

# Nesterov's Accelerated Gradient and Momentum as approximations to Regularised Update Descent

Aleksandar Botev

Department of Computer Science  
University College London  
Email: aleksandar.botev.14@ucl.ac.uk

Guy Lever

Department of Computer Science  
University College London  
Email: g.lever@ucl.ac.uk

David Barber

Department of Computer Science,  
University College London  
Alan Turing Institute  
Email: david.barber@ucl.ac.uk

**Abstract**—We present a unifying framework for adapting the update direction in gradient-based iterative optimization methods. As natural special cases we re-derive classical momentum and Nesterov’s accelerated gradient method, lending a new intuitive interpretation to the latter algorithm. We show that a new algorithm, which we term Regularised Gradient Descent, can converge more quickly than either Nesterov’s algorithm or the classical momentum algorithm.

## I. INTRODUCTION

We present a framework for optimisation by directly setting the parameter update to optimise the objective function. Under natural approximations, two special cases of this framework recover Nesterov’s Accelerated Gradient (NAG) descent [1] and the classical momentum method (MOM) [2]. This is particularly interesting in the case of NAG since, though popular and theoretically principled, it has largely defied intuitive interpretation. We show that (at least for the special quadratic objective case) that our algorithm can converge more quickly than either NAG or MOM.

Given a continuous objective  $J(\theta)$  we consider iterative algorithms to minimise  $J$ . We write  $J' \{\theta\}$  for the gradient of the function evaluated at  $\theta$ , and similarly  $J'' \{\theta\}$  for the second derivative<sup>1</sup>. Our focus is on first-order methods, namely those that form the parameter update on the basis of only first order gradient information.

### A. Gradient Descent

Perhaps the simplest optimisation approach is Gradient Descent (GD) which, starting from the current parameters, locally modifies the parameter  $\theta_t$  at iteration  $t$  to reduce  $J$ . Based on the Taylor series expansion

$$J(\theta_t + v_t) = J(\theta_t) + v_t J' \{\theta_t\} + O\{v_t^2\} \quad (1)$$

for a small learning rate  $\alpha_t > 0$ , setting  $v_t = -\alpha_t J' \{\theta_t\}$  reduces  $J$ . This motivates the GD update  $\theta_{t+1} = \theta_t + v_t$ . For convex Lipschitz  $J$  GD converges to the optimum value  $J^*$  as  $J(\theta_t) - J^* \sim 1/t$  [3]. Whilst gradient descent is universally popular, alternative methods such as momentum and Nesterov’s Accelerated Gradient (NAG) can result in significantly faster convergence to the optimum.

<sup>1</sup>These definitions extend in an obvious way to the gradient vector and Hessian in the vector  $\theta$  case.

### B. Momentum

The intuition behind momentum (MOM) is to continue updating the parameter along the previous update direction. This gives the algorithm (see for example [2])

$$\begin{aligned} v_{t+1} &= \mu_t v_t - \alpha_t J' \{\theta_t\} \\ \theta_{t+1} &= \theta_t + v_{t+1} \end{aligned} \quad (2)$$

The algorithm is particularly useful in the stochastic setting when only a sample of the gradient is available. By averaging the gradient over several minibatches/samples, the averaged gradient will better approximate the full batch gradient. In a different setting, when the objective function becomes flat, momentum is useful to maintain progress along directions of shallow gradient. As far as we are aware, relatively little is known about the convergence properties of momentum. We show below, at least for a special quadratic objective, that momentum indeed converges.

### C. Nesterov’s Accelerated Gradient

Nesterov’s Accelerated Gradient (NAG) [1] is given by

$$\begin{aligned} y_{t+1} &= (1 + \mu_t)\theta_t - \mu_t\theta_{t-1} \\ \theta_{t+1} &= y_{t+1} - \alpha_t J' \{y_{t+1}\} \end{aligned} \quad (3)$$

NAG has the interpretation that the previous two parameter values are smoothed and a gradient descent step is taken from this smoothed value. For Lipschitz convex functions (and a suitable schedule for  $\mu_t$  and  $\alpha_t$ ), NAG converges at rate  $1/t^2$ . Nesterov proved that this is the optimal possible rate for any method based on first order gradients<sup>2</sup> [1]. Nesterov proposed the schedule  $\mu_t = 1 - 3/(5+t)$  and fixed  $\alpha_t$ , which we adopt in the experiments.

