

A ADDITIONAL DEFINITIONS

Definition 3: A function $f(\cdot)$ is L -smooth, if there exists a positive constant L such that

$$f(y) \leq f(x) + \langle \nabla f(x), y - x \rangle + \frac{L}{2} \|y - x\|^2, \forall x, y \in \mathcal{X}.$$

Definition 4: A convex function $f(\cdot)$ is λ -strongly convex with respect to some norm $\|\cdot\|$, if there exists a positive constant λ such that

$$f(y) + \langle \nabla f(y), x - y \rangle + \frac{\lambda}{2} \|x - y\|^2 \leq f(x), \forall x, y \in \mathcal{X}.$$

Definition 5: A function $f(\cdot)$ is Lipschitz continuous with factor G if for all x and y in \mathcal{X} , the following holds:

$$|f(x) - f(y)| \leq G \|x - y\|, \forall x, y \in \mathcal{X}.$$

B PROOF OF LEMMA 1

Consider single-step mirror descent update as follows:

$$x = \operatorname{argmin}_{y \in \mathcal{X}} \left\{ f_t(x') + \langle \nabla f_t(x'), y - x' \rangle + \frac{1}{\alpha} D_r(y, x') \right\}. \quad (7)$$

Strong convexity of the above minimization objective implies

$$\begin{aligned} f_t(x') + \langle \nabla f_t(x'), x - x' \rangle + \frac{1}{\alpha} D_r(x, x') &\leq \\ f_t(x') + \langle \nabla f_t(x'), y - x' \rangle + \frac{1}{\alpha} D_r(y, x') - \frac{1}{\alpha} D_r(y, x), &\forall y \in \mathcal{X}. \end{aligned} \quad (8)$$

Furthermore, from the smoothness condition, we have

$$f_t(x) \leq f_t(x') + \langle \nabla f_t(x'), x - x' \rangle + \frac{L}{2} \|x - x'\|^2. \quad (9)$$

Substituting equation 9 into equation 8, and setting $y = x_t^*$, we obtain

$$\begin{aligned} f_t(x) - \frac{L}{2} \|x - x'\|^2 + \frac{1}{\alpha} D_r(x, x') &\leq \\ f_t(x') + \langle \nabla f_t(x'), x_t^* - x' \rangle + \frac{1}{\alpha} D_r(x_t^*, x') - \frac{1}{\alpha} D_r(x_t^*, x). \end{aligned} \quad (10)$$

Since $\alpha \leq \frac{1}{L}$, and regularization function $r(\cdot)$ is 1-strongly convex, we have

$$\frac{1}{\alpha} D_r(x, x') \geq L D_r(x, x') \geq \frac{L}{2} \|x - x'\|^2. \quad (11)$$

Next, we exploit the strong convexity of the cost function, i.e.,

$$f_t(x') + \langle \nabla f_t(x'), x_t^* - x' \rangle \leq f_t(x_t^*) - \lambda D_r(x_t^*, x'). \quad (12)$$

Combining equation 10, equation 11, and equation 12, we obtain

$$f_t(x) \leq f_t(x_t^*) - \lambda D_r(x_t^*, x') + \frac{1}{\alpha} D_r(x_t^*, x') - \frac{1}{\alpha} D_r(x_t^*, x). \quad (13)$$

Next, we use the result of (Hazan & Kale, 2014), which states that for every λ -strongly convex function $f_t(\cdot)$, the following bound holds:

$$f_t(x) - f_t(x_t^*) \geq \lambda D_r(x_t^*, x), \quad (14)$$

where $x_t^* = \operatorname{argmin}_{x \in \mathcal{X}} f_t(x)$. Combining the above with equation 13, we obtain

$$D_r(x_t^*, x) \leq \beta D_r(x_t^*, x'), \quad (15)$$

where $\beta = 1 - \frac{2\lambda\alpha}{1+\lambda\alpha}$. \square

C PROOF OF THEOREM 2

C.1 KEY LEMMAS

The following two lemmas pave the way for our regret analysis leading to Theorem 2. Lemma 7 presents an alternative form for the mirror descent update.

