# Polylogarithmic Approximation for Edit Distance and the Asymmetric Query Complexity 

Alexandr Andoni, Robert Krauthgamer, Krysztof Onak

Theorem 1.1 (Main): For every fixed $\varepsilon>0$, there is an algorithm that approximates the edit distance between two input strings $x, y \in \Sigma^{n}$ within a factor of $(\log n)^{\mathcal{O}(1 / \varepsilon)}$, and runs in $n^{1+\varepsilon}$ time.

Definition: Consider two strings $x, y \in \Sigma^{n}$ for some alphabet $\Sigma$, and let ed $(x, y)$ denote the edit distance between these two strings. The computational problem is the promise problem known as the Distance Threshold Estimation Problem (DTEP): distinguish whether ed $(x, y)>R$ or ed $(x, y) \leq$ $R / \alpha$, where $R>0$ is a parameter (known to the algorithm) and $\alpha \geq 1$ is the approximation factor. We use $\operatorname{DTEP}_{\beta}$ to denote the case of $R=n / \beta$, where $\beta \geq 1$ may be a function of $n$.
Definition: In the asymmetric query model, the algorithm knows in advance (has unrestricted access to) one of the strings, say $y$, and has only query access to the other string, $x$. The asymmetric query complexity of an algorithm is the number of coordinates in $x$ that the algorithm has to probe in order to solve DTEP with success probability at least $2 / 3$.
Theorem 1.2 (Query complexity upper bound): For every $\beta=\beta(n) \geq 2$ and fixed $0<\varepsilon<1$ there is an algorithm that solves $\mathrm{DTEP}_{\beta}$ with approximation $\alpha=(\log n)^{\mathcal{O}(1 / \varepsilon)}$, and makes $\beta n^{\varepsilon}$ asymmetric queries. This algorithm runs in time $\mathcal{O}\left(n^{1+\varepsilon}\right)$.

For every $\beta=\mathcal{O}(1)$ and fixed integer $t \geq 2$ there is an algorithm for $\mathrm{DTEP}_{\beta}$ achieving approximation $\alpha=\mathcal{O}\left(n^{1 / t}\right)$, with $\mathcal{O}\left(\log ^{t-1} n\right)$ queries into $x$.
Theorem 1.3 (Query complexity lower bound): For a sufficiently large constant $\beta>1$, every algorithm that solves $\mathrm{DTEP}_{\beta}$ with approximation $\alpha=\alpha(n)>2$ has asymmetric query complexity $2^{\Omega\left(\frac{\log n}{\log a+\log \log n}\right)}$. Moreover, for every fixed non-integer $t>1$, every algorithm that solves DTEP $_{\beta}$ with approximation $\alpha=n^{1 / t}$ has asymmetric query complexity $\Omega\left(\log ^{\lfloor t\rfloor} n\right)$.

Theorem 3.1: Let $n \geq 2, \beta=\beta(n) \geq 2$, and integer $b=b(n) \geq 2$ be such that $\left(\log _{b} n\right) \in \mathbb{N}$.
There is an algorithm solving $\mathrm{DTEP}_{\beta}$ with approximation $\alpha=\mathcal{O}\left(b \log _{b} n\right)$ and $\beta \cdot(\log n)^{\mathcal{O}\left(\log _{b} n\right)}$ queries into $x$. The algorithm runs in $n \cdot(\log n)^{\mathcal{O}\left(\log _{b} n\right)}$ time.

For every constant $\beta=\mathcal{O}(1)$ and integer $t \geq 2$, there is an algorithm for solving $\mathrm{DTEP}_{\beta}$ with $\mathcal{O}\left(n^{1 / t}\right)$ approximation and $\mathcal{O}\left(\log ^{t-1} n\right)$ queries. The algorithm runs in $\tilde{\mathcal{O}}(n)$ time.

