
Reparameterizing Mirror Descent as Gradient Descent

Ehsan Amid* and Manfred K. Warmuth
Google Research, Brain Team
Mountain View, CA
{eamid, manfred}@google.com

Abstract

Most of the recent successful applications of neural networks have been based on training with gradient descent updates. However, for some small networks, other mirror descent updates learn provably more efficiently when the target is sparse. We present a general framework for casting a mirror descent update as a gradient descent update on a different set of parameters. In some cases, the mirror descent reparameterization can be described as training a modified network with standard backpropagation. The reparameterization framework is versatile and covers a wide range of mirror descent updates, even cases where the domain is constrained. Our construction for the reparameterization argument is done for the continuous versions of the updates. Finding general criteria for the discrete versions to closely track their continuous counterparts remains an interesting open problem.

1 Introduction

Mirror descent (MD) [Nemirovski and Yudin, 1983, Kivinen and Warmuth, 1997] refers to a family of updates which transform the parameters $\mathbf{w} \in \mathcal{C}$ from a convex domain $\mathcal{C} \in \mathbb{R}^d$ via a link function (a.k.a. mirror map) $f : \mathcal{C} \rightarrow \mathbb{R}^d$ before applying the descent step. The *continuous-time mirror descent* (CMD) update, which can be seen as the limit case of (discrete-time) MD, corresponds to the solution of the following ordinary differential equation (ODE) [Nemirovski and Yudin, 1983, Warmuth and Jagota, 1998, Raginsky and Bouvrie, 2012]:

$$\frac{f(\mathbf{w}(t+h)) - f(\mathbf{w}(t))}{h} \stackrel{h \rightarrow 0}{\approx} \dot{f}(\mathbf{w}(t)) = -\nabla L(\mathbf{w}(t)), \quad (\text{CMD}) \quad (1)$$

$$\mathbf{w}(t+h) = f^{-1}\left(f(\mathbf{w}(t)) - h \nabla L(\mathbf{w}(t))\right). \quad (\text{MD}) \quad (2)$$

Here L denotes a differentiable real-valued loss and $\dot{f} := \frac{\partial f}{\partial t}$ is the time derivative of the link function. The vanilla discretized MD update is obtained by setting the step size to h . The main link functions investigated in the past are $f(\mathbf{w}) = \mathbf{w}$ and $f(\mathbf{w}) = \log(\mathbf{w})$ leading to the gradient descent (GD) and the unnormalized exponentiated gradient (EGU) family of updates.² These two link functions are associated with the squared Euclidean and the relative entropy divergences, respectively. For example, the classical Perceptron and Winnow algorithms are motivated using the identity and log links, respectively, when the loss is the hinge loss. A number of papers discuss the difference between the two updates [Kivinen and Warmuth, 1997, Kivinen et al., 2006, Nie et al., 2016, Ghai et al., 2020] and their rotational invariance properties have been explored in [Warmuth et al., 2014]. In

*An earlier version of this manuscript (with additional results on the matrix case) appeared as "Interpolating Between Gradient Descent and Exponentiated Gradient Using Reparameterized Gradient Descent" as a preprint.

²The normalized version is called EG and the two-sided version EGU^\pm . More about this later.

particular, the *Hadamard problem* is a paradigmatic linear problem which shows that EGU can converge dramatically faster than GD when the instances are dense and the target weight vector is sparse [Kivinen et al., 1997, Warmuth and Vishwanathan, 2005]. This property is linked to the strong-convexity of the relative entropy w.r.t. the L_1 -norm³ [Shalev-Shwartz et al., 2012], which motivates the discrete EGU update.

Contributions Although other MD updates can be drastically more efficient than GD updates on certain classes of problems, it was assumed that such MD updates are not realizable using GD. In this paper, we show that in fact a large number of MD updates (such as EGU, and those motivated by the Burg and Inverse divergences) can be reparameterized as GD updates. Concretely, our contributions can be summarized as follows.

- We cast continuous MD updates as minimizing a trade off between a *Bregman momentum* and the loss. We also derive the dual, natural gradient, and the constrained versions of the updates.
- We then provide a general framework that allows reparameterizing one CMD update by another. It requires the existence of a certain reparameterization function and a condition on the derivatives of the two link functions as well as the reparameterization function.
- Specifically, we show that on certain problems, the implicit bias of the GD updates can be controlled by considering a family of *tempered* updates (parameterized by a *temperature* $\tau \in \mathbb{R}$) that interpolate between GD (with $\tau = 0$) and EGU (with $\tau = 1$), while covering a wider class of updates.

We conclude the paper with a number of open problems for future research directions.

Previous work There has been an increasing amount of interest recently in determining the implicit bias of learning algorithms [Gunasekar et al., 2017, 2018, Vaskevicius et al., 2019]. Here, we mainly focus on the MD updates. The special case of reparameterizing continuous EGU as continuous GD was already known [Akin, 1979, Amid and Warmuth, 2020]. In this paper, we develop a more general framework for reparameterizing one CMD update by another. We give a large variety of examples for reparameterizing the CMD updates as continuous GD updates. The main new examples we consider are based on the tempered versions of the relative entropy divergence [Amid et al., 2019]. The main open problem regarding the CMD updates is whether the discretization of the reparameterized updates track the discretization of the original (discretized) MD updates. The strongest methodology for showing this would be to prove the same regret bounds for the discretized reparameterized update as for the original. This has been done in a case-by-case basis for the EG family [Amid and Warmuth, 2020]. For more discussion see the conclusion section, where we also discuss how our reparameterization method allows exploring the effect of the structure of the neural network on the implicit bias.

Some basic notation We use \odot , \oslash , and superscript \odot for element-wise product, division, and power, respectively. We let $\mathbf{w}(t)$ denote the weight or parameter vector as a function of time t . Learning proceeds in steps. During step s , we start with weight vector $\mathbf{w}(sh) = \mathbf{w}_s$ and go to $\mathbf{w}((s+1)h) = \mathbf{w}_{s+1}$ while processing a batch of examples. We also write the Jacobian of vector valued function q as \mathbf{J}_q and use \mathbf{H}_F to denote the Hessian of a scalar function F . Furthermore, we let $\nabla_{\mathbf{w}} F(\mathbf{w}(t))$ denote the gradient of function $F(\mathbf{w})$ evaluated at $\mathbf{w}(t)$ and often drop the subscript \mathbf{w} for conciseness.

