calculate objective function on this set, and if the decrease from the objective function in the previous epoch is below tolerance, will terminate procedure earlier. If 'valset\_frac' was provided and a validation set is passed, 'valset\_frac' will be ignored. Must provide objective function in order to use a validation set.

**Returns** **self** – This object.

**Return type** **obj**

**get\_x** ()

Get a copy of current values of the variables

**Returns** **x** – Current variable values.

**Return type** array(n, )

**niter**

**partial\_fit** (X, y, sample\_weight=None, additional\_kwargs={})

Update model with user-provided batches of data

---

**Note:** In SQN and adaQN, the data passed to all calls in partial fit will be stored in a limited-memory container which will be used to calculate Hessian-vector products or large-batch gradients. The size of this container is determined by the inputs 'batch\_size' and 'bfgs\_upd\_freq' passed in the constructor call.

---



---

**Note:** The step size in partial fit is determined by the number of optimizer iterations rather than the number of epochs, thus for a given amount of data, the default step size will be much smaller than when calling 'fit'. Recommended to provide a custom step size function ('decr\_step\_size' in the initialization), as otherwise the step size sequence will be too small.

---

### Parameters

- **X** (array(n\_samples, m)) – Sample data to with which to update the model.
- **y** (array(n\_samples, )) – Labels or target values for the sample data.
- **sample\_weight** (None or array(n\_samples, )) – Observations weights for the sample data.
- **additional\_kwargs** (dict) – Additional keyword arguments to pass to the objective, gradient, and Hessian-vector functions.

**Returns** self – This object.

**Return type** obj

**predict** (X, additional\_kwargs={})

Make predictions on new data

---

**Note:** Using this method requires passing 'pred\_fun' in the initialization.

---

### Parameters

- **X** (array(n\_samples, m)) – New data to pass to user-provided predict function.
- **additional\_kwargs** (dict) – Additional keyword arguments to pass to user-provided predict function.

```
class adaQN_free (mem_size=10,      fisher_size=100,      bfgs_upd_freq=20,      max_incr=1.01,
                  min_curvature=0.0001, scal_reg=0.0001, rmsprop_weight=None, y_reg=None,
                  use_grad_diff=False, check_nan=True, nthreads=-1, use_float=False)
adaQN optimizer (free mode)
```

Optimizes an empirical (perhaps non-convex) loss function over batches of sample data. Compared to class 'adaQN', this version lets the user do all the calculations from the outside, only interacting with the object by means of a function that returns a request type and is fed the required calculation through methods 'update\_gradient' and 'update\_function'.

Order in which requests are made:

```
===== loop ===== * calc_grad
```

... (repeat calc\_grad)

if max\_incr > 0:

- calc\_fun\_val\_batch

if 'use\_grad\_diff':

- calc\_grad\_big\_batch (skipped if below max\_incr)

### Parameters

- **mem\_size** (*int*) – Number of correction pairs to store for approximation of Hessian-vector products.
- **fisher\_size** (*int or None*) – Number of gradients to store for calculation of the empirical Fisher product with gradients. If passing 'None', will force 'use\_grad\_diff' to 'True'.
- **bfgs\_upd\_freq** (*int*) – Number of iterations (batches) after which to generate a BFGS correction pair.
- **max\_incr** (*float or None*) – Maximum ratio of function values in the validation set under the average values of 'x' during current epoch vs. previous epoch. If the ratio is above this threshold, the BFGS and Fisher memories will be reset, and 'x' values reverted to their previous average. If not using a validation set, will take a longer batch for function evaluations (same as used for gradients when using 'use\_grad\_diff=True').
- **min\_curvature** (*float or None*) – Minimum value of  $s^*y / s^*s$  in order to accept a correction pair.
- **scal\_reg** (*float*) – Regularization parameter to use in the denominator for AdaGrad and RMSProp scaling.
- **rmsprop\_weight** (*float(0,1) or None*) – If not 'None', will use RMSProp formula instead of AdaGrad for approximated inverse-Hessian initialization.
- **y\_reg** (*float or None*) – Regularizer for 'y' vector (gets added  $y\_reg * s$ ).
- **use\_grad\_diff** (*bool*) – Whether to create the correction pairs using differences between gradients instead of Fisher matrix. These gradients are calculated on a larger batch than the regular ones (given by  $batch\_size * bfgs\_upd\_freq$ ). If 'True', fisher\_size will be set to None, and empirical Fisher matrix will not be used.
- **check\_nan** (*bool*) – Whether to check for variables becoming NaN after each iteration, and reverting the step if they do (will also reset BFGS memory).
- **nthreads** (*int*) – Number of parallel threads to use. If set to -1, will determine the number of available threads and use all of them. Note however that not all the computations can be parallelized.
- **use\_float** (*bool*) – Whether to use C 'float' type (np.float32). If 'False' (the default), will use 'double' type (np.float64). The variables and gradient must be of this same dtype.