## Characterization of edit distance using $\mathcal{E}$-Distance

For a string $x, x[s: t]$ denotes the substring of $x$ comprising of $x[s], \ldots, x[t-1]$. The characterization may be viewed as a tree of arity $b$, where nodes correspond to substring $x[s: s+l]$. The root is the entire string $x[1: n+1]$. Let $h \stackrel{\text { def }}{=} \log _{b} n \in \mathbb{N}$. Then nodes on level $i$ for $0 \leq i \leq h$ correspond to substrings $x\left[s: s+l_{i}\right]$ of length $l_{i} \xlongequal{\text { def }} n / b^{i}$.
Definition 3.2 ( $\mathcal{E}$-distance): Consider two strings $x, y$ of length $n \geq 2$. Fix $i \in\{0,1, \ldots, h\}, s \in B_{i}=$ $\left\{1,1+l_{i}, \ldots\right\}$, and a position $u \in \mathbb{Z}$.

If $i=h$, the $\mathcal{E}$-distance of $x[s: s+1]$ to the position $u$ is 1 if $u \notin[n]$ or $x[s] \neq y[u]$, and 0 otherwise.
For $i \in\{0,1, \ldots, h-1\}$, we recursively define the $\mathcal{E}$-distance $\mathcal{E}_{x, y}(i, s, u)$ of $x\left[s: s+l_{i}\right]$ to the position $u$ as follows. Partition $x\left[s: s+l_{i}\right]$ into $b$ blocks of length $l_{i+1}=l_{i} / b$, starting at positions $s+j l_{i+j}$, where $j \in\{0,1, \ldots, b-1\}$. Then

$$
\mathcal{E}_{x, y}(i, s, u) \stackrel{\text { def }}{=} \sum_{j=0}^{b-1} \min _{r_{j} \in \mathbb{Z}} \mathcal{E}_{x, y}\left(i+1, s+j l_{i+1}, u+j l_{i+1}+r_{j}\right)+\left|r_{j}\right| .
$$

The $\mathcal{E}$-distance from $x$ to $y$ is $\mathcal{E}_{x, y}(0,1,1)$.

Theorem 3.3 (Characterization): For every $b \geq 2$ and two strings $x, y \in \Sigma^{n}$, the $\mathcal{E}$-distance between $x$ and $y$ is $a 6 \cdot \frac{b}{\log b} \cdot \log n$ approximation to the edit distance between $x$ and $y$.
Definition (Alternative): Consider all the matching positions of $\mathcal{E}$ during the computation. Denote by $Z$ a vector of integers $z_{i, s}$ indexed by $i \in\{0,1, \ldots, h\}$ and $s \in B_{i}=\left\{1,1+l_{i}, \ldots\right\}$, where $z_{0,1}=1$ by convention. The coordinate $z_{i, s}$ should be understood as the position to which we match the substring $x\left[s: s+l_{i}\right]$. Then we define the cost of $Z$ as

$$
\operatorname{cost}(Z) \stackrel{\text { def }}{=} \sum_{i=0}^{h-1} \sum_{s \in B_{i}} \sum_{j=0}^{b-1}\left|z_{i, s}+j l_{i+1}-z_{i+1, s+j l_{i+1}}\right| .
$$

Claim 3.4 (Alternative definition of $\mathcal{E}$-distance): The $\mathcal{E}$-distance between $x$ and $y$ is the minimum of

$$
\operatorname{cost}(Z)+\sum_{s \in[n]} \mathrm{H}\left(x[s], y\left[z_{h, s}\right]\right)
$$

over all choices of the vector $Z=\left(z_{i, s}\right)_{i \in\{0,1, \ldots, h\}, s \in B_{i}}$ with $z_{0,1}=1$, where $\mathrm{H}(\cdot, \cdot)$ is the Hamming distance.

Lemma 3.5: The $\mathcal{E}$-distance between $x$ and $y$ is at most $3 h b \cdot \operatorname{ed}(x, y)$.
Lemma 3.6: The edit distance $\operatorname{ed}(x, y)$ is at most twice the $\mathcal{E}$-distance between $x$ and $y$.

