Federated Learning of Oligonucleotide Drug Molecule Thermodynamics with Differentially Private ADMM-Based SVM

Shirin Tavara[1(✉)], Alexander Schliep[1], and Debabrota Basu[2]

[1] CSE, University of Gothenburg | Chalmers University of Technology, Gothenburg, Sweden
tavara@chalmers.se

[2] Équipe Scool, Inria, UMR 9189-CRIStAL, CNRS, Univ. Lille, Centrale Lille, Lille, France

Abstract. A crucial step to assure drug safety is predicting off-target binding. For oligonucleotide drugs this requires learning the relevant thermodynamics from often large-scale data distributed across different organisations. This process will respect data privacy if distributed and private learning under limited and private communication between local nodes is used. We propose an ADMM-based SVM with differential privacy for this purpose. We empirically show that this approach achieves accuracy comparable to the non-private one, i.e. ∼86%, while yielding tight empirical privacy guarantees even after convergence.

Keywords: Differential privacy · Distributed learning · Federated learning · Oligonucleotide drug molecules · ADMM · SVM

1 Introduction

Machine learning (ML) and AI have been a remarkable success in improving small molecule drug discovery over the past decade. Advanced ML models have been applied to generate novel molecular structures, suggest synthesis pathways, or predict biochemical properties or interactions of those molecules [6,15,16]. The advances have led to open-source software [3] and cross-industry co-operations centered around federated learning [20]. While small molecules interacting with proteins are the most frequently used drug modality, alternatives are explored to address unmet clinical needs. In particular, advances in chemistry for the novel class of drugs based on antisense oligonucleotides (ASOs) have shown promise to deliver successful therapies to the clinic [2]. In ASOs, similar to modalities like small interfering RNAs (siRNAs), the fundamental reaction is one of hybridization, or binding, of two oligonucleotides. The main difference to most small molecule drugs is that oligonucleotide-based drugs modulate the expression of genes post transcription by leveraging inter-cellular mechanisms to degrade the

© Springer Nature Switzerland AG 2021
M. Kamp et al. (Eds.): ECML PKDD 2021 Workshops, CCIS 1525, pp. 459–467, 2021.
https://doi.org/10.1007/978-3-030-93733-1_34

gene's transcript, i.e. its mRNA, before it is translated to a protein. Thermodynamics of oligonucleotides, i.e. structure formation of individual RNA molecules, and hybridization of two oligonucleotides, have been studied extensively: the former to improve understanding of RNA molecules, and the latter partly because DNA microarrays, the main platform for gene expression measurements before the advent of high-throughput sequencing, rely on hybridization. Consequently, there is a wide range of data sets, biophysical models and software available [19].

Thermodynamics are a crucial factor *both* in drug efficacy—binding to the intended target—and drug safety—assuring no binding to unintended targets [33]. These models however do not reflect the many different chemical modifications proposed for ASOs, which often have a stabilizing effect [22] and thus alter the binding energy even without changes to the oligonucleotide sequences. Experiments to elucidate the binding energies are performed by many pharmaceutical companies and academic laboratories, but due to intellectual property interests in both oligonucleotide sequences and gene targets, this information cannot be shared easily. This suggests that a privacy-preserving federated learning approach would allow development of improved ML models for prediction of thermodynamics from the entirety of data, while protecting respective IP rights.

For our proof of concept, the machine learning task we consider is that of predicting off-target hybridization between an oligonucleotide and an mRNA binding site with mismatches in its sequence. We consider the free energy as the indicator of hybridization [33]. The mismatches are represented as features capturing different aspects of sequence dissimilarity. In lieu of experimental data, we use the widely popular nearest-neighbor model [38] as implemented in the Vienna RNA package [19] to simulate data.