**run\_optimizer** (*x, step\_size*)

Continue optimization process after supplying the calculation requested from the last run

Continue the optimization process from where it was left since the last calculation was requested. Will internally do all the updates that are possible until the moment some calculation of function/gradient/hessian-vector is required.

---

**Note:** The first time this is run, no calculation needs to be supplied.

---

### Parameters

- **x** (*array(m, )*) – Current values of the variables. Will be modified in-place. Do NOT modify the values between runs.
- **step\_size** (*float*) – Step size for the next update (note that variables are not updated during all runs).

### Returns

**request** – Dictionary with the calculation required to proceed and iteration information. Structure:

- **task** : str - one of “calc\_grad”, “calc\_grad\_same\_batch” (oLBFGS w. ‘min\_curvature’ or ‘check\_nan’), “calc\_hess\_vec” (SQN wo. ‘use\_grad\_diff’), “calc\_fun\_val\_batch” (adaQN w. ‘max\_incr’), “calc\_grad\_big\_batch” (SQN and adaQN w. ‘use\_grad\_diff’). \* **requested\_on** : array(m, ) or tuple(array(m, ), array(m, )), containing the values on which the request in “task” has to be evaluated. In the case of Hessian-vector products (SQN), the first vector is the values of ‘x’ and the second is the vector with which the product is required. \* **info** : dict(x\_changed\_in\_run : bool, iteration\_number : int, iteration\_info : str), iteration\_info can be one of “no\_problems\_encountered”, “search\_direction\_was\_nan”, “func\_increased”, “curvature\_too\_small”.

**Return type** dict

### update\_function (*fun*)

Pass requested function evaluation to optimizer (task = “calc\_fun\_val\_batch”)

**Parameters** **fun** (*float*) – Function evaluated at “requested\_on” under a validation set or a larger batch, perhaps including all the cases from the last such calculation.

### update\_gradient (*gradient*)

Pass requested gradient to optimizer

**Parameters** **gradient** (*array(m, )*) – Gradient calculated as requested, evaluated at values given in “requested\_on”, calculated either in a regular batch (task = “calc\_grad”), same batch as before (task = “calc\_grad\_same\_batch” - oLBFGS only), or a larger batch of data (task = “calc\_grad\_big\_batch”), perhaps including all the cases from the last such calculation (SQN and adaQN with ‘use\_grad\_diff=True’).

**class oLBFGS** (*x0, grad\_fun, obj\_fun=None, pred\_fun=None, batches\_per\_epoch=25, step\_size=0.001, decr\_step\_size='auto', shuffle\_data=True, random\_state=1, nepochs=25, valset\_frac=None, tol=0.1, callback\_epoch=None, callback\_iter=None, kwargs\_cb={}, verbose=True, mem\_size=10, hess\_init=None, min\_curvature=0.0001, y\_reg=None, check\_nan=True, nthreads=-1, use\_float=False*)

oLBFGS optimizer

Optimizes an empirical (convex) loss function over batches of sample data.