2 Continuous-time Mirror Descent

For a strictly convex, continuously-differentiable function $F : \mathcal{C} \rightarrow \mathbb{R}$ with convex domain $\mathcal{C} \subseteq \mathbb{R}^d$, the *Bregman divergence* between $\tilde{\mathbf{w}}, \mathbf{w} \in \mathcal{C}$ is defined as

$$D_F(\tilde{\mathbf{w}}, \mathbf{w}) := F(\tilde{\mathbf{w}}) - F(\mathbf{w}) - f(\mathbf{w})^\top (\tilde{\mathbf{w}} - \mathbf{w}),$$

where $f := \nabla F$ denotes the gradient of F , sometimes called the *link function*.⁴ Trading off the divergence to the last parameter \mathbf{w}_s with the current loss lets us motivate the iterative *mirror descent*

³Whereas the squared Euclidean divergence (which motivates GD) is strongly-convex w.r.t. the L_2 -norm.

⁴The gradient of a scalar function is a special case of a Jacobian, and should therefore be denoted by a row vector. However, in this paper we use the more common column vector notation for gradients, i.e. $\nabla F(\mathbf{w}) := (\frac{\partial F}{\partial \mathbf{w}})^\top$.

(MD) updates [Nemirovski and Yudin, 1983, Kivinen and Warmuth, 1997]:

$$\mathbf{w}_{s+1} = \underset{\mathbf{w}}{\operatorname{argmin}} \ 1/h D_F(\mathbf{w}, \mathbf{w}_s) + L(\mathbf{w}), \quad (3)$$

where $h > 0$ is often called the *learning rate*. Solving for \mathbf{w}_{s+1} yields the so-called *prox* or *implicit update* [Rockafellar, 1976, Nemirovski and Yudin, 1983, Kivinen et al., 2006]:

$$f(\mathbf{w}_{s+1}) = f(\mathbf{w}_s) - h \nabla L(\mathbf{w}_{s+1}). \quad (4)$$

This update is typically approximated by the following *explicit* update that uses the gradient at the old parameter \mathbf{w}_s instead (denoted here as the MD update):

$$f(\mathbf{w}_{s+1}) = f(\mathbf{w}_s) - h \nabla L(\mathbf{w}_s). \quad (\text{MD}) \quad (5)$$

We now show that the CMD update (1) can be motivated similarly by replacing the Bregman divergence in the minimization problem (3) with a ‘‘momentum’’ version which quantifies the rate of change in the value of Bregman divergence as $\mathbf{w}(t)$ varies over time. For the convex function F , we define the *Bregman momentum* between $\mathbf{w}(t)$, $\mathbf{w}_0 \in \mathcal{C}$ as the time differential of the Bregman divergence induced by F ,

$$\dot{D}_F(\mathbf{w}(t), \mathbf{w}_0) = \dot{F}(\mathbf{w}(t)) - f(\mathbf{w}_0)^\top \dot{\mathbf{w}}(t) = (f(\mathbf{w}(t)) - f(\mathbf{w}_0))^\top \dot{\mathbf{w}}(t).$$

Theorem 1 (Main result #1). *The CMD update* (5)

$$\dot{f}(\mathbf{w}(t)) = -\nabla L(\mathbf{w}(t)), \text{ with initial condition } \mathbf{w}(0) = \mathbf{w}_0,$$

is the solution of the following functional (6)

$$\min_{\text{curve } \mathbf{w}(t)} \left\{ \dot{D}_F(\mathbf{w}(t), \mathbf{w}_0) + L(\mathbf{w}(t)) \right\}. \quad (6)$$

Proof. Setting the derivatives w.r.t. $\mathbf{w}(t)$ to zero, we have

$$\begin{aligned} & \frac{\partial}{\partial \mathbf{w}(t)} \left((f(\mathbf{w}(t)) - f(\mathbf{w}_0))^\top \dot{\mathbf{w}}(t) + L(\mathbf{w}(t)) \right) \\ &= \mathbf{H}_F(\mathbf{w}(t)) \dot{\mathbf{w}}(t) + \frac{\partial \dot{\mathbf{w}}(t)}{\partial \mathbf{w}(t)} (f(\mathbf{w}(t)) - f(\mathbf{w}_0)) + \nabla L(\mathbf{w}(t)) \\ &= \dot{f}(\mathbf{w}(t)) + \nabla L(\mathbf{w}(t)) = \mathbf{0}, \end{aligned}$$

where we use the fact that $\mathbf{w}(t)$ and $\dot{\mathbf{w}}(t)$ are independent variables (7) [Burke, 1985], thus $\frac{\partial \dot{\mathbf{w}}(t)}{\partial \mathbf{w}(t)} = \mathbf{0}$. \square

Note that the implicit update (4) and the explicit update (5) can both be realized as the backward and the forward Euler approximations of (1), respectively, with step size h . Alternatively, (3) can be obtained from (6) via a simple discretization of the momentum term (see Appendix C).

We can provide an alternative definition of Bregman momentum in terms of the dual of F function. If $F^*(\mathbf{w}^*) = \sup_{\tilde{\mathbf{w}} \in \mathcal{C}} (\tilde{\mathbf{w}}^\top \mathbf{w}^* - F(\tilde{\mathbf{w}}))$ denotes the Fenchel dual of F and $\mathbf{w} = \arg \sup_{\tilde{\mathbf{w}} \in \mathcal{C}} (\tilde{\mathbf{w}}^\top \mathbf{w}^* - F(\tilde{\mathbf{w}}))$, then the following relation holds between the pair of dual variables $(\mathbf{w}, \mathbf{w}^*)$:

$$\mathbf{w} = f^*(\mathbf{w}^*), \quad \mathbf{w}^* = f(\mathbf{w}), \quad \text{and} \quad f^* = f^{-1}. \quad (7)$$

Taking the derivative of $\mathbf{w}(t)$ and $\mathbf{w}^*(t)$ w.r.t. t yields:

$$\dot{\mathbf{w}}(t) = \dot{f}^*(\mathbf{w}^*(t)) = \mathbf{H}_{F^*}(\mathbf{w}^*(t)) \dot{\mathbf{w}}^*(t), \quad (8) \quad \dot{\mathbf{w}}^*(t) = \dot{f}(\mathbf{w}(t)) = \mathbf{H}_F(\mathbf{w}(t)) \dot{\mathbf{w}}(t). \quad (9)$$

This pairing allows rewriting the Bregman momentum in its dual form:

$$\dot{D}_F(\mathbf{w}(t), \mathbf{w}_0) = \dot{D}_{F^*}(\mathbf{w}_0^*, \mathbf{w}^*(t)) = (\mathbf{w}^*(t) - \mathbf{w}_0^*)^\top \mathbf{H}_{F^*}(\mathbf{w}^*(t)) \dot{\mathbf{w}}^*(t). \quad (10)$$

⁵An equivalent integral form of the CMD update is $\mathbf{w}(t) = f^{-1} \left(f(\mathbf{w}_s) - \int_{z=s}^t \nabla L(\mathbf{w}(z)) dz \right)$.