We propose to use an SVM [24] to classify the occurrence of binding and non-binding, and an ADMM-based [29] framework to train in a distributed manner. We adopt a differentially private mechanism [9] that respects privacy of individual nodes and local data while learning and communicating the final SVM model (Sect. 3). This conglomeration allows us to operate over a large number of nodes and data, while experimentally achieving comparable accuracy as the non-private SVM and a reasonably tight privacy after global convergence (Sect. 4).

2 Background

In this section, we provide a brief overview of the methods: SVM, ADMM, and Differential Privacy, which are essential for developing our methodology.

Support Vector Machines (SVM). Support vector machines (SVM) are a supervised machine learning framework to solve classification [24,30,32,36] and regression problems [7,11,26]. Let us consider a binary classification problem with dataset D, i.e. training samples $\{\mathbf{x}_i\}_{i=1}^N$ and corresponding class labels $\{y_i\}_{i=1}^N \in \{+1,-1\}^N$. SVMs separate two classes of training data by finding a *hyperplane*, $h(\mathbf{x}) \triangleq \mathbf{w}^T\Phi(\mathbf{x})+b$, with the maximum distance, or *margin*, from the closest points on either side of $h(\mathbf{x})$. In training, a primal optimization problem as in Eq. (1) is solved to determine the hyperplane with the maximum margin.

$$\min_{\mathbf{w},b} \frac{1}{2}||\mathbf{w}||^2 + C\sum_{i=1}^{N} \xi_i \quad \text{s.t. } y_i(\mathbf{w}^T \Phi(\mathbf{x}_i) + b) \geq 1 - \xi_i, \xi_i \geq 0, \forall i \in [N]. \quad (1)$$

Here, $\mathbf{w} \in \mathbb{R}^D$ represents the weight vector of the model which is orthogonal to the hyperplane, ξ_i's are the classification error, b is the bias parameter, and $\Phi(\mathbf{x}) : \mathbb{R}^N \rightarrow \mathbb{R}^D$ is a map function. If the training samples are linearly separable, the map function is linear, i.e. $\Phi(\mathbf{x}) = \mathbf{x}$. In more complex and nonlinear scenarios, SVM uses a kernel-based mapping function Φ, e.g. Radial Basis Functions (RBFs), to map training data into a high-dimensional space, where linear separation of data is possible [18,27].

Alternating Direction Method of Multipliers (ADMM). In many real-life applications, including the application we consider, the data is often large and distributed across different organisations at different locations. This naturally calls for incorporating a distributed framework to train the SVM where local models are built in different local servers and the local servers internally communicate to form a global model. ADMM [13,29] provides a robust, scalable, and parallelizable framework to train SVM models on such distributed data, where each node has limited access to the local datasets $\{D_i\}_{i=1}^M$ and can only communicate to certain neighbouring nodes[1]. In this case, an individual node $i \in [M]$ solves a local problem

$$\min_{\mathbf{w}_i} f_i(\mathbf{w}_i) \quad \text{s.t. } \mathbf{w}_i - \mathbf{w}_j = 0, \forall i \in \mathcal{N}_i, \quad (2)$$

such that the global objective function to optimize is $f(\mathbf{w}) \triangleq \sum_{i=1}^M f_i(\mathbf{w}_i)$, and $\mathcal{N}_i$ is the set of one-hop neighboring nodes of the node i. The constraints in (2) require that local variables $\mathbf{w}_i$'s should agree across one-hop neighbors $\mathbf{w}_j$'s. This asks for communication between neighbouring nodes. Detailed information regarding convergence and robustness of ADMM can be found in [4,12]. ADMM's dependence on the network between nodes and their communications can be found in [29,31]. We use ADMM to train our SVM model over distributed data.

Differential Privacy (DP). The final SVM model learnt using ADMM and the internal communication between nodes during the interactive procedure of ADMM posit new challenges for protecting data privacy of involved users. Specially, considering privacy is imperative if the data contains sensitive and classified information such as financial, medical records etc. [37], which should be protected against unauthorized disclosure. DP [10] is a widely-studied and scalably deployed statistical framework that enables data privacy. A differentially private algorithm aims to generate outputs that remain almost indistinguishable when an individual's data is added or deleted from the training dataset [14].