### Parameters

- **x0** (*array (m, )*) – Initial values of the variables to optimize (referred hereafter as ‘x’).
- **grad\_fun** (*function(x, X, y, sample\_weight, \*\*kwargs) -> array(m, )*) – Function that calculates the empirical gradient at values ‘x’ on data ‘X’ and ‘y’. Note: output must be

one-dimensional and with the same number of entries as 'x', otherwise the Python session might segfault. (The extra keyword arguments are passed in the 'fit' method, not here)

- **obj\_fun** (*function(x, X, y, sample\_weight, \*\*kwargs) -> float*) – Function that calculates the empirical objective value at values 'x' on data 'X' and 'y'. Only used when using a validation set ('valset\_frac' not None, or 'valset' passed to fit). Ignored when fitting the data in user-provided batches. (The extra keyword arguments are passed in the 'fit' method, not here)
- **pred\_fun** (*None or function(xopt, X)*) – Prediction function taking as input the optimal 'x' values as obtained by the optimization procedure, and new observation 'X' on which to make predictions. If passed, will have an additional method `oLBFGS.predict(X, *args)` that calls this function with current values of 'x'.
- **batches\_per\_epoch** (*int*) – Number of batches per epoch (each batch will have the same number of observations except for the last one which might be smaller).
- **step\_size** (*float*) – Initial step size to use. (Can be modified after object is already initialized)
- **decr\_step\_size** (*str "auto", None, or function(initial\_step\_size, epoch) -> float*) – Function that determines the step size during each epoch, taking as input the initial step size and the epoch number (starting at zero). If "auto", will use  $1/\sqrt{\text{iteration}}$ . If None, will use constant step size. For 'partial\_fit', it will take as input the number of iterations of the algorithm rather than epoch, so it's very recommended to provide a custom function when passing data in user-provided batches. Can be modified after the object has been initialized (`oLBFGS.decr_step_size`)
- **shuffle\_data** (*bool*) – Whether to shuffle the data at the beginning of each epoch.
- **random\_state** (*int*) – Random seed to use for shuffling data and selecting validation set. The algorithm is deterministic so it's not used for anything else.
- **nepochs** (*int*) – Number of epochs for which to run the optimization procedure. Might terminate earlier if using a validation set for monitoring.
- **valset\_frac** (*float(0, 1) or None*) – Percent of the data to use as validation set for early stopping. Can also pass a user-provided validation set to 'fit', in which case it will be ignored. If passing None, will run for the number of epochs passed in 'nepochs'.
- **tol** (*float*) – If the objective function calculated on the validation set decrease by less than 'tol' upon completion of an epoch, will terminate the optimization procedure. Ignored when not using a validation set.
- **callback\_epoch** (*None or function\*(x, \*\*kwargs)*) – Callback function to call at the end of each epoch
- **callback\_iter** (*None or function\*(x, \*\*kwargs)*) – Callback function to call at the end of each iteration
- **kwargs\_cb** (*tuple*) – Additional arguments to pass to 'callback' and 'stop\_crit'. (Can be modified after object is already initialized)
- **verbose** (*bool*) – Whether to print messages when there is some problem during an iteration (e.g. correction pair not meeting minum curvature).
- **mem\_size** (*int*) – Number of correction pairs to store for approximation of Hessian-vector products.
- **hess\_init** (*float or None*) – value to which to initialize the diagonal of H0. If passing 0, will use the same initialization as for SQN ( $s_{\text{last}}*y_{\text{last}} / y_{\text{last}}*y_{\text{last}}$ ).

- **min\_curvature** (*float or None*) – Minimum value of  $s^*y / s^*s$  in order to accept a correction pair.
- **y\_reg** (*float or None*) – regularizer for ‘y’ vector (gets added  $y\_reg * s$ )
- **check\_nan** (*bool*) – Whether to check for variables becoming NaN after each iteration, and reverting the step if they do (will also reset BFGS memory).
- **nthreads** (*int*) – Number of parallel threads to use. If set to -1, will determine the number of available threads and use all of them. Note however that not all the computations can be parallelized.
- **use\_float** (*bool*) – Whether to use C ‘float’ type (np.float32). If ‘False’ (the default), will use ‘double’ type (np.float64). The variables and gradient must be of this same dtype.