## SAmpling Algorithm

Chernoff bound: Let $Z_{i} \in[0,1]$ be $n$ independent random variables from possibly different distributions. Let $Z=\sum_{i} Z_{i}$ and $\mu=\mathbb{E}[Z]$. Then for any $\varepsilon>0$ :

$$
\mathbb{P}[Z \geq(1+\varepsilon) \mu] \leq e^{-\frac{\varepsilon^{2} \mu}{2+\varepsilon}} \quad \text { and } \quad \mathbb{P}[Z \leq(1-\varepsilon) \mu] \leq e^{-\frac{\varepsilon^{2} \mu}{2}}
$$

Hoeffding bound: Let $Z_{i} \in[0,1]$ be $n$ independent random variables from possibly different distributions. Let $Z=\sum_{i} Z_{i}$ and $\mu=\mathbb{E}[Z]$. Then for any $t>0$, we have that

$$
\mathbb{P}[Z \geq t] \leq e^{-(t-2 \mu)}
$$

Definition 3.8: Fix $\rho>0$ and some $f \in[1,2]$. For a quantity $\tau \geq 0$, we call its $(\rho, f)$-approximator any quantity $\hat{\tau}$ such that $\tau / f-\rho \leq \hat{\tau} \leq f \tau+\rho$.

If $\hat{\tau}_{1}, \hat{\tau}_{2}$ are $(\rho, f)$-approximators to $\tau_{1}, \tau_{2}$ respectively, $\hat{\tau}_{1}+\hat{\tau}_{2}$ is a $(2 \rho, f)$-approximator to $\tau_{1}+\tau_{2}$.
If $\hat{\tau}^{\prime}$ is a $\left(\rho^{\prime}, f^{\prime}\right)$-approximator to $\hat{\tau}$, which itself is a $(\rho, f)$-approximator to $\tau$, then $\hat{\tau}^{\prime}$ is a ( $\rho^{\prime}+$ $f^{\prime} \rho, f f^{\prime}$ )-approximator to $\tau$.
Lemma 3.9 (Sum of random variables): Fix $n \in \mathbb{N}, \rho>0$ and error probability $\delta$. Let $Z_{i} \in[0, \rho]$ be independent random variables, and let $\zeta>0$ be a sufficiently large absolute constant. Then for every $\varepsilon \in[0,1]$, the summation $\sum_{i} Z_{i}$ is a $\left(\zeta \frac{\log 1 / \delta}{\varepsilon^{2}}, e^{\varepsilon}\right)$-approximator to $\mathbb{E}\left[\sum_{i} Z_{i}\right]$, with probability $\geq 1-\delta$.
Lemma 3.11 (Uniform Sampling): Fix $b \in \mathbb{N}, \varepsilon>0$, and error probability $\delta>0$. Consider some $a_{j}$, $j \in[b]$, such that $a_{j} \in[0,1 / b]$. For arbitrary $w \in[1, \infty)$, construct the set $J \subseteq[b]$ by subsampling each $j \in[b]$ with probability $p_{w}=\min \left(1, \frac{w}{b} \cdot \zeta \frac{\log 1 / \delta}{\varepsilon^{2}}\right)$. Then, with probability at least $1-\delta$, the value $\frac{1}{p_{w}} \sum_{j \in J} a_{j}$ is a $\left(1 / w, e^{\varepsilon}\right)$-approximator to $\sum_{j \in[b]} a_{j}$, and $|J| \leq \mathcal{O}\left(w \cdot \frac{\log 1 / \delta}{\varepsilon^{2}}\right)$.
Lemma 3.12 (Non-uniform Sampling): Fix integers $n \leq N$, approximation $\varepsilon>0$, factor $1<f<1.1$, error probability $\delta>0$, and an "additive error bound" $\rho>6 n / \varepsilon / N^{3}$. There exists a distribution $\mathcal{W}$ on the real interval $\left[1, N^{3}\right]$ with $\mathbb{E}_{w \in \mathcal{W}}[w] \leq \mathcal{O}\left(\frac{1}{\rho} \cdot \frac{\log 1 / \delta}{\varepsilon^{3}} \cdot \log N\right)$, as well as a "reconstruction algorithm" $R$, with the following property.

Take arbitrary $a_{i} \in[0,1]$, for $i \in[n]$, and let $\sigma=\sum_{i} a_{i}$. Suppose one draws $w_{i}$ i.i.d. from $\mathcal{W}$ and let $\hat{a}_{i}$ be an $\left(1 / w_{i}, f\right)$-approximator of $a_{i}$. Then, given $\hat{a}_{i}$ and $w_{i}$ for all $i \in[n]$, the algorithm $R$ generates a $\left(\rho, f \cdot e^{\varepsilon}\right)$-approximator to $\sigma$, with probability at least $1-\delta$.