Definition 1 (((ϵ, δ) − DP) [10]). *A randomized algorithm $\mathcal{M}$ is (ϵ, δ)-differential private if for any two neighboring datasets $\mathcal{D}$ and $\mathcal{D}'$ that only differ at most in one data point, and for any subset of output $O \in Range(\mathcal{M})$, the*

[1] Here, M is the number of nodes across which the data is distributed and $M \leq N$.

following holds, $\Pr[\mathcal{M}(\mathcal{D}) = O] \leq e^{\epsilon}\Pr[\mathcal{M}(\mathcal{D}') = O] + \delta$, *given privacy level* $\epsilon \geq 0$, *slack* $\delta \geq 0$.

The lower the privacy level ϵ, the harder it is to distinguish the neighbouring datasets, and thus a higher degree of DP is ensured. DP has been used to publicly release medical and genetic data while respecting the data privacy of involved individuals [25,28,34]. This motivates the use of DP as the privacy framework for distributed learning of the hybridization prediction.

3 Methodology

In this section, we present the differentially private ADMM-based SVM algorithm and the thermodynamics data that it is used to learn.

ADMM-Based Training of SVM. We use the ADMM-based SVM algorithm proposed by Forero et al. [12] as the non-private base algorithm. The communication model in the algorithm is decentralized such that each node computes the local optimization problem (Eq. (2)) and then broadcasts the local results to all the one-hop neighboring nodes ($\mathcal{N}_i$). In the aggregation phase, each node receives and gathers the neighboring results to conduct the corresponding updates. The process continues until the result of all nodes reach the consensus defined by the constraints in (2). For further details, please check [31].

Differentially Private ADMM-Based SVM (DP-ADMM-SVM). To incorporate differential privacy (DP) into the ADMM-based SVM algorithm, we adopt the PP-ADMM algorithm [9]. Following PP-ADMM, we add calibrated Gaussian noise into the local objective function $f_i(\mathbf{w}_i)$ at each ADMM iteration. Specifically, each node i generates a random vector η_{i1} from a Gaussian distribution $\mathcal{N}(\mathbf{0}, \sigma_{i1}^2\mathbf{I}_d)$ and adds $(\eta_{i1})^T\mathbf{w}_i$ into the local objective function and obtains an approximated solution $\hat{\mathbf{w}}_i^t$ at iteration t. While communicating the local model $\hat{\mathbf{w}}_i^t$ to the neighboring nodes, another random noise vector η_{i2} is generated from a Gaussian distribution $\mathcal{N}(\mathbf{0}, \sigma_{i2}^2\mathbf{I}_d)$ and added to the approximate solution of the perturbed objective function, i.e., $\mathbf{w}_i^t = \hat{\mathbf{w}}_i^t + \eta_{i2}$. Here, σ_{i1}^2 and σ_{i2}^2 are proportional to ϵ^{-2}, where ϵ is the desired privacy level at every iteration. For further details regarding noise variances σ_{i1} and σ_{i2}, we refer to [9].

Algorithm 1. DP-ADMM-SVM

Input: Dataset $\{D_i\}_{i=1}^M$, initial models $\mathbf{w}_i^0$, privacy parameters ϵ, an optimizer
Initialization: Compute noise variances σ_{i1}^2 & σ_{i2}^2, regularization parameters
while #Iterations $\leq$ MAX_ITER **do**
  For each node, add noise $\eta_{i1} \sim \mathcal{N}(\mathbf{0}, \sigma_{i1}^2\mathbf{I}_d)$ to the local objective $f_i(\mathbf{w}_i)$
  For each node, use the optimizer to obtain updated model $\hat{\mathbf{w}}_i^{t+1}$
  For each node, add noise $\eta_{i2} \sim \mathcal{N}(\mathbf{0}, \sigma_{i2}^2\mathbf{I}_d)$ to $\hat{\mathbf{w}}_i^{t+1}$ & get the local model $\mathbf{w}_i^{t+1}$
  For each node, broadcast the local model $\mathbf{w}_i^{t+1}$ to neighbours $\mathcal{N}_i$
  Update local models of each node using the communicated neighbouring models
end while