⁶The objective of (3) is essentially a discretization of the objective of (6). See Appendix C

⁷That is, the value of one variable does not depend on changes in the other.

An expanded derivation is given in Appendix A. Using (9), we can rewrite the CMD update (1) as

$$\dot{\mathbf{w}}(t) = -\mathbf{H}_F^{-1}(\mathbf{w}(t)) \nabla L(\mathbf{w}(t)), \quad (\text{NGD}) \quad (11)$$

i.e. a natural gradient descent (NGD) update [Amari, 1998] w.r.t. the Riemannian metric \mathbf{H}_F . Using $\nabla L(\mathbf{w}) = \mathbf{H}_{F^*}(\mathbf{w}^*) \nabla_{\mathbf{w}^*} L \circ f^*(\mathbf{w}^*)$ and $\mathbf{H}_F(\mathbf{w}) = \mathbf{H}_{F^*}^{-1}(\mathbf{w}^*)$, the CMD update (1) can be written equivalently in the dual domain \mathbf{w}^* as an NGD update w.r.t. the Riemannian metric \mathbf{H}_{F^*} , or by applying (8) as a CMD with the link f^* :

$$\dot{\mathbf{w}}^*(t) = -\mathbf{H}_{F^*}^{-1}(\mathbf{w}^*(t)) \nabla_{\mathbf{w}^*} L \circ f^*(\mathbf{w}^*(t)), \quad (12) \quad \dot{f}^*(\mathbf{w}^*(t)) = -\nabla_{\mathbf{w}^*} L \circ f^*(\mathbf{w}^*(t)). \quad (13)$$

The equivalence of the primal-dual updates was already shown in [Warmuth and Jagota, 1998] for the continuous case and in [Raskutti and Mukherjee, 2015] for the discrete case (where it only holds in one direction). We will show that the equivalence relation is a special case of the reparameterization theorem, introduced in the next section. In the following, we discuss the projected CMD updates for the constrained setting.

Proposition 1. *The CMD update with the additional constraint $\psi(\mathbf{w}(t)) = \mathbf{0}$ for some function $\psi : \mathbb{R}^d \rightarrow \mathbb{R}^m$ s.t. $\{\mathbf{w} \in \mathcal{C} \mid \psi(\mathbf{w}(t)) = \mathbf{0}\}$ is non-empty, amounts to the projected gradient update*

$$\dot{f}(\mathbf{w}(t)) = -\mathbf{P}_\psi(\mathbf{w}(t)) \nabla L(\mathbf{w}(t)) \ \& \ \dot{f}^*(\mathbf{w}^*(t)) = -\mathbf{P}_\psi(\mathbf{w}(t))^\top \nabla L \circ f^*(\mathbf{w}^*(t)), \quad (14)$$

where $\mathbf{P}_\psi := \mathbf{I}_d - \mathbf{J}_\psi^\top (\mathbf{J}_\psi \mathbf{H}_F^{-1} \mathbf{J}_\psi^\top)^{-1} \mathbf{J}_\psi \mathbf{H}_F^{-1}$ is the projection matrix onto the tangent space of F at $\mathbf{w}(t)$ and $\mathbf{J}_\psi(\mathbf{w}(t))$. Equivalently, the update can be written as a projected natural gradient descent update

$$\dot{\mathbf{w}}(t) = -\mathbf{P}_\psi^\top(\mathbf{w}(t)) \mathbf{H}_F^{-1}(\mathbf{w}(t)) \nabla L(\mathbf{w}(t)) \ \& \ \dot{\mathbf{w}}^*(t) = -\mathbf{P}_\psi \mathbf{H}_{F^*}^{-1}(\mathbf{w}^*(t)) \nabla L \circ f^*(\mathbf{w}^*(t)). \quad (15)$$

Example 1 (Normalized) EG). *The unnormalized EG update is motivated using the link function $f(\mathbf{w}) = \log \mathbf{w}$. Adding the linear constraint $\psi(\mathbf{w}) = \mathbf{w}^\top \mathbf{1} - 1$ to the unnormalized EG update results in the (normalized) EG update [Kivinen and Warmuth, 1997]. Since $\mathbf{J}_\psi(\mathbf{w}) = \mathbf{1}^\top$ and $\mathbf{H}_F(\mathbf{w})^{-1} = \text{diag}(\mathbf{w})$, $\mathbf{P}_\psi = \mathbf{I} - \frac{\mathbf{1}\mathbf{1}^\top \text{diag}(\mathbf{w})}{\mathbf{1}^\top \text{diag}(\mathbf{w}) \mathbf{1}} = \mathbf{I} - \mathbf{1}\mathbf{w}^\top$ and the projected CMD update (15) (the continuous EG update) and its NGD form become*

$$\begin{aligned} \log(\mathbf{w}) &= -(\mathbf{I} - \mathbf{1}\mathbf{w}^\top) \nabla L(\mathbf{w}) = -(\nabla L(\mathbf{w}) - \mathbf{1}\mathbf{w}^\top \nabla L(\mathbf{w})), \\ \dot{\mathbf{w}} &= -(\text{diag}(\mathbf{w}) \nabla L(\mathbf{w}) - \mathbf{w}\mathbf{w}^\top \nabla L(\mathbf{w})). \end{aligned}$$

3 Reparameterization

We now establish the second main result of the paper.

Theorem 2 (Main result #2). *Let F and G be strictly convex, continuously-differentiable functions with domains in \mathbb{R}^d and \mathbb{R}^k , respectively, s.t. $k \geq d$. Let $q : \mathbb{R}^k \rightarrow \mathbb{R}^d$ be a reparameterization function expressing parameters \mathbf{w} of F uniquely as $q(\mathbf{u})$ where \mathbf{u} lies in the domain of G . Then the CMD update on parameter \mathbf{w} for the convex function F (with link $f(\mathbf{w}) = \nabla F(\mathbf{w})$) and loss $L(\mathbf{w})$,*

$$\dot{f}(\mathbf{w}(t)) = -\nabla L(\mathbf{w}(t)),$$

coincides with the CMD update on parameters \mathbf{u} for the convex function G (with link $g(\mathbf{u}) := \nabla G(\mathbf{u})$) and the composite loss $L \circ q$,

$$\dot{g}(\mathbf{u}(t)) = -\nabla_{\mathbf{u}} L \circ q(\mathbf{u}(t)),$$

provided that $\mathbf{w}(0) = q(\mathbf{u}(0))$ and $\text{range}(q) \subseteq \text{dom}(F)$ hold, and we have

$$\mathbf{H}_F^{-1}(\mathbf{w}) = \mathbf{J}_q(\mathbf{u}) \mathbf{H}_G^{-1}(\mathbf{u}) \mathbf{J}_q(\mathbf{u})^\top, \text{ for all } \mathbf{w} = q(\mathbf{u}).$$