Recently, [4] showed that by setting  $v_{t+1} = \theta_{t+1} - \theta_t$ , equation(3) can be rewritten as:

$$\begin{aligned} v_{t+1} &= \mu_t v_t - \alpha_t J' \{\theta_t + \mu_t v_t\} \\ \theta_{t+1} &= \theta_t + v_{t+1} \end{aligned} \quad (4)$$

This formulation reveals the relation of NAG to the classical momentum algorithm equation(2) which uses  $J' \{\theta_t\}$  in place

<sup>2</sup>This is a ‘worst case’ result. For example for quadratic functions, convergence is exponentially fast, leaving open the possibility that other algorithms may have superior convergence on ‘benign’ problems.

of  $J'\{\theta_t + \mu_t v_t\}$  in equation(4). In both cases, NAG and MOM tend to continue updating the parameters along the previous update direction.

In the machine learning community, NAG is largely viewed as somewhat mysterious and explained as performing a lookahead gradient evaluation and then performing a correction [4]. The closely related momentum is often loosely motivated by analogy with a physical system [2]. One contribution of our work, presented in section(II), shows that these algorithms can be intuitively understood from the perspective of optimising the objective with respect to the update  $v_t$  itself.

## II. REGULARISED UPDATE DESCENT

The central idea of our optimisation method is to consider the objective as a function of the update  $v_t$  and minimise the objective with respect to the update<sup>3</sup>. In order to prevent overly confident updates<sup>4</sup> we require that the update  $v_t$  is small and therefore add a regularisation term to the objective:

$$\tilde{J}(\theta_t, v_t) \equiv J(\theta_t + v_t) + \frac{\gamma_t}{2} v_t^2 \quad (5)$$

We first note that the optimum of  $\tilde{J}$  occurs when

$$\frac{\partial \tilde{J}}{\partial \theta_t} = 0 \quad (6)$$

$$\frac{\partial \tilde{J}}{\partial v_t} = 0 \quad (7)$$

These two conditions give

$$J'(\theta_t + v_t) = 0 \quad (8)$$

$$J'(\theta_t + v_t) + \gamma_t v_t = 0 \quad (9)$$

which implies that at the optimum  $v_t = 0$  and therefore that  $J'(\theta_t) = 0$  when we have found the optimum of  $\tilde{J}$ . Hence, the  $\theta_t$  that minimises  $\tilde{J}$  also minimises  $J$ .

Differentiating with respect to  $v_t$  we obtain

$$J'(\theta_t + v_t) + \gamma_t v_t \quad (10)$$

We thus make a gradient descent update in the direction that lowers the regularised objective:

$$v_{t+1} = v_t - \alpha_t \{J'(\theta_t + v_t) + \gamma_t v_t\} \quad (11)$$

A natural thing to do is a joint gradient descent step in  $\theta_t, v_t$  space. This would give the update

$$\theta_{t+1} = \theta_t - \alpha_t J'(\theta_t + v_t) \quad (12)$$

which can be written

$$\theta_{t+1} = \theta_t + v_{t+1} - \mu_t v_t \quad (13)$$

where we defined the momentum parameter  $\mu_t = 1 - \alpha_t \gamma_t$ . Since  $\alpha_t$  should be small,  $\mu_t$  will be less than 1. The learning

<sup>3</sup>Previous authors have also considered optimising the update, for example [5].

<sup>4</sup>Even when  $J'(\theta_t + v_t) = 0$ , without the regularising term,  $v_{t+1}$  would be equal to  $v_t$  and the parameter update would not converge.

rate is given by  $0 < \alpha_t$ . As we converge to the minimum,  $v_t$  tends to zero. In this limit the above joint gradient descent in  $(\theta_t, v_t)$  space approaches the standard gradient descent step. This form of joint descent therefore is asymptotically equivalent to standard gradient descent. We therefore seek another approach to update  $\theta_t$ .