Lemma 7 *Suppose there exists z_{t+1} that satisfies $\nabla r(z_{t+1}) = \nabla r(x_t) - \alpha \nabla f_t(x_t)$, for some strongly convex function $r(\cdot)$, and step size α . Then, the following updates are equivalent*

$$x_{t+1} = \operatorname{argmin}_{x \in \mathcal{X}} D_r(x, z_{t+1}), \quad (16)$$

$$x_{t+1} = \operatorname{argmin}_{x \in \mathcal{X}} \left\{ \langle \nabla f_t(x_t), x \rangle + \frac{1}{\alpha} D_r(x, x_t) \right\}. \quad (17)$$

Proof. We begin by expanding equation 16 as follows:

$$\begin{aligned} x_{t+1} &= \operatorname{argmin}_{x \in \mathcal{X}} \{r(x) - r(z_{t+1}) - \langle \nabla r(z_{t+1}), x - z_{t+1} \rangle\} \\ &= \operatorname{argmin}_{x \in \mathcal{X}} \{r(x) - \langle \nabla r(z_{t+1}), x \rangle\} \\ &= \operatorname{argmin}_{x \in \mathcal{X}} \{r(x) - \langle \nabla r(x_t) - \alpha \nabla f_t(x_t), x \rangle\} \\ &= \operatorname{argmin}_{x \in \mathcal{X}} \{\alpha \langle \nabla f_t(x_t), x \rangle + r(x) - r(x_t) - \langle \nabla r(x_t), x - x_t \rangle\} \\ &= \operatorname{argmin}_{x \in \mathcal{X}} \left\{ \langle \nabla f_t(x_t), x \rangle + \frac{1}{\alpha} D_r(x, x_t) \right\}. \end{aligned} \quad (18)$$

Thus, the update in equation 16 is equivalent to equation 17. \square

Lemma 8 *Under the same convexity and smoothness condition stated in Theorem 2, let x_t be the sequence of decisions generated by OMMD. Then, the following bound holds:*

$$\|x_{t+1} - x_t^*\| \leq \sqrt{L_r \beta^M} \|x_t - x_t^*\|, \quad (19)$$

where L_r is the smoothness factor of the regularization function $r(\cdot)$, and β is the shrinking factor obtained in Lemma 1.

Proof. Using the result of Lemma 1, OMMD with M mirror descent steps guarantees

$$D_r(x_t^*, x_{t+1}) \leq \beta^M D_r(x_t^*, x_t). \quad (20)$$

Since the regularization function $r(\cdot)$ is 1-strongly convex, we have

$$\frac{\|x_t^* - x_{t+1}\|^2}{2} \leq r(x_t^*) - r(x_{t+1}) - \langle \nabla r(x_{t+1}), x_t^* - x_{t+1} \rangle. \quad (21)$$

Next, we exploit the smoothness condition of the regularization function $r(\cdot)$, i.e.,

$$r(x_t^*) - r(x_t) - \langle \nabla r(x_t), x_t^* - x_t \rangle \leq \frac{L_r}{2} \|x_t^* - x_t\|^2. \quad (22)$$

By combining the above with equation 20, and equation 21, and using the definition of Bregman divergence, we obtain

$$\|x_{t+1} - x_t^*\|^2 \leq L_r \beta^M \|x_t - x_t^*\|^2. \quad (23)$$

Taking the square root on both sides of equation 23 completes the proof. \square

C.2 PROOF OF THE THEOREM

Now, we are ready to present the proof of Theorem 2. In this proof, we will use the following properties of Bregman divergence.

(a) By direct substitution, the following equality holds for any $x, y, z \in \mathcal{X}$,

$$\langle \nabla r(z) - \nabla r(y), x - y \rangle = D_r(x, y) - D_r(x, z) + D_r(y, z). \quad (24)$$

(b) If $x = \operatorname{argmin}_{x' \in \mathcal{X}} D_r(x', z)$, i.e., x is the Bregman projection of z into the set \mathcal{X} , then for any arbitrary point $y \in \mathcal{X}$, we have

$$D_r(y, z) \geq D_r(y, x) + D_r(x, z). \quad (25)$$

To bound the dynamic regret, we begin by using the strong convexity of the cost function $f_t(\cdot)$, i.e.,