Proof. Note that (dropping t for simplicity) we have $\dot{\mathbf{w}} = \frac{\partial \mathbf{w}}{\partial \mathbf{u}} \dot{\mathbf{u}} = \mathbf{J}_q(\mathbf{u}) \dot{\mathbf{u}}$ and $\nabla_{\mathbf{u}} L \circ q(\mathbf{u}) = \mathbf{J}_q(\mathbf{u})^\top \nabla L(\mathbf{w})$. The CMD update on \mathbf{u} with the link function $g(\mathbf{u})$ can be written in the NGD form as $\dot{\mathbf{u}} = -\mathbf{H}_G^{-1}(\mathbf{u}) \nabla_{\mathbf{u}} L \circ q(\mathbf{u})$. Thus,

$$\dot{\mathbf{w}} = -\mathbf{H}_G^{-1}(\mathbf{u}) \mathbf{J}_q(\mathbf{u})^\top \nabla_{\mathbf{u}} L(\mathbf{w}).$$

Multiplying by $\mathbf{J}_q(\mathbf{u})$ from the left yields

$$\dot{\mathbf{w}} = -\mathbf{J}_q(\mathbf{u}) \mathbf{H}_G^{-1}(\mathbf{u}) \mathbf{J}_q(\mathbf{u})^\top \nabla_{\mathbf{u}} L(\mathbf{w}).$$

Comparing the result to (11) concludes the proof. \square

In the following examples, we will mainly consider reparameterizing a CMD update with the link function $f(\mathbf{w})$ as a GD update on \mathbf{u} , for which we have $\mathbf{H}_G = \mathbf{I}_k$.

Example 2 (EGU as GD). *The continuous-time EGU can be reparameterized as continuous GD with the reparameterization function $\mathbf{w} = q(\mathbf{u}) = 1/4 \mathbf{u} \odot \mathbf{u} = 1/4 \mathbf{u}^{\odot 2}$, i.e.*

$$\dot{\log}(\mathbf{w}) = -\nabla L(\mathbf{w}) \text{ equals } \dot{\mathbf{u}} = -\frac{\nabla L \circ q(\mathbf{u})}{\nabla_{\mathbf{u}} L(1/4 \mathbf{u}^{\odot 2})} = -1/2 \mathbf{u} \odot \nabla L(\mathbf{w}).$$

This is proven by verifying the condition of Theorem 2:

$$\mathbf{J}_q(\mathbf{u})\mathbf{J}_q(\mathbf{u})^\top = 1/2 \text{diag}(\mathbf{u}) (1/2 \text{diag}(\mathbf{u}))^\top = \text{diag}(1/4 \mathbf{u}^{\odot 2}) = \text{diag}(\mathbf{w}) = \mathbf{H}_F^{-1}(\mathbf{w}).$$

Example 3 (Reduced EG in 2-dimension). *Consider the 2-dimensional normalized weights $\mathbf{w} = [\omega, 1 - \omega]^\top$ where $0 \leq \omega \leq 1$. The normalized reduced EG update [Warmuth and Jagota, 1998] is motivated by the link function $f(w) = \log \frac{w}{1-w}$, thus $H_F(w) = \frac{1}{w} + \frac{1}{1-w} = \frac{1}{w(1-w)}$. This update can be reparameterized as a GD update on $u \in \mathbb{R}$ via $\omega = q(u) = 1/2(1 + \sin(u))$ i.e.*

$$\dot{\log}\left(\frac{w}{1-w}\right) = -\nabla_w L(w) \text{ equals } \dot{u} = -\frac{\nabla_u L \circ q(u)}{\nabla_u L(1/2(1+\sin(u)))} = -\frac{\cos(u)}{2} \nabla L(w).$$

This is verified by checking the condition of Theorem 2: $J_q(u) = 1/2 \cos(u)$ and

$$J_q(u)J_q(u)^\top = \frac{1}{4} \cos^2(u) = \frac{1}{2} (1 + \sin(u)) \frac{1}{2} (1 - \sin(u)) = w(1-w) = H_F^{-1}(w).$$

Open problem The generalization of the reduced EG link function to $d > 2$ dimensions becomes $f(\mathbf{w}) = \log \frac{w}{1 - \sum_{i=1}^{d-1} w_i}$ which utilizes the first $(d-1)$ -dimensions \mathbf{w} s.t. $[\mathbf{w}^\top, w_d]^\top \in \Delta^{d-1}$. Reparameterizing the CMD update using this link as CGD is open. The update can be reformulated as

$$\dot{\mathbf{w}} = -\left(\text{diag}(\mathbf{1} \otimes \mathbf{w}) + \frac{1}{1 - \sum_{i=1}^{d-1} w_i} \mathbf{1}\mathbf{1}^\top\right)^{-1} \nabla L(\mathbf{w}) = -(\text{diag}(\mathbf{w}) - \mathbf{w}\mathbf{w}^\top) \nabla L(\mathbf{w}).$$

We will give a d -dimensional version of EG using a projection onto a constraint in Example 6.

Example 4 (Burg updates as GD). *The update associated with the negative Burg entropy $F(\mathbf{w}) = -\sum_{i=1}^d \log w_i$ and link $f(\mathbf{w}) = -\mathbf{1} \otimes \mathbf{w}$ is reparameterized as GD with $\mathbf{w} = q(\mathbf{u}) := \exp(\mathbf{u})$, i.e.*

$$(-\mathbf{1} \otimes \mathbf{w}) = -\nabla L(\mathbf{w}) \text{ equals } \dot{\mathbf{u}} = -\frac{\nabla L \circ q(\mathbf{u})}{\nabla_{\mathbf{u}} L(\exp(\mathbf{u}))} = -\exp(\mathbf{u}) \odot \nabla L(\mathbf{w}),$$

This is verified by the condition of Theorem 2: $\mathbf{H}_F(\mathbf{w}) = \text{diag}(\mathbf{1} \otimes \mathbf{w})^2$, $\mathbf{J}_q(\mathbf{u}) = \text{diag}(\exp(\mathbf{u}))$, and

$$\mathbf{J}_q(\mathbf{u})\mathbf{J}_q(\mathbf{u})^\top = \text{diag}(\exp(\mathbf{u}))^2 = \text{diag}(\mathbf{w})^2 = \mathbf{H}_F^{-1}(\mathbf{w}).$$

Example 5 (EGU as Burg). *The reparameterization step can be chained, and applied in reverse, when the reparameterization function q is invertible. For instance, we can first apply the inverse reparameterization of the Burg update as GD from Example 4 i.e. $\mathbf{u} = q^{-1}(\mathbf{w}) = \log \mathbf{w}$. Subsequently, applying the reparameterization of EGU as GD from Example 2 i.e. $\mathbf{v} = \tilde{q}(\mathbf{u}) = 1/4 \mathbf{u}^{\odot 2}$, results in the reparameterization of EGU update on \mathbf{v} as Burg update on \mathbf{w} , that is,*

$$\dot{\log}(\mathbf{v}) = -\nabla L(\mathbf{v}) \text{ equals } (-\mathbf{1} \otimes \mathbf{w}) = -\frac{\nabla_{\mathbf{w}} L \circ \tilde{q} \circ q^{-1}(\mathbf{w})}{\nabla_{\mathbf{w}} L(1/4(\log \mathbf{w})^{\odot 2})} = -(\log(\mathbf{w}) \otimes (2\mathbf{w})) \odot \nabla L(\mathbf{v}).$$

For completeness, we also provide the constrained reparameterized updates (proof in Appendix B).