The complete Regularised Update Descent (RUD) algorithm is thus given by

$$\begin{aligned} v_{t+1} &= \mu_t v_t - \alpha_t J'(\theta_t + v_t) \\ \theta_{t+1} &= \theta_t + v_{t+1} \end{aligned} \quad (14)$$

As we converge towards a minimum, the update  $v_t$  will become small (since the gradient is small) and the regularisation term can be safely tuned down. This means that  $\mu_t$  should be set so that it tends to 1 with increasing iterations. As we will show below one can view momentum and Nesterov as approximations to RUD based on a first order expansion (for momentum) and a more accurate second order expansion (for Nesterov).

### A. Deriving Momentum from RUD

We consider an update  $v_t$  at the current  $\theta_t$ . Assuming  $v_t$  is small:

$$J(\theta_t + v_t) = J(\theta_t) + v_t J'(\theta_t) + O\{v_t^2\} \quad (15)$$

Under this first order approximation, the RUD objective becomes

$$J(\theta_t) + v_t J'(\theta_t) + \frac{\gamma_t}{2} v_t^2 \quad (16)$$

Differentiating wrt  $v_t$  we get

$$J'(\theta_t) + \gamma_t v_t \quad (17)$$

We thus make an update in this direction:

$$v_{t+1} = v_t - \alpha_t \{J'(\theta_t) + \gamma_t v_t\} \quad (18)$$

$$= \mu_t v_t - \alpha_t J'(\theta_t) \quad (19)$$

where  $\mu_t$  should be close to 1. We then make a parameter update

$$\theta_{t+1} = \theta_t + v_{t+1} \quad (20)$$

which recovers the momentum algorithm. We can therefore view momentum as optimising, with respect to the update, a first order approximation of the RUD objective.

### B. Deriving Nesterov from RUD

Expanding  $J(\theta_t + v_t)$  to the next order, we obtain

$$J(\theta_t + v_t) = J(\theta_t) + v_t J'(\theta_t) + \frac{1}{2} v_t^2 J''(\theta_t) + O\{v_t^3\} \quad (21)$$

Since  $v_t$  is not infinitesimally small, we cannot ‘trust’ the higher order terms as we move away from  $v_t = 0$ ; as we move further from  $\theta_t$  we are trying to approximate the function based on curvature information at  $\theta_t$ , rather than the current point  $v_t + \theta_t$ . This is analogous to the idea of trust regions in Quasi-Newton approaches which limit the extent to which the Taylor expansion is trusted away from the origin [3]. To

encode this lack of trust, we reduce the second order term by a factor  $\mu_t < 1$  and add another term to encourage  $v_t$  to be small. This gives the modified approximate RUD objective

$$J(\theta) + v_t J'(\theta_t) + \frac{\mu_t}{2} v_t^2 J''(\theta) + \frac{\gamma_t}{2} v_t^2 \quad (22)$$

Differentiating with respect to  $v_t$  we get

$$J'(\theta_t) + \mu_t v_t J''(\theta_t) + \gamma_t v_t = J'(\theta_t + \mu_t v_t) + \gamma_t v_t + O\{v^2\} \quad (23)$$

We then update  $v_t$  to reduce this approximate RUD objective:

$$v_{t+1} = v_t - \alpha_t \{ J'(\theta_t + \mu_t v_t) + \gamma_t v_t \} \quad (24)$$

$$= (1 - \alpha_t \gamma_t) v_t - \alpha_t J'(\theta_t + \mu_t v_t) \quad (25)$$

We are free to choose  $\alpha_t$ , and  $\gamma_t$  which should both be small. Ideally  $\mu_t$  should be close to 1. Hence, it is reasonable to set  $1 - \alpha_t \gamma_t = \mu_t$  and choose  $\mu_t$  to be close to 1. This setting recovers the NAG algorithm:

$$v_{t+1} = \mu_t v_t - \alpha_t J'(\theta_t + \mu_t v_t) \quad (26)$$

$$\theta_{t+1} = \theta_t + v_{t+1} \quad (27)$$

and explains why we want  $\mu_t$  to tend to 1 as we converge, since as we zoom in to the minimum, we can trust more a quadratic approximation to the objective.