4 Experimental Analysis: Results and Discussions

In order to establish that the proposed methodology can scale to a large number of academic and industrial participants, we have decided to use the relatively large number of 224 entities with private data, which maps directly to the available computational resources. In order to study the privacy and accuracy of DP-ADMM-SVM, we conducted three experiments on 479, 136 data points privately distributed over 224 nodes. We implemented DP-ADMM-SVM in C++ and performed the computations on 224 nodes provided by Tetralith, the Swedish National Supercomputer Centre's largest HPC cluster at Linköping University [21]. In this section, we describe the data under investigation, and corresponding insights from the experimental results.

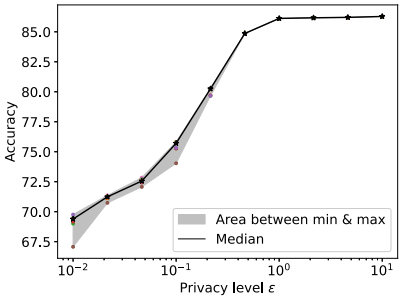

Fig. 1. Accuracy of 10 runs for ADMM algorithm with 50 iterations in logarithmic scale. The grey area shows the area between the minimum and maximum accuracy of 10 runs and the black line is the median accuracy of 10 runs.

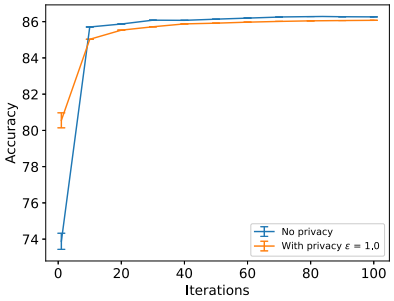

Fig. 2. Average of accuracy per iteration for distributed ADMM-based SVM using 224 nodes, $\epsilon = 1.0$ and 50 ADMM iterations with and without privacy. The error bars show the standard deviation of accuracy.

Data Description. We sample 10,000 binding sites of length 15, i.e. 15-mers, uniformly from the Human transcriptome, i.e. the set of all expressed sequences (GRCh38 reference genome). Their perfect Watson-Crick complements define the oligonucleotides. For each oligo, we simulate 50 binding sites at an edit distance to the perfect match site of up to 5. Note that larger edit distances are highly unlikely to lead to binding. We compute the difference in Gibbs free energy ΔG° between the perfect match duplex and the mismatch duplex using RNAcofold [19]. We use $\Delta G^\circ < \tau$ as the binary indicator variable for the hybridization with $\tau = 7$. The sequence dissimilarity was represented through the following features: edit distance, length of longest common factor, 2- and 3-gram distance, Jaquard distance on 2- and 3-grams, and weighted versions of the q-gram distances, with weights[2] derived from the nearest neighbor model [38]. Randomly

[2] E.g. the 3-gram "GCG" has larger weight due to higher binding affinity than "ATA".

selecting data for each of the 224 nodes resulted in a total of 479,136 data points in the training data set.