Theorem 3. *The constrained CMD update (14) coincides with the reparameterized projected gradient update on the composite loss,*

$$\dot{\mathbf{g}}(\mathbf{u}(t)) = -\mathbf{P}_{\psi \circ q}(\mathbf{u}(t)) \nabla_{\mathbf{u}} L \circ q(\mathbf{u}(t)),$$

where $\mathbf{P}_{\psi \circ q} := \mathbf{I}_k - \mathbf{J}_{\psi \circ q}^\top (\mathbf{J}_{\psi \circ q} \mathbf{H}_G^{-1} \mathbf{J}_{\psi \circ q}^\top)^{-1} \mathbf{J}_{\psi \circ q} \mathbf{H}_G^{-1}$ is the projection matrix onto the tangent space at $\mathbf{u}(t)$ and $\mathbf{J}_{\psi \circ q}(\mathbf{u}) := \mathbf{J}_q^\top(\mathbf{u}) \mathbf{J}_\psi(\mathbf{w})$.

Example 6 (EG as GD). We now extend the reparameterization of the EGU update as GD in Example 2 to the normalized case in terms of a projected GD update. Combining $q(\mathbf{u}) = 1/4 \mathbf{u}^{\odot 2}$ with $\psi(\mathbf{w}) = \mathbf{1}^\top \mathbf{w} - 1$, we have $\mathbf{J}_{\psi \circ q}(\mathbf{u}) = 1/2 \text{diag}(\mathbf{u}) \mathbf{1}^\top = 1/2 \mathbf{u}^\top$ and $\mathbf{P}_{\psi \circ q}(\mathbf{u}) = \mathbf{I} - \frac{1/4 \mathbf{u} \mathbf{u}^\top}{1/4 \|\mathbf{u}\|^2} = \mathbf{I} - 1/4 \mathbf{u} \mathbf{u}^\top$. Thus,

$$\dot{\mathbf{u}} = -(\mathbf{I} - 1/4 \mathbf{u} \mathbf{u}^\top) \nabla_{\mathbf{u}} L(1/4 \mathbf{u}^{\odot 2}) \text{ with } \mathbf{w}(t) = 1/4 \mathbf{u}(t)^{\odot 2},$$

equals the normalized EG update in Example 2. Note that similar ideas were explored in an evolutionary game theory context in [Sandholm 2010].

4 Tempered Updates

In this section, we consider a richer class of examples derived using the tempered relative entropy divergence [Amid et al. 2019], parameterized by a *temperature* $\tau \in \mathbb{R}$. As we will see, the tempered updates allow interpolating between many well-known cases. We start with the tempered logarithm link function [Naudts, 2002]:

$$f_\tau(\mathbf{w}) = \log_\tau(\mathbf{w}) = \frac{1}{1-\tau} (\mathbf{w}^{1-\tau} - 1), \quad (16)$$

for $\mathbf{w} \in \mathbb{R}_{\geq 0}^d$ and $\tau \in \mathbb{R}$. The \log_τ function is shown in Figure 1 for different values of $\tau \geq 0$. Note that $\tau = 1$ recovers the standard log function as a limit point. The $\log_\tau(\mathbf{w})$ link function is the gradient of the convex function

$$F_\tau(\mathbf{w}) = \sum_i (w_i \log_\tau w_i + \frac{1}{2-\tau} (1 - w_i^{2-\tau})) = \sum_i \left(\frac{1}{(1-\tau)(2-\tau)} w_i^{2-\tau} - \frac{1}{1-\tau} w_i + \frac{1}{2-\tau} \right).$$

The convex function F_τ induces the following tempered Bregman divergence⁸:

$$\begin{aligned} D_{F_\tau}(\tilde{\mathbf{w}}, \mathbf{w}) &= \sum_i \left(\tilde{w}_i \log_\tau \tilde{w}_i - \tilde{w}_i \log_\tau w_i - \frac{\tilde{w}_i^{2-\tau} - w_i^{2-\tau}}{2-\tau} \right) \\ &= \frac{1}{1-\tau} \sum_i \left(\frac{\tilde{w}_i^{2-\tau} - w_i^{2-\tau}}{2-\tau} - (\tilde{w}_i - w_i) w_i^{1-\tau} \right). \end{aligned} \quad (17)$$

For $\tau = 0$, we obtain the squared Euclidean divergence $D_{F_0}(\tilde{\mathbf{w}}, \mathbf{w}) = \frac{1}{2} \|\tilde{\mathbf{w}} - \mathbf{w}\|_2^2$ and for $\tau = 1$, the relative entropy $D_{F_1}(\tilde{\mathbf{w}}, \mathbf{w}) = \sum_i (\tilde{w}_i \log(\tilde{w}_i/w_i) - \tilde{w}_i + w_i)$ (See [Amid et al. 2019] for an extensive list of examples).

In the following, we derive the CMD updates using the time derivative of (17) as the tempered Bregman momentum. Notice that the link function $\log_\tau(x)$ is only defined for $x \geq 0$ when $\tau > 0$. In order to have a weight $\mathbf{w} \in \mathbb{R}^d$, we use the \pm -trick [Kivinen and Warmuth, 1997] by maintaining two non-negative weights \mathbf{w}_+ and \mathbf{w}_- and setting $\mathbf{w} = \mathbf{w}_+ - \mathbf{w}_-$. We call this the *tempered EGU* \mathbb{E}^\pm_τ updates, which contain the standard EGU $^\pm$ updates as a special case of $\tau = 1$. As our final main result, we show that that continuous tempered EGU $^\pm_\tau$ updates interpolate between continuous-time GD and continuous EGU (for $\tau \in [0, 1]$). Furthermore, these updates can be simulated by continuous GD on a new set of parameters \mathbf{u} using a simple reparameterization. We show that reparameterizing the tempered updates as GD updates on the composite loss $L \circ q$ changes the implicit bias of GD, making the updates converge to the solution with the smallest $L_{2-\tau}$ -norm for arbitrary $\tau \in [0, 1]$.