From the perspective that NAG and MOM are approximations to RUD, NAG is preferable to MOM since it is based on a more accurate expansion. In terms of RUD versus NAG, the difference between NAG and RUD is the use of  $\mu_t$  in the argument of  $J'(\theta_t + \mu_t v_t)$  in NAG, whereas we use  $J'(\theta_t + v_t)$  in RUD. This means that RUD ‘looks further forward’ than NAG (since  $\mu_t < 1$ ) in a manner more consistent with the eventual parameter update  $\theta_t + v_{t+1}$ . This tentatively explains why RUD can outperform NAG.

### C. Alternative NAG derivation

We consider the regularised objective

$$\tilde{J}(\theta_t, v_t) = J(\theta_t + v_t) + \frac{1}{2} \gamma_t v_t^2 \quad (28)$$

We consider a two stage process of optimizing  $\tilde{J}$ . The algorithm proceeds as follows: given  $\theta_t$  and  $v_t$  we first perform a descent step only on the regularizer, followed by a descent step on the ‘lookahead’  $J(\theta_t + v)$ . After this we perform the usual step on  $\theta_t$  based on the final updated  $v$ . The procedure is summarized below:

$$\begin{aligned} \tilde{v}_{t+1} &= v_t - \alpha_t \gamma_t v_t = (1 - \alpha_t \gamma_t) v_t \\ g_t &= J'(\theta_t + \tilde{v}_{t+1}) \\ v_{t+1} &= \tilde{v}_{t+1} - \alpha_t g_t = (1 - \alpha_t \gamma_t) v_t - \alpha_t g_t \\ \theta_{t+1} &= \theta_t + v_{t+1} \end{aligned} \quad (29)$$

Setting  $\gamma_t = \frac{1-\mu_t}{\alpha_t}$  recovers the original NAG formulation as in [4].

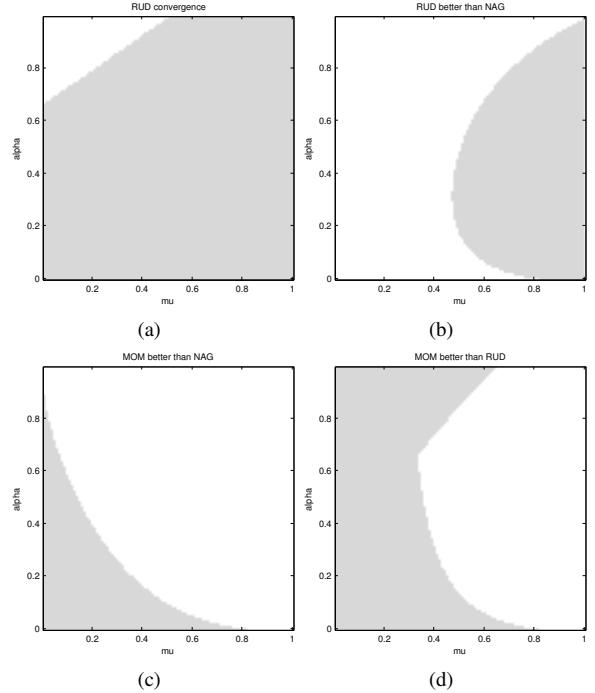


Fig. 1. (a) Shaded is the parameter region  $(\mu, \alpha)$  for which RUD converges for the simple quadratic function  $f(\theta) = 0.5\theta^2$ . (b) Shaded is the parameter region  $(\mu, \alpha)$  for which RUD converges more quickly than NAG. (c) Shaded is the region in which MOM converges more quickly than NAG. (d) Shaded is the region in which MOM converges more quickly than RUD.

### III. COMPARISON ON A QUADRATIC FUNCTION

An interesting question is whether and under what conditions RUD may converge more quickly than NAG for convex Lipschitz functions. To date we have not been able to fully analyse this. In lieu of a more complete understanding we consider the simple quadratic objective<sup>5</sup>

$$J(\theta) = \frac{1}{2} \theta^2 \quad (30)$$

For this simple function the gradient is given by  $\theta$  and, for fixed  $\alpha_t$ ,  $\mu_t$ , we are able to fully compute the update trajectories for NAG and RUD and MOM.