$$\begin{aligned} f_t(x_t) - f_t(x_t^*) &\leq \langle \nabla f_t(x_t), x_t - x_t^* \rangle - \lambda D_r(x_t^*, x_t) \\ &\leq \frac{1}{\alpha} \langle \nabla r(x_t) - \nabla r(z_{t+1}), x_t - x_t^* \rangle - \lambda D_r(x_t^*, x_t) \\ &\leq \frac{1}{\alpha} \left(D_r(x_t^*, x_t) - D_r(x_t^*, z_{t+1}) + D_r(x_t, z_{t+1}) \right) - \lambda D_r(x_t^*, x_t) \\ &\leq \frac{1}{\alpha} \left(D_r(x_t^*, x_t) - D_r(x_t^*, x_{t+1}) - D_r(x_{t+1}, z_{t+1}) + D_r(x_t, z_{t+1}) \right) - \lambda D_r(x_t^*, x_t) \\ &\leq \left(\frac{1}{\alpha} - \lambda \right) D_r(x_t^*, x_t) + \frac{1}{\alpha} \left(D_r(x_t, z_{t+1}) - D_r(x_{t+1}, z_{t+1}) \right) \\ &\leq \left(\frac{1}{\alpha} - \lambda \right) \left(\frac{f_t(x_t) - f_t(x_t^*)}{\lambda} \right) + \frac{1}{\alpha} \left(D_r(x_t, z_{t+1}) - D_r(x_{t+1}, z_{t+1}) \right), \end{aligned} \quad (26)$$

where in the second line we have used the alternative mirror descent update stated in Lemma 7, i.e., $\nabla f_t(x_t) = (1/\alpha)(\nabla r(x_t) - \nabla r(z_{t+1}))$. To obtain the third line, we have utilized the Bregman divergence property in equation 24. We have used the Bregman projection property in equation 25 in the fourth line. By omitting some negative terms, and using equation 14, we obtain the right-hand side of equation 26.

Thus, if $\alpha > \frac{1}{2\lambda}$, we have

$$\begin{aligned} f_t(x_t) - f_t(x_t^*) &\leq \frac{\lambda}{2\alpha\lambda - 1} \left(D_r(x_t, z_{t+1}) - D_r(x_{t+1}, z_{t+1}) \right) \\ &\stackrel{(a)}{\leq} \frac{\lambda K}{2\alpha\lambda - 1} \|x_{t+1} - x_t\| \\ &\leq \frac{\lambda K}{2\alpha\lambda - 1} \left(\|x_{t+1} - x_t^*\| + \|x_t - x_t^*\| \right) \\ &\stackrel{(b)}{\leq} \frac{\lambda K}{2\alpha\lambda - 1} (1 + \sqrt{L_r \beta^M}) \|x_t - x_t^*\|, \end{aligned} \quad (27)$$

where we have used the Lipschitz continuity of Bregman divergence to obtain inequality (a), and we have applied Lemma 8 to obtain inequality (b). Summing equation 27 over time, we have

$$\operatorname{Reg}_T^d = \sum_{t=1}^T f_t(x_t) - f_t(x_t^*) \leq \frac{\lambda K}{2\alpha\lambda - 1} (1 + \sqrt{L_r \beta^M}) \sum_{t=1}^T \|x_t - x_t^*\|. \quad (28)$$

Now, we proceed to bound $\sum_{t=1}^T \|x_t - x_t^*\|$ as follows:

$$\begin{aligned} \sum_{t=1}^T \|x_t - x_t^*\| &= \|x_1 - x_1^*\| + \sum_{t=2}^T \|x_t - x_t^*\| \\ &\leq \|x_1 - x_1^*\| + \sum_{t=2}^T \|x_t - x_{t-1}^*\| + \|x_{t-1}^* - x_t^*\| \\ &\stackrel{(a)}{\leq} \|x_1 - x_1^*\| + \sum_{t=2}^T \sqrt{L_r \beta^M} \|x_{t-1} - x_{t-1}^*\| + \sum_{t=2}^T \|x_t^* - x_{t-1}^*\|, \end{aligned} \quad (29)$$

where we used the result of Lemma 8 to obtain inequality (a). If $M \geq \lceil (\frac{1}{2} + \frac{1}{2\alpha\lambda}) \log L_r \rceil$, we have

$$\beta^M = \left(1 - \frac{2\alpha\lambda}{1 + \alpha\lambda}\right)^M \leq \exp\left(\frac{-2M\alpha\lambda}{1 + \alpha\lambda}\right) < \frac{1}{L_r}, \quad (30)$$

which implies $L_r\beta^M < 1$. Therefore, by combining equation 29 and equation 30, we have

$$\sum_{t=1}^T \|x_t - x_t^*\| \leq \frac{\|x_1 - x_1^*\|}{1 - \sqrt{L_r\beta^M}} + \frac{\sum_{t=2}^T \|x_t^* - x_{t-1}^*\|}{1 - \sqrt{L_r\beta^M}}. \quad (31)$$