## References

**fit** (*X, y, sample\_weight=None, additional\_kwargs={}, valset=None*)

Fit model to sample data

### Parameters

- **X** (*array(n\_samples, m)*) – Sample data to which to fit the model.
- **y** (*array(n\_samples, )*) – Labels or target values for the sample data.
- **sample\_weight** (*None or array(n\_samples, )*) – Observations weights for the sample data.
- **additional\_kwargs** (*dict*) – Additional keyword arguments to pass to the objective, gradient, and Hessian-vector functions.
- **valset** (*tuple(3)*) – User-provided validation set containing (X\_val, y\_val, sample\_weight\_val). At the end of each epoch, will calculate objective function on this set, and if the decrease from the objective function in the previous epoch is below tolerance, will terminate procedure earlier. If ‘valset\_frac’ was provided and a validation set is passed, ‘valset\_frac’ will be ignored. Must provide objective function in order to use a validation set.

**Returns** **self** – This object.

**Return type** **obj**

**get\_x** ()

Get a copy of current values of the variables

**Returns** **x** – Current variable values.

**Return type** **array(n, )**

**niter**

**partial\_fit** (*X, y, sample\_weight=None, additional\_kwargs={}*)

Update model with user-provided batches of data

---

**Note:** In SQN and adaQN, the data passed to all calls in partial fit will be stored in a limited-memory container which will be used to calculate Hessian-vector products or large-batch gradients. The size of this container is determined by the inputs ‘batch\_size’ and ‘bfgs\_upd\_freq’ passed in the constructor call.

---



---

**Note:** The step size in partial fit is determined by the number of optimizer iterations rather than the number of epochs, thus for a given amount of data, the default step size will be much smaller than when

---



calling 'fit'. Recommended to provide a custom step size function ('decr\_step\_size' in the initialization), as otherwise the step size sequence will be too small.

### Parameters

- **X** (*array(n\_samples, m)*) – Sample data to with which to update the model.
- **y** (*array(n\_samples, )*) – Labels or target values for the sample data.
- **sample\_weight** (*None or array(n\_samples, )*) – Observations weights for the sample data.
- **additional\_kwargs** (*dict*) – Additional keyword arguments to pass to the objective, gradient, and Hessian-vector functions.

**Returns** **self** – This object.

**Return type** **obj**

**predict** (*X, additional\_kwargs={}*)  
Make predictions on new data

---

**Note:** Using this method requires passing 'pred\_fun' in the initialization.

---

### Parameters

- **X** (*array(n\_samples, m)*) – New data to pass to user-provided predict function.
- **additional\_kwargs** (*dict*) – Additional keyword arguments to pass to user-provided predict function.

**class oLBFGS\_free** (*mem\_size=10, hess\_init=None, min\_curvature=0.0001, y\_reg=None, check\_nan=True, nthreads=-1, use\_float=False*)  
oLBFGS optimizer (free mode)

Optimizes an empirical (convex) loss function over batches of sample data. Compared to class 'oLBFGS', this version lets the user do all the calculations from the outside, only interacting with the object by means of a function that returns a request type and is fed the required calculation through a method 'update\_gradient'.

Order in which requests are made:

```
===== loop ===== * calc_grad * calc_grad_same_batch (might skip if using
check_nan) =====
```

### Parameters

- **mem\_size** (*int*) – Number of correction pairs to store for approximation of Hessian-vector products.
- **hess\_init** (*float or None*) – value to which to initialize the diagonal of H0. If passing 'None', will use the same initialization as for SQN ( $s_{last} * y_{last} / y_{last} * y_{last}$ ).
- **min\_curvature** (*float or None*) – Minimum value of  $s * y / s * s$  in order to accept a correction pair.
- **y\_reg** (*float or None*) – Regularizer for 'y' vector (gets added  $y_{reg} * s$ ).
- **check\_nan** (*bool*) – Whether to check for variables becoming NaN after each iteration, and reverting the step if they do (will also reset BFGS memory).