#### A. NAG

For NAG, the algorithm is given by

$$v_{t+1} = \mu v_t - \alpha(\theta_t + \mu v_t) \quad (31)$$

$$\theta_{t+1} = \theta_t + v_{t+1} \quad (32)$$

Assuming  $v_1 = 0$ , and a given value for  $\theta_1$ , this gives  $\theta_2 = (1 - \alpha)\theta_1$ . Similarly, for both MOM and NAG,  $\theta_2$  is given by the same value.

<sup>5</sup>For the simple quadratic objective, the convergence is exponentially fast in terms of the number of iterations. This is clearly a very special case compared to the more interesting convex Lipschitz scenario. Nevertheless, the analysis gives some insight that some improvement over NAG might be possible.

We can write equations (31,32) as a single second order difference equation

$$\theta_{t+1} + b\theta_t + c\theta_{t-1} = 0 \quad (33)$$

where

$$b \equiv -1 - \mu + \alpha + \alpha\mu \quad (34)$$

$$c \equiv \mu - \alpha\mu \quad (35)$$

For the scalar case  $\dim(\theta) = 1$ , assuming a solution of the form  $\theta_t = Aw^t$  gives

$$w = \frac{-b \pm \sqrt{b^2 - 4c}}{2} \quad (36)$$

which defines two values  $w_+$  and  $w_-$ , so that the general solution is given by

$$\theta_t = Aw_+^t + Bw_-^t \quad (37)$$

where  $A$  and  $B$  are determined by the linear equations

$$\theta_1 = Aw_+ + Bw_- \quad (38)$$

$$\theta_2 = Aw_+^2 + Bw_-^2 \quad (39)$$

A sufficient condition for NAG to converge is that  $|w_+| < 1$  and  $|w_-| < 1$  which is equivalent to the conditions  $|b| < 1+c$ ,  $c < 1$  [6]. For any learning rate  $0 < \alpha < 1$  and momentum  $0 < \mu < 1$ , it is straightforward to show that these conditions hold and thus that NAG converges to the minimum  $\theta_* = 0$ .

### B. MOM

The above analysis carries over directly for the MOM algorithm, with the only change being

$$b \equiv -1 - \mu + \alpha \quad (40)$$

$$c \equiv \mu \quad (41)$$

It is straightforward to show that for any learning rate  $0 < \alpha < 1$  and momentum  $0 < \mu < 1$ , the corresponding conditions  $|w_+| < 1$  and  $|w_-| < 1$  are always satisfied. Therefore the MOM algorithm (at least for this problem) always converges. For MOM to have better asymptotic convergence rate than NAG, we need  $\max(|w_+^{MOM}|, |w_-^{MOM}|) < \max(|w_+^{NAG}|, |w_-^{NAG}|)$ . From fig(1) we see that MOM only outperforms NAG (and RUD) when the momentum is small. This is essentially the uninteresting regime since, in practice, we will typically use a value of momentum that is close to 1. For this simple quadratic case, for practical purposes, MOM therefore performs worse than RUD or NAG.