Finally, substituting equation 31 into equation 28 completes the proof. \square

D PROOF OF THEOREM 3

In order to bound the dynamic regret, we begin by the smoothness condition of the cost function $f_t(\cdot)$, i.e.,

$$\begin{aligned} f_t(x_t) - f_t(x_t^*) &\leq \langle \nabla f_t(x_t^*), x_t - x_t^* \rangle + \frac{L}{2} \|x_t - x_t^*\|^2 \\ &\leq \|\nabla f_t(x_t^*)\|_* \|x_t - x_t^*\| + \frac{L}{2} \|x_t - x_t^*\|^2. \end{aligned} \quad (32)$$

Next, we use the fact

$$\|\nabla f_t(x_t^*)\|_* \|x_t - x_t^*\| \leq \frac{\|\nabla f_t(x_t^*)\|_*^2}{2\theta} + \frac{\theta \|x_t - x_t^*\|^2}{2}, \quad (33)$$

for any arbitrary positive constant $\theta > 0$. Thus, we have

$$f_t(x_t) - f_t(x_t^*) \leq \frac{\|\nabla f_t(x_t^*)\|_*^2}{2\theta} + \frac{(L + \theta) \|x_t - x_t^*\|^2}{2}. \quad (34)$$

Summing equation 34 over time, we obtain

$$\text{Reg}_T^d = \sum_{t=1}^T f_t(x_t) - f_t(x_t^*) \leq \sum_{t=1}^T \frac{\|\nabla f_t(x_t^*)\|_*^2}{2\theta} + \frac{L + \theta}{2} \sum_{t=1}^T \|x_t - x_t^*\|^2. \quad (35)$$

Now, we proceed by bounding $\sum_{t=1}^T \|x_t - x_t^*\|^2$ as follows:

$$\begin{aligned} \sum_{t=1}^T \|x_t - x_t^*\|^2 &= \|x_1 - x_1^*\|^2 + \sum_{t=2}^T \|x_t - x_{t-1}^* + x_{t-1}^* - x_t^*\|^2 \\ &\leq \|x_1 - x_1^*\|^2 + \sum_{t=2}^T (2\|x_t - x_{t-1}^*\|^2 + 2\|x_{t-1}^* - x_t^*\|^2) \\ &\leq \|x_1 - x_1^*\|^2 + 2\beta^M L_r \sum_{t=1}^T \|x_t - x_{t-1}^*\|^2 + 2 \sum_{t=2}^T \|x_{t-1}^* - x_t^*\|^2. \end{aligned} \quad (36)$$

We note that if $M \geq \lceil (\frac{1}{2} + \frac{1}{2\alpha\lambda}) \log 2L_r \rceil$, then $2\beta^M L_r < 1$. Therefore, from equation 36 we can obtain

$$\sum_{t=1}^T \|x_t - x_t^*\|^2 \leq \frac{\|x_1 - x_1^*\|^2}{1 - 2\beta^M L_r} + \frac{2}{1 - 2\beta^M L_r} \sum_{t=2}^T \|x_t^* - x_{t-1}^*\|^2. \quad (37)$$

Substituting equation 37 into equation 35 completes the proof. \square

E PROOF OF THEOREM 5

The proof of Theorem 5 initially follows the first half of the proof of Theorem 2, which is repeated here for completeness.

To analyze the dynamic regret, we first use the strong convexity of the cost function $f_t(\cdot)$, i.e.,

$$\begin{aligned}
f_t(x_t) - f_t(x_t^*) &\leq \langle \nabla f_t(x_t), x_t - x_t^* \rangle - \lambda D_r(x_t^*, x_t) \\
&\leq \frac{1}{\alpha} \langle \nabla r(x_t) - \nabla r(z_{t+1}), x_t - x_t^* \rangle - \lambda D_r(x_t^*, x_t) \\
&\leq \frac{1}{\alpha} \left(D_r(x_t^*, x_t) - D_r(x_t^*, z_{t+1}) + D_r(x_t, z_{t+1}) \right) - \lambda D_r(x_t^*, x_t) \\
&\leq \frac{1}{\alpha} \left(D_r(x_t^*, x_t) - D_r(x_t^*, z_{t+1}) - D_r(x_{t+1}, z_{t+1}) + D_r(x_t, z_{t+1}) \right) - \lambda D_r(x_t^*, x_t) \\
&\leq \left(\frac{1}{\alpha} - \lambda \right) D_r(x_t^*, x_t) + \frac{1}{\alpha} \left(D_r(x_t, z_{t+1}) - D_r(x_{t+1}, z_{t+1}) \right) \\
&\leq \left(\frac{1}{\alpha} - \lambda \right) \left(\frac{f_t(x_t) - f_t(x_t^*)}{\lambda} \right) + \frac{1}{\alpha} \left(D_r(x_t, z_{t+1}) - D_r(x_{t+1}, z_{t+1}) \right), \quad (38)
\end{aligned}$$