Accuracy vs. Privacy Trade-Off. In order to study the effect of privacy on accuracy of DP-ADMM-SVM, we compute the test accuracy of DP-ADMM-SVM after 50 ADMM iterations and repeat 10 times. In each run, we fix $\delta = 10^{-4}$ and vary the privacy level ϵ logarithmically (with base 10) in $[0.01, 10]$. Figure 1 shows the average accuracy of 10 runs on 224 nodes containing a grey area, which is the area between the minimum and maximum average accuracy along with the median of average accuracy of 10 runs in logarithmic scale. Figure 1 shows that the accuracy increases with increase in ϵ for $0.01 \leq \epsilon \leq 1.0$.[3] For $1.0 \leq \epsilon \leq 10.0$, the accuracy remains stable around 86%. Figure 1 also illustrates that the grey area between the maximum and minimum average accuracy over 224 nodes gets tighter as ϵ increases. *This experiment shows that accuracy increases as privacy decreases but after* $\epsilon = 1.0$*, the privacy has negligible effect anymore.*

Cost of Privacy. Now we investigate the cost of DP on ADMM-based SVM by comparing convergence of non-private ADMM-SVM and DP–ADMM-SVM with $\epsilon = 1.0$ over 100 iterations. Figure 2 shows the average test accuracy over 224 nodes and corresponding error bars indicating the standard deviation of accuracies over these nodes. Figure 2 shows that the average accuracy of both DP-ADMM-SVM and ADMM-SVM is almost the same, i.e. 86% after 100 iterations. We observe that though the initial accuracy of DP-ADMM-SVM is higher than that of ADMM-SVM, following iteration 10, it is lower than that of ADMM-SVM. *Both of these are due to the randomization introduced by the DP mechanism. This leads to model stability over iterations and lower accuracy after convergence* [23].

Empirical Estimate of Privacy Level. From the naïve [10] and adaptive [17] composition theorems of differential privacy, we know that the effective DP decreases, i.e. the effective privacy level of the ML model increases, as the number of iterations increase. Following the invent of moment accountant [1], we know that often the privacy level achieved by the dataset under experiment is much tighter than that of these data-independent and worst-case composition theorems. In order to investigate the increase in effective privacy level of DP-ADMM-SVM with increasing number of iterations, we plot the privacy levels ϵ computed using naive and adaptive compositions and a budget accountant as

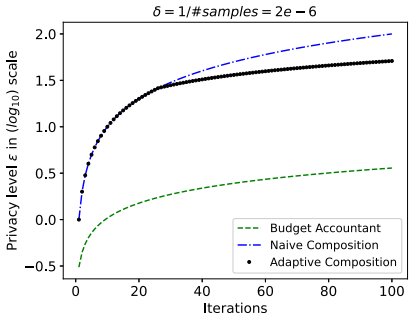

Fig. 3. Data-dependent (Budget Accountant) and data-independent (Naive and Adaptive Composition) privacy levels achieved per iteration.

[3] Increase in privacy level ϵ indicates decrease in differential privacy.

a function of the number of ADMM iterations. We fix the ϵ per-step, i.e. the ϵ used to calibrate the noise, as 1, and $\delta = 1/N$, where the number of samples $N = 479,136$ [5]. Figure 3 shows that the naïve and adaptive compositions predict the worst-case privacy level of the SVM model obtained after 100 iterations as 100.0 and 51.15 respectively. In contrast, *while considering the data, the empirical privacy level of the SVM model is* 3.59*, which is significantly tighter and indicates moderate loss (i.e.* $\epsilon \in [1, 10]$*) in DP.*

5 Conclusion

We propose to use Differentially Private ADMM-based SVM to learn the thermodynamics of oligonucleotide drug molecules from data distributed across multiple nodes and with communication limited to immediate neighbours. We show that DP-ADMM-SVM achieves ∼86% accuracy for privacy level $\epsilon = 1$ at predicting off-target hybridization, a determinant of drug safety. This accuracy level is comparable with that of non-private ADMM-based SVM and does not improve for lower DP, i.e. $\epsilon \geq 1$. We empirically show that the data-dependent privacy level of the final SVM model is $(3.59, 1/\#\text{samples})$. This is significantly tight as a global privacy level where a distributed learning problem over 224 nodes and 479,136 data points evolves over iterations. In future work we want to improve the present DP mechanism such that it is not worst-case [8], does gradient perturbation [35], and is affected and gets stronger by the network between nodes.