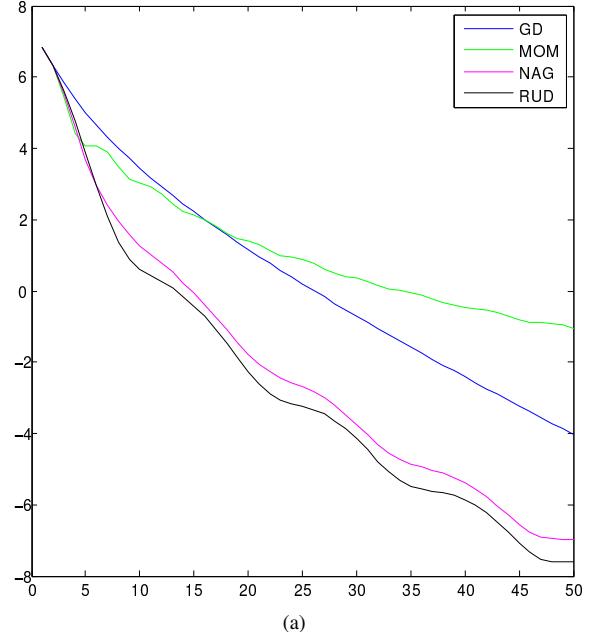
### C. RUD

For the RUD algorithm the corresponding solutions are given by setting

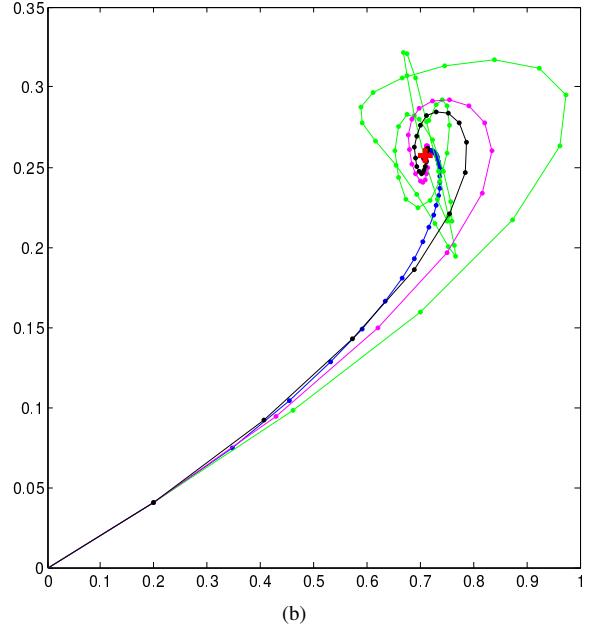
$$b \equiv -1 - \mu + 2\alpha \quad (42)$$

$$c \equiv \mu - \alpha \quad (43)$$

RUD has more complex convergence behaviour than NAG or MOM. The conditions  $|w_+| < 1$  and  $|w_-| < 1$  are satisfied



(a)



(b)

Fig. 2. Optimising a 1000 dimensional quadratic function  $J(\theta)$  using different algorithms, all with the same learning rate  $\alpha_t$  and  $\mu_t$  schedule. (a) The log objective  $\log J(\theta_t)$  for Gradient Descent, Momentum, Nesterov's Accelerated Gradient and Regularised Update Descent. (b) Trajectories of the different algorithms plotted for the first two components  $(\theta_{1t}, \theta_{2t})$ . The behaviour demonstrated is typical in that momentum tends to more significantly overshoot the minimum than RUD or NAG, with RUD typically outperforming NAG.

only within the region as shown in fig(1a), which is determined by

$$1 + \mu > \frac{3}{2}\alpha \quad (44)$$

The main requirement is that the learning rate should not be too high, at least for values of momentum  $\mu$  less than 0.5.

Unlike NAG and MOM, RUD has therefore the possibility to diverge.

In fig(1b) we show the region for which the asymptotic convergence of RUD is faster than NAG. The main requirement is that the momentum needs to be high (say above 0.8) and is otherwise largely independent of the learning rate (provided  $\alpha < 1$ ).

#### IV. EXPERIMENTS

##### A. A toy high dimensional quadratic function

In fig(2) we show the progress for different algorithms using the same learning rate  $\alpha_t = 0.2$  and  $\mu_t = 1 - 3/(5 + t)$  for a toy 1000 dimension quadratic function  $\frac{1}{2}\theta^\top A\theta - \theta^\top b$  for randomly chosen  $A$  and  $b$ . This simple experiment shows that the theoretical property derived in section(III) that RUD can outperform NAG and MOM carries over to the more general quadratic setting. Indeed, in our experience, the improved convergence of RUD over NAG for the quadratic objective function is typical behaviour.

##### B. Deep Learning: MNIST

Whilst RUD has interesting convergence for quadratic functions, in practice of course it is interesting to see how it behaves in the case of more general non-convex functions. In fig(3) we look at a classical deep learning problem of training an 784 – 1000 – 500 – 25 – 30 autoencoder for handwritten digit reconstruction [7]. The dataset consists of black and white images of size 28x28 and we used 50000 training images, with the images scaled to lie in the 0 to 1 range. The target is for the network to learn to reduce the dimensionality of the input to a 30 dimensional vector and then to reconstruct the input. The nonlinearity at each layer is the hyperbolic tangent<sup>6</sup> and for the last layer we used the binary cross entropy loss.

