The hyperbolic perceptron algorithm as given in the original version of the paper is not guaranteed to converge. Counterexamples can be constructed for data sets that are only separable with a small margin. A recent paper [1] proposes a hyperbolic perceptron algorithm with a modified update rule, which is guaranteed to converge on data that is separable with any margin γ. [1] gives a sample complexity of $O(1/\sinh(\gamma)^2)$, instead of $O(1/\sinh(\gamma))$, as we gave in the original version of our paper.

The corrected version of our paper gives an adversarial perceptron algorithm with the corrected update rule introduced in [1]. Consequently, we corrected the analysis of the adversarial margin and the sample complexity of the adversarial perceptron. We further modified the analysis of the adversarial gradient descent algorithm, which relies on the guarantees for the adversarial perceptron. Here, we retain the polynomial-time convergence guarantee, albeit at a slightly different rate. Formally, the corrected theoretical results are as follows:

- Adversarial margin: $\gamma - \alpha$ (previously $\gamma/\cosh(\alpha)$), see Lemma 4.5.
- Sample complexity of adversarial perceptron: $O(1/(\gamma-\alpha)^2)$ (previously $O(\cosh(\alpha)/\sinh(\gamma))$), see Theorem 4.6.
- Convergence of adversarial gradient descent: $\Omega\left(\text{poly}\left(\frac{1}{\sinh(\gamma-\alpha)}\right)\right)$ (previously $\Omega\left(\text{poly}\left(\frac{\cosh(\alpha)}{\sinh(\gamma)}\right)\right)$), see Theorem 4.4.

The updated arXiv version presents a corrected version for all affected results.

[1] Tabaghi, Puoya, et al. "Linear Classifiers in Mixed Constant Curvature Spaces." arXiv preprint arXiv:2102.10204 (2021).