where in the second line we have used the alternative mirror descent update stated in Lemma 7, i.e., $\nabla f_t(x_t) = (1/\alpha)(\nabla r(x_t) - \nabla r(z_{t+1}))$. To obtain the third line, we have utilized the Bregman divergence property in equation 24. We have used the Bregman projection property in equation 25 in the fourth line. By omitting some negative terms, and using equation 14, we obtain the right-hand side of equation 38.

Therefore, if $\alpha > \frac{1}{2\lambda}$, we have

$$f_t(x_t) - f_t(x_t^*) \leq \frac{\lambda}{2\alpha\lambda - 1} (D_r(x_t, z_{t+1}) - D_r(x_{t+1}, z_{t+1})). \quad (39)$$

Now we continue to bound $D_r(x_t, z_{t+1})$. By the definition of Bregman divergence, we have

$$\begin{aligned}
D_r(x_t, z_{t+1}) + D_r(z_{t+1}, x_t) &= \langle \nabla r(x_t) - \nabla r(z_{t+1}), x_t - z_{t+1} \rangle \\
&= \langle \alpha \nabla f_t(x_t), x_t - z_{t+1} \rangle \\
&\leq \|\alpha \nabla f_t(x_t)\|_* \|x_t - z_{t+1}\| \\
&\leq \frac{\alpha^2}{2} \|\nabla f_t(x_t)\|_*^2 + \frac{\|x_t - z_{t+1}\|^2}{2}. \quad (40)
\end{aligned}$$

The strong convexity of the regularization function implies

$$\frac{\|x_t - z_{t+1}\|^2}{2} \leq r(z_{t+1}) - r(x_t) - \langle \nabla r(x_t), z_{t+1} - x_t \rangle = D_r(z_{t+1}, x_t). \quad (41)$$

Combining the above with equation 40, we obtain

$$D_r(x_t, z_{t+1}) \leq \frac{\alpha^2}{2} \|\nabla f_t(x_t)\|_*^2. \quad (42)$$

By substituting equation 42 into equation 40, and summing over time, we have

$$\text{Reg}_T^d = \sum_{t=1}^T f_t(x_t) - f_t(x_t^*) \leq \frac{\alpha^2 \lambda}{4\alpha\lambda - 2} \|\nabla f_t(x_t)\|_*^2. \quad (43)$$

□

F CLOSED-FORM UPDATE FOR MIRROR DESCENT

In this section, we derive the close-form mirror descent update in equation 6.

Let $r(y) = \sum_{j=1}^d y_j \log(y_j)$ be the negative entropy. Then, we have

$$\begin{aligned} D_r(y, y_t^i) &= \sum_{j=1}^d [y_j \log(y_j) - y_{t,j}^i \log(y_{t,j}^i) - (\log(y_{t,j}^i) + 1)(y_j - y_{t,j}^i)] \\ &= \sum_{j=1}^d y_j \log\left(\frac{y_j}{y_{t,j}^i}\right) + \langle 1, y - y_t^i \rangle = D_{KL}(y, y_t^i), \end{aligned} \quad (44)$$

where $y_{t,j}^i$ denotes the j -th component of the decision vector y_t^i , and $D_{KL}(y, y_t^i)$ represents the KL divergence between y and y_t^i .

Now consider the update in equation 5, which can be written as follows:

$$\begin{aligned} &\text{minimize}_{y \in \mathcal{X}} \quad \langle \nabla f_t(y_t^i), y \rangle + \frac{1}{\alpha} \sum_{j=1}^d y_j \log\left(\frac{y_j}{y_{t,j}^i}\right) \\ &\text{subject to} \quad \langle 1, y \rangle = 1, y \geq 0. \end{aligned} \quad (45)$$