- **nthreads** (*int*) – Number of parallel threads to use. If set to -1, will determine the number of available threads and use all of them. Note however that not all the computations can be parallelized.
- **use\_float** (*bool*) – Whether to use C ‘float’ type (np.float32). If ‘False’ (the default), will use ‘double’ type (np.float64). The variables and gradient must be of this same dtype.

**run\_optimizer** (*x, step\_size*)

Continue optimization process after supplying the calculation requested from the last run

Continue the optimization process from where it was left since the last calculation was requested. Will internally do all the updates that are possible until the moment some calculation of function/gradient/hessian-vector is required.

---

**Note:** The first time this is run, no calculation needs to be supplied.

---

### Parameters

- **x** (*array(m, )*) – Current values of the variables. Will be modified in-place. Do NOT modify the values between runs.
- **step\_size** (*float*) – Step size for the next update (note that variables are not updated during all runs).

### Returns

**request** – Dictionary with the calculation required to proceed and iteration information. Structure:

- **task** : str - one of “calc\_grad”, “calc\_grad\_same\_batch” (oLBFGS w. ‘min\_curvature’ or ‘check\_nan’), “calc\_hess\_vec” (SQN wo. ‘use\_grad\_diff’), “calc\_fun\_val\_batch” (adaQN w. ‘max\_incr’), “calc\_grad\_big\_batch” (SQN and adaQN w. ‘use\_grad\_diff’). \* requested\_on : array(m, ) or tuple(array(m, ), array(m, )), containing the values on which the request in “task” has to be evaluated. In the case of Hessian-vector products (SQN), the first vector is the values of ‘x’ and the second is the vector with which the product is required. \* info : dict(x\_changed\_in\_run : bool, iteration\_number : int, iteration\_info : str), iteration\_info can be one of “no\_problems\_encountered”, “search\_direction\_was\_nan”, “func\_increased”, “curvature\_too\_small”.

**Return type** dict

**update\_gradient** (*gradient*)

Pass requested gradient to optimizer

**Parameters** **gradient** (*array(m, )*) – Gradient calculated as requested, evaluated at values given in “requested\_on”, calculated either in a regular batch (task = “calc\_grad”), same batch as before (task = “calc\_grad\_same\_batch” - oLBFGS only), or a larger batch of data (task = “calc\_grad\_big\_batch”), perhaps including all the cases from the last such calculation (SQN and adaQN with ‘use\_grad\_diff=True’).

**class StochasticLogisticRegression** (*reg\_param=0.001, fit\_intercept=True, random\_state=1, optimizer='SQN', step\_size=0.1, valset\_frac=0.1, verbose=False, \*\*optimizer\_kwargs*)

Logistic Regression fit with stochastic quasi-Newton optimizer

### Parameters

- **reg\_param** (*float*) – Strength of l2 regularization. Note that the loss function has an average log-loss over observations, so the optimal regularization will likely be a lot smaller than for scikit-learn’s (which uses sum instead).
- **step\_size** (*float*) – Initial step size to use. Note that it will be decreased after each epoch when using ‘fit’, but will not be decreased after calling ‘partial\_fit’.
- **fit\_intercept** (*bool*) – Whether to add an intercept to the model parameters.
- **random\_state** (*int*) – Random seed to use.
- **optimizer** (*str, one of ‘oLBFGS’, ‘SQN’, ‘adaQN’*) – Optimizer to use.
- **optimizer\_kwargs** (*dict, optional*) – Additional options to pass to the optimizer (see each optimizer’s documentation).

**coef\_**

**fit** (*X, y, sample\_weight=None*)

Fit Logistic Regression model in stochastic batches

**Parameters**

- **X** (*array(n\_samples, n\_features)*) – Covariates (features).
- **y** (*array(n\_samples, ) or array(n\_samples, n\_classes)*) – Labels for each observation (must be already one-hot encoded).
- **sample\_weight** (*array(n\_samples, ) or None*) – Observation weights for each data point.