Since NAG and RUD are closely related, we use the same schedule  $\mu_t = 1 - 3/(5 + t)$  for both algorithms. All remaining hyperparameters for each method (learning rates) were set optimally based on a grid search over a set of reasonable parameters for each algorithm. For this problem, there is little difference between NAG and RUD, with RUD slightly outperforming NAG.

#### V. CONCLUSION

We described a general approach to first order optimisation based on optimising the objective with respect to the updates. This gives a simple optimisation algorithm which we termed Regularised Update Descent; we showed that his algorithm can converge more quickly than Nesterov's Accelerated Gradient. In addition to being a potentially useful optimisation algorithm in its own right, the main contribution of this work is to show that the Nesterov and momentum algorithms can be viewed as approximations to the Regularised Update Descent algorithm.

<sup>6</sup>We tried also rectifier linear units and leaky rectifier linear units, but they did not affect the relative performance of any of the algorithms.

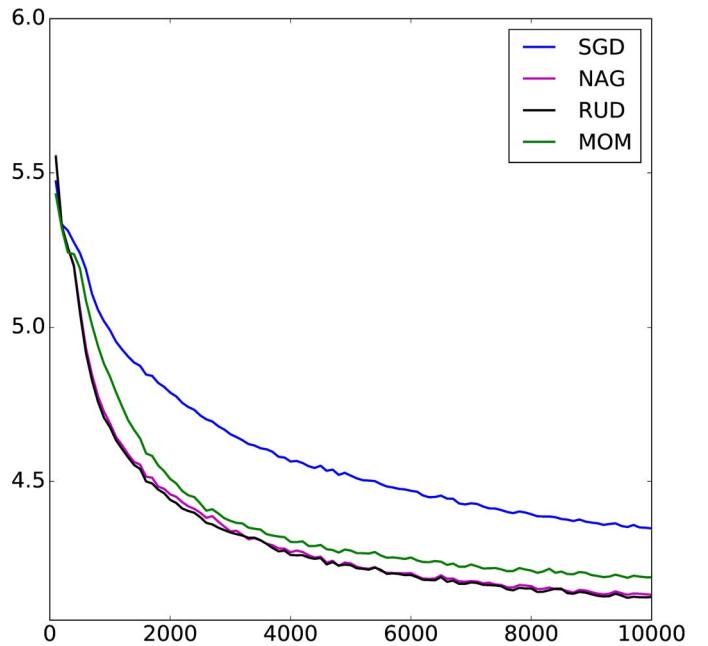


Fig. 3. The negative log loss for the classical MNIST 784 – 1000 – 500 – 25 – 30 autoencoder network [7] trained using minibatches contains 200 examples. Similar to the small quadratic objective experiments, we see that on this much larger problem, as expected, NAG and RUD perform very similarly (with RUD slightly outperforming NAG). All methods used the same learning rate and momentum parameter  $\mu_t$  schedule.

#### REFERENCES

- [1] Y. Nesterov, “A method of solving a convex programming problem with convergence rate  $O(1/k^2)$ ,” in *Soviet Mathematics Doklady*, vol. 27, no. 2, 1983, pp. 372–376.
- [2] N. Qian, “On the momentum term in gradient descent learning algorithms,” *Neural networks*, vol. 12, no. 1, pp. 145–151, 1999.
- [3] J. Nocedal and S. J. Wright, *Numerical optimization*, ser. Springer Series in Operations Research and Financial Engineering. Berlin: Springer, 2006. [Online]. Available: <http://opac.inria.fr/record=b1120179>
- [4] I. Sutskever, J. Martens, G. Dahl, and G. Hinton, “On the importance of initialization and momentum in deep learning,” in *Proceedings of the 30th international conference on machine learning (ICML-13)*, 2013, pp. 1139–1147.
- [5] P.-Y. Massé and Y. Ollivier, “Speed learning on the fly,” *arXiv preprint arXiv:1511.02540*, 2015.
- [6] K. Sydsæter and P. Hammond, *Essential Mathematics for Economic Analysis*. Prentice Hall, 2008.
- [7] G. E. Hinton and R. R. Salakhutdinov, “Reducing the dimensionality of data with neural networks,” *Science*, vol. 313, no. 5786, pp. 504–507, 2006.