The Lagrangian of the above problem is given by

$$L(y, \lambda, \gamma) = \langle \nabla f_t(y_t^i), y \rangle + \sum_{j=1}^d \left[\frac{1}{\alpha} y_j \log\left(\frac{y_j}{y_{t,j}^i}\right) + \lambda y_j - \gamma_j y_j \right] - \lambda, \quad (46)$$

where $\lambda \in \mathbb{R}$ and $\gamma \in \mathbb{R}_+^d$ are Lagrange multipliers corresponding to the constraints. Next, we take derivative with respect to y to obtain

$$\frac{\partial}{\partial y_j} L(y, \lambda, \gamma) = \nabla f_t(y_t^i)_j + \frac{1}{\alpha} \log(y_j) + \frac{1}{\alpha} - \frac{1}{\alpha} \log(y_{t,j}^i) + \lambda - \gamma_j. \quad (47)$$

Setting the above to zero results in the following closed-form update:

$$y_{t,j}^{i+1} = \frac{y_{t,j}^i \exp(-\alpha \nabla f_t(y_t^i))}{\sum_{j=1}^d y_{t,j}^i \exp(-\alpha \nabla f_t(y_t^i))}. \quad (48)$$

□

G ADDITIONAL EXPERIMENTS

In this section, we present additional experiments to study the performance of OMMD. In the first experiment, we use the MNIST dataset. In the second experiment, we consider a switching problem where the cost function switches between two quadratic functions after a specific number of rounds.

First, we consider the well-known MNIST digits dataset, where every data sample ω is an image of size 28×28 pixel that can be represented by a 784-dimensional vector, i.e., $d = 784$. Each sample corresponds to one of the digits in $\{0, 1, \dots, 9\}$, and thus, there are $c = 10$ different classes. The goal of the learner is to classify streaming digit images in an online fashion.

We consider a robust regression problem, where the cost function for each data sample is given by

$$f(x, (\omega_i, z_i)) = \|\omega_i^T x - z_i\|_1^2,$$

where x is the optimization variable, belonging to the constraint set is $\mathcal{X} = \{x : x \in \mathbb{R}_+^n, \|x\|_1 = 1\}$.

We use the negative entropy regularization function, i.e., $r(x) = \sum_{i=1}^d x_i \log(x_i)$, which is strongly convex with respect to the $l1$ -norm. We set the step size $\alpha = 0.1$.

From Fig. 3, we again observe that OMMD consistently outperforms the other alternatives. In particular, compared with DMD, applying $M = 10$ steps of mirror descent can reduce the dynamic regret up to 20%. We also see that the dynamic regret grows linearly with the number of rounds, which is a natural consequence of steady fluctuation in the sequence of dynamic minimizers x_t^* as explained before.




Figure 3: Dynamic regret comparison on MNIST dataset.




Figure 4: Dynamic regret comparison for switching cost.

Next, we consider the case where the cost function switches between two functions. Both functions are in the quadratic form $f_t(x) = \|A_t x - b_t\|_2^2$, where $A_t \in \mathbb{R}^{d \times d}$, and $b_t \in \mathbb{R}^d$. In particular, we assume that the parameter A_t is chosen among

$$A_t^{(1)} = \text{diag}\left(\underbrace{\frac{1}{t^{p_1}}, \frac{1}{t^{p_1}}, \dots, \frac{1}{t^{p_1}}}_{d_1}, \underbrace{0, 0, \dots, 0}_{d_2}\right), \text{ and } A_t^{(2)} = \text{diag}\left(\underbrace{0, 0, \dots, 0}_{d_1}, \underbrace{\frac{1}{t^{p_1}}, \frac{1}{t^{p_1}}, \dots, \frac{1}{t^{p_1}}}_{d_2}\right),$$

such that $d_1 + d_2 = d$, and $b_t = [\frac{1}{t^{p_2}}, \dots, \frac{1}{t^{p_2}}]'$. Therefore, at each round the cost function is either $f_t^{(1)}(x) = \|A_t^{(1)}x - b_t\|_2^2$ or $f_t^{(2)}(x) = \|A_t^{(2)}x - b_t\|_2^2$. We assume that the cost function switches between $f_t^{(1)}(\cdot)$ and $f_t^{(2)}(\cdot)$ every τ rounds. In our experiment, we set $d_1 = 10$, $d = 1000$, $p_1 = 0.9$, and $p_2 = 0.1$. We further set the switching period $\tau = 10$, and parameter $\alpha = 0.02$. The dynamic regret roughly reflects the accumulated mismatch error over time.

In Fig. 4, we compare the performance of OMMD with that of other alternatives in terms of the dynamic regret. OMMD with $M = 10$ nearly halves the dynamic regret of DMD after 300 rounds. Furthermore, the benefit of applying multiple steps of mirror descent can be significant even for smaller values of M .