4.1 Tempered EGU and Reparameterization

We first introduce the generalization of the EGU update using the tempered Bregman divergence (17). Let $\mathbf{w}(t) \in \mathbb{R}_{\geq 0}^d$. The tempered EGU update is motivated by

$$\operatorname{argmin}_{\text{curve } \mathbf{w}(t) \in \mathbb{R}_{\geq 0}^d} \left\{ \dot{D}_{F_\tau}(\mathbf{w}(t), \mathbf{w}_0) + L(\mathbf{w}(t)) \right\}.$$

This results in the CMD update

$$\log_\tau \dot{\mathbf{w}}(t) = -\nabla L(\mathbf{w}(t)). \quad (18)$$

⁸The second form is more commonly known as β -divergence [Cichocki and Amari, 2010] with $\beta = 2 - \tau$.




Figure 1: $\log_\tau(x)$, for different $\tau \geq 0$.

An equivalent integral version of this update is

$$\mathbf{w}(t) = \exp_{\tau} \left(\log_{\tau} \mathbf{w}_0 - \int_0^t \nabla_{\mathbf{w}} L(\mathbf{w}(z)) dz \right), \quad (19)$$

where $\exp_{\tau}(x) := [1 + (1 - \tau)x]_{+}^{\frac{1}{1-\tau}}$ is the inverse of tempered logarithm (16). Note that $\tau = 1$ is a limit case which recovers the standard exp function and the update (18) becomes the standard EGU update. Additionally, the GD update (on the non-negative orthant) is recovered at $\tau = 0$. As a result, the tempered EGU $_{\tau}$ update (18) interpolates between GD and EGU for $\tau \in [0, 1]$ and generalizes beyond for values of $\tau > 1$ and $\tau < 0$.⁹ We now show the reparameterization of the tempered EGU $_{\tau}$ update (18) as GD. This corresponds to continuous-time gradient descent on the network of Figure 2

Proposition 2 (Main result #3). *The tempered continuous EGU $_{\tau}$ update can be reparameterized continuous-time GD with the reparameterization function*

$$\mathbf{w} = q_{\tau}(\mathbf{u}) = \left(\frac{2 - \tau}{2} \right)^{\frac{2}{2-\tau}} |\mathbf{u}|^{\odot \frac{2}{2-\tau}}, \text{ for } \mathbf{u} \in \mathbb{R}^d \text{ and } \tau \neq 2. \quad (20)$$

That is

$$\log_{\tau}(\dot{\mathbf{w}}) = -\nabla L(\mathbf{w}) \text{ equals } \dot{\mathbf{u}} = -\underbrace{\nabla L \circ q_{\tau}}_{\nabla_{\mathbf{u}} L \left(\left(\frac{2-\tau}{2} \right)^{\frac{2}{2-\tau}} |\mathbf{u}|^{\odot \frac{2}{2-\tau}} \right)} = -\text{sign}(\mathbf{u}) \odot \left(\frac{2 - \tau}{2} \right)^{\frac{\tau}{2-\tau}} |\mathbf{u}|^{\odot \frac{\tau}{2-\tau}} \odot \nabla L(\mathbf{w}).$$

Proof. This is verified by checking the condition of Theorem 2. The lhs is

$$(\mathbf{H}_{F_{\tau}(\mathbf{w})}(\mathbf{w}))^{-1} = (\mathbf{J}_{\log_{\tau}}(\mathbf{w}))^{-1} = (\text{diag}(\mathbf{w})^{-\tau})^{-1} = \text{diag}(\mathbf{w})^{\tau}.$$

Note that the Jacobian of q_{τ} is

$$\mathbf{J}_{q_{\tau}}(\mathbf{u}) = \left(\frac{2 - \tau}{2} \right)^{\frac{\tau}{2-\tau}} \text{diag}(\text{sign}(\mathbf{u}) \odot |\mathbf{u}|^{\odot \frac{\tau}{2-\tau}}) = \text{diag}(\text{sign}(\mathbf{u}) \odot q_{\tau}(\mathbf{u})^{\odot \frac{\tau}{2}}).$$

Thus the rhs $\mathbf{J}_{q_{\tau}}(\mathbf{u}) \mathbf{J}_{q_{\tau}}^{\top}(\mathbf{u})$ of the condition equals $\text{diag}(\mathbf{w}^{\odot \tau})$ as well. \square

4.2 Minimum-norm Solutions

We apply the (reparameterized) tempered EGU $_{\tau}^{\pm}$ update on the underdetermined linear regression problem. For this, we first consider the \pm -trick on (18), in which we set $\mathbf{w}(t) = \mathbf{w}_{+}(t) - \mathbf{w}_{-}(t)$ where

$$\log_{\tau} \dot{\mathbf{w}}_{+}(t) = -\nabla_{\mathbf{w}} L(\mathbf{w}(t)), \quad \log_{\tau} \dot{\mathbf{w}}_{-}(t) = +\nabla_{\mathbf{w}} L(\mathbf{w}(t)). \quad (21)$$

Note that using the \pm -trick, we have $\mathbf{w}(t) \in \mathbb{R}^n$. We call the updates (21) the *tempered EGU $_{\tau}^{\pm}$* . The reparameterization of the tempered EGU $_{\tau}^{\pm}$ updates as GD can be written by applying Proposition 2

$$\dot{\mathbf{u}}_{+}(t) = -\nabla_{\mathbf{u}_{+}} L(q_{\tau}(\mathbf{u}_{+}(t)) - q_{\tau}(\mathbf{u}_{-}(t))), \quad \dot{\mathbf{u}}_{-}(t) = -\nabla_{\mathbf{u}_{-}} L(q_{\tau}(\mathbf{u}_{+}(t)) - q_{\tau}(\mathbf{u}_{-}(t))), \quad (22)$$

and setting $\mathbf{w}(t) = q_{\tau}(\mathbf{u}_{+}(t)) - q_{\tau}(\mathbf{u}_{-}(t))$.