Acknowledgments. SSF Strategic Mobility Grant "Drug Discovery for Antisense Oligos" (A.S.), Swedish National Supercomputer Centre (A.S. & S.T.).

References

1. Abadi, M., et al.: Deep learning with differential privacy. In: Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security, pp. 308–318 (2016)
2. Bennett, C.F.: Therapeutic antisense oligonucleotides are coming of age. Annu. Rev. Med. **70**, 307–321 (2019)
3. Blaschke, T., et al.: Reinvent 2.0: an AI tool for de novo drug design. J. Chem. Inf. Mod. **60**(12), 5918–5922 (2020)
4. Boyd, S., Parikh, N., Chu, E., Peleato, B., Eckstein, J.: Distributed optimization and statistical learning via the alternating direction method of multipliers. Found. Trends® Mach. Learn. **3**(1), 1–122 (2011)
5. Canonne, C.: What is δ, and what δ difference does it make? DifferentialPrivacy.org, March 2021. https://differentialprivacy.org/flavoursofdelta/
6. Chen, H., Engkvist, O., Wang, Y., Olivecrona, M., Blaschke, T.: The rise of deep learning in drug discovery. Drug Discov. Today **23**(6), 1241–1250 (2018)
7. Collobert, R., Bengio, S.: Svmtorch: support vector machines for large-scale regression problems. J. Mach. Learn. Res. **1**, 143–160 (2001)
8. Dandekar, A., Basu, D., Bressan, S.: Differential privacy at risk: bridging randomness and privacy budget. In: Proceedings on Privacy Enhancing Technologies, vol. 1, pp. 64–84 (2021)

9. Ding, J., Wang, J., Liang, G., Bi, J., Pan, M.: Towards plausible differentially private ADMM based distributed machine learning. In: Proceedings of the 29th ACM International Conference on Information & Knowledge Management, pp. 285–294 (2020)
10. Dwork, C., McSherry, F., Nissim, K., Smith, A.: Calibrating noise to sensitivity in private data analysis. In: Halevi, S., Rabin, T. (eds.) TCC 2006. LNCS, vol. 3876, pp. 265–284. Springer, Heidelberg (2006). https://doi.org/10.1007/11681878_14
11. Flake, G.W., Lawrence, S.: Efficient SVM regression training with SMO. Mach. Learn. **46**(1), 271–290 (2002)
12. Forero, P.A., Cano, A., Giannakis, G.B.: Consensus-based distributed support vector machines. J. Mach. Learn. Res. **11**, 1663–1707 (2010)
13. França, G., Bento, J.: How is distributed ADMM affected by network topology? ArXiv e-prints, October 2017
14. Harvard: Differential privacy (2021). https://privacytools.seas.harvard.edu/differential-privacy
15. Johansson, S., et al.: AI-assisted synthesis prediction. Drug Discov. Today Technol. **32–33**, 65–72 (2020)
16. Johansson, S.V., et al.: Using active learning to develop machine learning models for reaction yield prediction. ChemRxiv (2021). https://doi.org/10.33774/chemrxiv-2021-bpv0c. Under review
17. Kairouz, P., Oh, S., Viswanath, P.: The composition theorem for differential privacy. In: International Conference on Machine Learning, pp. 1376–1385. PMLR (2015)
18. Leslie, C., Eskin, E., Noble, W.S.: The spectrum kernel: a string kernel for SVM protein classification. In: Biocomputing 2002, pp. 564–575. World Scientific (2001)
19. Lorenz, R., et al.: ViennaRNA package 2.0. Algorithms Mol. Biol. **6**(1), 1–14 (2011)
20. Martin, E.J., Zhu, X.W.: Collaborative profile-QSAR: a natural platform for building collaborative models among competing companies. J. Chem. Inf. Mod. **61**(4), 1603–1616 (2021)
21. NSC: Tetralith (2021). https://www.nsc.liu.se/systems/tetralith/, https://www.nsc.liu.se/systems/tetralith/
22. Papargyri, N., Pontoppidan, M., Andersen, M.R., Koch, T., Hagedorn, P.H.: Chemical diversity of locked nucleic acid-modified antisense oligonucleotides allows optimization of pharmaceutical properties. Mol. Ther. Nucleic Acids **19**, 706–717 (2020)
23. Pinot, R., Yger, F., Gouy-Pailler, C., Atif, J.: A unified view on differential privacy and robustness to adversarial examples (2019)
24. Platt, J.: Sequential minimal optimization: A fast algorithm for training support vector machines. Technical Report MSR-TR-98-14, Microsoft Research, April 1998
25. Raisaro, J.L., et al.: Protecting privacy and security of genomic data in i2b2 with homomorphic encryption and differential privacy. IEEE/ACM Trans. Comput. Biol. Bioinform. **15**(5), 1413–1426 (2018)
26. Shevade, S.K., Keerthi, S.S., Bhattacharyya, C., Murthy, K.R.K.: Improvements to the SMO algorithm for SVM regression. IEEE Trans. Neural Netw. **11**(5), 1188–1193 (2000)
27. Soman, K., Loganathan, R., Ajay, V.: Machine learning with SVM and other kernel methods. PHI Learning Pvt. Ltd. (2009)
28. Sun, Z., Wang, Y., Shu, M., Liu, R., Zhao, H.: Differential privacy for data and model publishing of medical data. IEEE Access **7**, 152103–152114 (2019)
29. Tavara, S.: Parallel computing of support vector machines: a survey. ACM Comput. Surv. (CSUR) **51**(6), 1–38 (2019)