**Returns** **self** – This object

**Return type** **obj**

**intercept\_**

**partial\_fit** (*X, y, sample\_weight=None, classes=None, decr\_step\_size=False*)

Fit Logistic Regression model in stochastic batches

**Parameters**

- **X** (*array(n\_samples, n\_features)*) – Covariates (features).
- **y** (*array(n\_samples, ) or array(n\_samples, n\_classes)*) – Labels for each observation (must be already one-hot encoded).
- **sample\_weight** (*array(n\_samples, ) or None*) – Observation weights for each data point.
- **classes** (*None*) – Not used. Kept there for compatibility with other packages that assume scikit-learn’s API.
- **decr\_step\_size** (*bool*) – Whether to decrease or not decrease the step size after the update is done, according to the function ‘decr\_step\_size’ passed at initialization.

**Returns** **self** – This object

**Return type** **obj**

**predict** (*X*)

Predict the class of new observations

**Parameters** **X** (*array(n\_samples, n\_features)*) – Input data on which to predict classes.

**Returns** **pred** – Predicted class for each observation

**Return type** **array(n\_samples, )**

**predict\_proba** (*X*)

Predict class probabilities for new observations

**Parameters** *X* (*array(n\_samples, n\_features)*) – Input data on which to predict class probabilities.

**Returns** *pred* – Predicted class probabilities for each observation

**Return type** *array(n\_samples, n\_classes)*

# CHAPTER 1

---

## Indices and tables

---

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



### S

`stochqn._logistic`, [14](#)  
`stochqn._optimizers`, [1](#)  
`stochqn.tf`, [14](#)





## A

adaQN (*class in stochqn.\_optimizers*), 5  
adaQN\_free (*class in stochqn.\_optimizers*), 8

## C

coef\_ (*StochasticLogisticRegression attribute*), 15

## F

fit() (*adaQN method*), 7  
fit() (*oLBFGS method*), 12  
fit() (*SQN method*), 2  
fit() (*StochasticLogisticRegression method*), 15

## G

get\_x() (*adaQN method*), 7  
get\_x() (*oLBFGS method*), 12  
get\_x() (*SQN method*), 3

## I

intercept\_ (*StochasticLogisticRegression attribute*), 15

## N

niter (*adaQN attribute*), 8  
niter (*oLBFGS attribute*), 12  
niter (*SQN attribute*), 3

## O

oLBFGS (*class in stochqn.\_optimizers*), 10  
oLBFGS\_free (*class in stochqn.\_optimizers*), 13

## P

partial\_fit() (*adaQN method*), 8  
partial\_fit() (*oLBFGS method*), 12  
partial\_fit() (*SQN method*), 3  
partial\_fit() (*StochasticLogisticRegression method*), 15  
predict() (*adaQN method*), 8  
predict() (*oLBFGS method*), 13

predict() (*SQN method*), 3  
predict() (*StochasticLogisticRegression method*), 15  
predict\_proba() (*StochasticLogisticRegression method*), 15

## R

run\_optimizer() (*adaQN\_free method*), 9  
run\_optimizer() (*oLBFGS\_free method*), 14  
run\_optimizer() (*SQN\_free method*), 4

## S

SQN (*class in stochqn.\_optimizers*), 1  
SQN\_free (*class in stochqn.\_optimizers*), 3  
StochasticLogisticRegression (*class in stochqn.\_logistic*), 14  
stochqn.\_logistic (*module*), 14  
stochqn.\_optimizers (*module*), 1  
stochqn.tf (*module*), 14

## U

update\_function() (*adaQN\_free method*), 10  
update\_gradient() (*adaQN\_free method*), 10  
update\_gradient() (*oLBFGS\_free method*), 14  
update\_gradient() (*SQN\_free method*), 5  
update\_hess\_vec() (*SQN\_free method*), 5