The strong convexity of the F_{τ} function w.r.t. the $L_{2-\tau}$ -norm (see [Amid et al., 2019]) suggests that the updates motivated by the tempered Bregman divergence (17) yield the minimum $L_{2-\tau}$ -norm solution in certain settings. We verify this by considering the following underdetermined linear regression problem. Let $\{\mathbf{x}_n, y_n\}_{n=1}^N$ where $\mathbf{x}_n \in \mathbb{R}^d$, $y_n \in \mathbb{R}$ denote the set of input-output pairs and let $\mathbf{X} \in \mathbb{R}^{N \times d}$ with $N < d$ be the design matrix for which the n -th row is equal to \mathbf{x}_n^{\top} . Also, let $\mathbf{y} \in \mathbb{R}^N$ denote the vector of targets. Consider the tempered EGU $_{\tau}^{\pm}$ updates (21) on the weights $\mathbf{w}(t) = \mathbf{w}_{+}(t) - \mathbf{w}_{-}(t)$ where $\mathbf{w}_{+}(t), \mathbf{w}_{-}(t) \geq \mathbf{0}$ and $\mathbf{w}_{+}(0) = \mathbf{w}_{-}(0) = \mathbf{w}_0$. Following (19), we have

$$\mathbf{w}_{+}(t) = \exp_{\tau} \left(\log_{\tau} \mathbf{w}_0 - \int_0^t \mathbf{X}^{\top} \delta(z) dz \right), \quad \mathbf{w}_{-}(t) = \exp_{\tau} \left(\log_{\tau} \mathbf{w}_0 + \int_0^t \mathbf{X}^{\top} \delta(z) dz \right),$$

where $\delta(t) = \mathbf{X}(\mathbf{w}_{+}(t) - \mathbf{w}_{-}(t))$.

⁹For example, $\tau = 2$ corresponds to the Burg updates (Example 4).




Figure 2: A reparameterized linear neuron where $w_i = |u_i|^{\frac{2}{2-\tau}}$ as a two-layer sparse network: value of $\tau = 0$ reduces to GD while $\tau = 1$ simulates the EGU update.

Theorem 4. Consider the underdetermined linear regression problem where $N < d$. Let $\mathcal{E} = \{\mathbf{w} \in \mathbb{R}^d \mid \mathbf{X}\mathbf{w} = \mathbf{y}\}$ be the set of solutions with zero error. Given $\mathbf{w}(\infty) \in \mathcal{E}$, then the tempered EGU_τ^\pm updates (21) with temperature $0 \leq \tau \leq 1$ and initial solution $\mathbf{w}_0 = \alpha \mathbf{1} \geq 0$ converge to the minimum $L_{2-\tau}$ -norm solution in \mathcal{E} in the limit $\alpha \rightarrow 0$.

Proof. We show that the solution of the tempered EGU_τ^\pm satisfies the dual feasibility and complementary slackness KKT conditions for the following optimization problem (omitting t for simplicity):

$$\min_{\mathbf{w}_+, \mathbf{w}_-} \|\mathbf{w}_+ - \mathbf{w}_-\|_{2-\tau}^{2-\tau}, \quad \text{for } 0 \leq \tau \leq 1, \quad \text{s.t. } \mathbf{X}(\mathbf{w}_+ - \mathbf{w}_-) = \mathbf{y} \quad \text{and} \quad \mathbf{w}_+, \mathbf{w}_- \geq \mathbf{0}.$$

Imposing the constraints using a set of Lagrange multipliers $\boldsymbol{\nu}_+, \boldsymbol{\nu}_- \geq \mathbf{0}$ and $\boldsymbol{\lambda} \in \mathbb{R}$, we have

$$\min_{\mathbf{w}} \sup_{\boldsymbol{\nu}_+, \boldsymbol{\nu}_- \geq \mathbf{0}, \boldsymbol{\lambda}} \left\{ \|\mathbf{w}_+ - \mathbf{w}_-\|_{2-\tau}^{2-\tau} + \boldsymbol{\lambda}^\top (\mathbf{X}(\mathbf{w}_+ - \mathbf{w}_-) - \mathbf{y}) - \mathbf{w}_+^\top \boldsymbol{\nu}_+ - \mathbf{w}_-^\top \boldsymbol{\nu}_- \right\}.$$

The set of KKT conditions are

$$\begin{cases} \mathbf{w}_+, \mathbf{w}_- \geq \mathbf{0}, \quad \mathbf{X}\mathbf{w} = \mathbf{y}, \\ + \text{sign}(\mathbf{w}) \odot |\mathbf{w}|^{\odot(1-\tau)} - \mathbf{X}^\top \boldsymbol{\lambda} \geq \mathbf{0}, \quad - \text{sign}(\mathbf{w}) \odot |\mathbf{w}|^{\odot(1-\tau)} + \mathbf{X}^\top \boldsymbol{\lambda} \geq \mathbf{0}, \\ (\text{sign}(\mathbf{w}) \odot |\mathbf{w}|^{\odot(1-\tau)} - \mathbf{X}^\top \boldsymbol{\lambda}) \odot \mathbf{w}_+ = \mathbf{0}, \quad (\text{sign}(\mathbf{w}) \odot |\mathbf{w}|^{\odot(1-\tau)} - \mathbf{X}^\top \boldsymbol{\lambda}) \odot \mathbf{w}_- = \mathbf{0}, \end{cases}$$

where $\mathbf{w} = \mathbf{w}_+ - \mathbf{w}_-$. The first condition is imposed by the form of the updates and the second condition is satisfied by the assumption at $t \rightarrow \infty$. Using $\mathbf{w}_0 = \alpha \mathbf{1}$ with $\alpha \rightarrow 0$, we have

$$\begin{aligned} \mathbf{w}_+(t) &= \exp_\tau \left(-\frac{1}{1-\tau} - \int_0^t \mathbf{X}^\top \boldsymbol{\delta}(z) dz \right) = \left[-(1-\tau) \mathbf{X}^\top \int_0^t \boldsymbol{\delta}(z) dz \right]_+^{\odot \frac{1}{1-\tau}}, \\ \mathbf{w}_-(t) &= \exp_\tau \left(-\frac{1}{1-\tau} + \int_0^t \mathbf{X}^\top \boldsymbol{\delta}(z) dz \right) = \left[+(1-\tau) \mathbf{X}^\top \int_0^t \boldsymbol{\delta}(z) dz \right]_+^{\odot \frac{1}{1-\tau}}. \end{aligned}$$

Setting $\boldsymbol{\lambda} = -(1-\tau) \int_0^\infty \boldsymbol{\delta}(z) dz$ satisfies the remaining KKT conditions. \square

Corollary 1. Under the assumptions of Theorem 4 the reparameterized tempered EGU_τ^\pm updates (22) also recover the minimum $L_{2-\tau}$ -norm solution where $\mathbf{w}(t) = q_\tau(\mathbf{u}_+(t)) - q_\tau(\mathbf{u}_-(t))$.

This corollary shows that reparameterizing the loss in terms of the parameters \mathbf{u} changes the implicit bias of the GD updates. Similar results were observed before in terms of sparse signal recovery [Vaskevicius et al., 2019] and matrix factorization [Gunasekar et al., 2017]. Here, we show that this is a direct result of the dynamics induced by the reparameterization Theorem 2.