30. Tavara, S., Schliep, A.: Effect of network topology on the performance of ADMM-based SVMs. In: 2018 30th International Symposium on Computer Architecture and High Performance Computing (SBAC-PAD), pp. 388–393. IEEE (2018)
31. Tavara, S., Schliep, A.: Effects of network topology on the performance of consensus and distributed learning of SVMs using ADMM. PeerJ Comput. Sci. **7**, e397 (2021)
32. Tavara, S., Sundell, H., Dahlbom, A.: Empirical study of time efficiency and accuracy of support vector machines using an improved version of PSVM. In: Proceedings of the International Conference on Parallel and Distributed Processing Techniques and Applications (PDPTA), p. 177. The Steering Committee of The World Congress in Computer Science, Computer (2015)
33. Watt, A.T., Swayze, G., Swayze, E.E., Freier, S.M.: Likelihood of nonspecific activity of gapmer antisense oligonucleotides is associated with relative hybridization free energy. Nucleic Acid Ther. **30**(4), 215–228 (2020)
34. Wei, J., Lin, Y., Yao, X., Zhang, J., Liu, X.: Differential privacy-based genetic matching in personalized medicine. IEEE Trans. Emerg. Top. Comput. (2020)
35. Yu, D., Zhang, H., Chen, W., Liu, T.Y., Yin, J.: Gradient perturbation is underrated for differentially private convex optimization. arXiv preprint arXiv:1911.11363 (2019)
36. Zhang, R., Ma, J.: An improved SVM method P-SVM for classification of remotely sensed data. Int. J. Remote Sens. **29**(20), 6029–6036 (2008)
37. Zhang, X., Khalili, M.M., Liu, M.: Improving the privacy and accuracy of ADMM-based distributed algorithms. In: International Conference on Machine Learning, pp. 5796–5805. PMLR (2018)
38. Zuker, M., Mathews, D.H., Turner, D.H.: Algorithms and thermodynamics for RNA secondary structure prediction: a practical guide. In: Barciszewski, J., Clark, B.F.C. (eds.) RNA Biochemistry and Biotechnology. NATO Science Series (Series 3: High Technology), vol. 70, pp. 11–43. Springer, Dordrecht (1999). https://doi.org/10.1007/978-94-011-4485-8_2