5 Conclusion and Future Work

In this paper, we motivated the continuous-time mirror descent updates and provided a general framework for reparameterizing these updates. We also introduced the tempered EGU_τ^\pm updates and its reparameterizations. The tempered EGU_τ^\pm updates include the two commonly used gradient descent and exponentiated gradient updates, and interpolations between them. For the underdetermined linear regression problem we showed that under certain conditions, the tempered EGU_τ^\pm updates converge to the minimum $L_{2-\tau}$ -norm solution. The current work leads to many interesting future directions:

- The focus in this paper was to develop the reparameterization method in full generality. Our reparameterization equivalence theorem holds only in the continuous-time and the equivalence relation breaks down after discretization. However, in many important cases the discretized reparameterized updates closely track the discretized original updates [Amid and Warmuth, 2020]. This was done by proving the same on-line worst case regret bounds for the discretized reparameterized updates as the originals. A key research direction is to find general conditions for which this is true.
- Perhaps the most important application of the current work is reparameterizing the weights of deep neural networks for achieving sparse solutions or obtaining an implicit form of regularization that mimics a trade-off between the ridge and lasso methods (e.g. elastic net regularization [Zou and Hastie, 2005]). Here the key open question is the following: Are sparse networks (as in Figure 2) necessary, if the goal is to learn sample efficiently when the solution is sparse?¹⁰
- A general treatment of the convergence results for underdetermined linear regression should allow any start vector. Also, developing a matrix form of the reparameterization theorem is left open.

¹⁰This has recently been proven in [Warmuth et al., 2020].

Acknowledgement

Thanks to Zakaria Mhammedi for valuable feedback and to NSF grant IIS-1546459 for partial support.

Broader Impact

The result of the paper suggests that the mirror descent updates can be effectively used in neural networks by running backpropagation on the reparameterized form of the neurons. This may have a potential use case for training these networks more efficiently. This is a theoretical paper and the broader ethical impact discussion is not applicable.

References

- Ethan Akin. *The geometry of population genetics*, volume 31 of *Lecture Notes in Biomathematics*. Springer-Verlag, Berlin-New York, 1979.
- Shun-ichi Amari. Natural gradient works efficiently in learning. *Neural Computation*, 10(2):251–276, 1998.
- Ehsan Amid and Manfred K Warmuth. Winnowing with gradient descent. In *Conference on Learning Theory (COLT)*, pages 163–182. PMLR, 2020.
- Ehsan Amid, Manfred K. Warmuth, Rohan Anil, and Koren Tomer. Robust bi-tempered logistic loss based on Bregman divergences. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems, NeurIPS’19*, Cambridge, MA, USA, 2019.
- William L. Burke. *Applied Differential Geometry*. Cambridge University Press, 1985.
- Andrzej Cichocki and Shun-ichi Amari. Families of alpha-beta-and gamma-divergences: Flexible and robust measures of similarities. *Entropy*, 12(6):1532–1568, 2010.
- Udaya Ghai, Elad Hazan, and Yoram Singer. Exponentiated gradient meets gradient descent. In *Conference on Learning Theory (COLT)*, volume 117, pages 1–23. PMLR, 2020.
- Suriya Gunasekar, Blake E. Woodworth, Srinadh Bhojanapalli, Behnam Neyshabur, and Nati Srebro. Implicit regularization in matrix factorization. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 6151–6159, 2017.
- Suriya Gunasekar, Jason D. Lee, Daniel Soudry, and Nati Srebro. Implicit bias of gradient descent on linear convolutional networks. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 9461–9471, 2018.
- Jyrki Kivinen and Manfred K. Warmuth. Exponentiated gradient versus gradient descent for linear predictors. *Information and Computation*, 132(1):1–63, 1997.
- Jyrki Kivinen, Manfred K. Warmuth, and Peter Auer. The Perceptron algorithm vs. Winnow: linear vs. logarithmic mistake bounds when few input variables are relevant. *Artificial Intelligence*, 97: 325–343, December 1997.
- Jyrki Kivinen, Manfred K Warmuth, and Babak Hassibi. The p-norm generalization of the LMS algorithm for adaptive filtering. *IEEE Transactions on Signal Processing*, 54(5):1782–1793, 2006.
- Nick Littlestone and Manfred K. Warmuth. The weighted majority algorithm. *Information and Computation*, 108(2):212–261, 1994.
- Jan Naudts. Deformed exponentials and logarithms in generalized thermostatics. *Physica A*, 316 (1–4):323–334, 2002.
- Arkadi Nemirovski and David B. Yudin. *Problem Complexity and Method Efficiency in Optimization*. Wiley & Sons, New York, 1983.

- Jiazhong Nie, Wojciech Kotłowski, and Manfred K Warmuth. Online PCA with optimal regret. *The Journal of Machine Learning Research*, 17(1):6022–6070, 2016.
- Maxim Raginsky and Jake Bouverie. Continuous-time stochastic mirror descent on a network: Variance reduction, consensus, convergence. In *2012 IEEE 51st IEEE Conference on Decision and Control (CDC)*, pages 6793–6800. IEEE, 2012.
- Garvesh Raskutti and Sayan Mukherjee. The information geometry of mirror descent. *IEEE Transactions on Information Theory*, 61(3):1451–1457, 2015.
- R Tyrrell Rockafellar. Monotone operators and the proximal point algorithm. *SIAM journal on control and optimization*, 14(5):877–898, 1976.
- William H Sandholm. *Population games and evolutionary dynamics*. MIT Press, 2010.
- Shai Shalev-Shwartz et al. Online learning and online convex optimization. *Foundations and Trends® in Machine Learning*, 4(2):107–194, 2012.
- Tomas Vaskevicius, Varun Kanade, and Patrick Rebeschini. Implicit regularization for optimal sparse recovery. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 2968–2979, 2019.
- Manfred K. Warmuth and Arun Jagota. Continuous and discrete time nonlinear gradient descent: relative loss bounds and convergence. In R. Greiner E. Boros, editor, *Electronic Proceedings of Fifth International Symposium on Artificial Intelligence and Mathematics*. Electronic, <http://rutcor.rutgers.edu/~amai>, 1998.
- Manfred K. Warmuth and S.V.N. Vishwanathan. Leaving the span. In *Proceedings of the 18th Annual Conference on Learning Theory (COLT)*, pages 366–381, 2005.
- Manfred K. Warmuth, Wojciech Kotłowski, and Shuisheng Zhou. Kernelization of matrix updates. *Journal of Theoretical Computer Science*, 558:159–178, 2014. Special issue for the 23rd International Conference on Algorithmic Learning Theory (ALT).
- Manfred K Warmuth, Wojciech Kotłowski, and Ehsan Amid. A case where a spindly two-layer linear network whips any neural network with a fully connected input layer. *arXiv preprint arXiv:2010.08625*, 2020.
- Hui Zou and Trevor Hastie. Regularization and variable selection via the elastic net. *Journal of the royal statistical society: series B (statistical methodology)*, 67(2):301–320, 2005.