

DISTRIBUTION OF ALIGNED LETTER PAIRS IN OPTIMAL ALIGNMENTS OF RANDOM SEQUENCES

RAPHAEL HAUSER AND HEINRICH MATZINGER

ABSTRACT. Considering the optimal alignment of two i.i.d. random sequences of length n , we show that when the scoring function is chosen randomly, almost surely the empirical distribution of aligned letter pairs in all optimal alignments converges to a unique limiting distribution as n tends to infinity. This result is interesting because it helps understanding the microscopic path structure of a special type of last passage percolation problem with correlated weights, an area of long-standing open problems. Characterizing the microscopic path structure yields furthermore a robust alternative to optimal alignment scores for testing the relatedness of genetic sequences.

1. INTRODUCTION

1.1. Basic Definitions and Overview. Let $\mathcal{A}$ denote a finite alphabet, and let us consider alignments with gaps of two strings x and y of equal length n consisting of letters from $\mathcal{A}$. For each such alignment π we may count the number of letter pairs of different types aligned with each other and divide this number by n . Collecting these ratios for each possible pair of letters results in what we call the *empirical distribution vector* and denote by $\vec{p}_\pi(x, y)$. We remark that this results in the empirical probability distribution in the classical sense scaled by a factor $\tau \geq 1$ that is due to the presence of gaps. The classical distribution can of course be recovered by normalizing our notion of

1991 *Mathematics Subject Classification.* Primary 60K35; Secondary 52A40, 60C05, 60D05, 65K10, 60F10, 62E20, 90C27.

Key words and phrases. Sequence alignment, percolation theory, convex geometry, large deviations.

Raphael Hauser was supported by the Engineering and Physical Sciences Research Council [grant number EP/H02686X/1].

Heinrich Matzinger was supported by the Engineering and Physical Sciences Research Council [grant number EP/I01893X/1], IMA Grant SGS29/11, and by Pembroke College Oxford.

distribution. The set of all such empirical distribution vectors for given strings x and y will be denoted by

$$SET(x, y) := \{\vec{p}_\pi(x, y) : \pi \text{ is an alignment with gaps of } x \text{ and } y\}.$$

Let us give an example to illustrate the concepts we just introduced: Take $n = 4$, consider the strings $x = \mathbf{aabb}$ and $y = \mathbf{abab}$, where the alphabet consists of two letters $\mathcal{A} = \{\mathbf{a}, \mathbf{b}\}$, and let us look at a few alignments with gaps of x and y . First, let π be given by

$$\begin{array}{c|c|c|c|c} x & \mathbf{a} & \mathbf{a} & \mathbf{b} & \mathbf{b} \\ \hline y & \mathbf{a} & & \mathbf{b} & \mathbf{a} & \mathbf{b} \end{array}$$

There is one $\mathbf{a}$ aligned with $\mathbf{a}$ and hence the coefficient $p_{\mathbf{aa}} = 1/4$. Two pairs of letters $\mathbf{b}$ are aligned with each other, so $p_{\mathbf{bb}} = 2/4 = 0.5$. One letter $\mathbf{a}$ from x is aligned with a gap, so that $p_{\mathbf{a}\mathfrak{G}} = 1/4$. Here and elsewhere we use the symbol $\mathfrak{G}$ for a gap. And finally, one letter $\mathbf{a}$ from y is aligned with a gap, so that $p_{\mathfrak{G}\mathbf{a}} = 1/4$. The empirical distribution vector of the alignment π is now given by

$$\vec{p}_\pi(x, y) = (p_{\mathbf{aa}}, p_{\mathbf{ab}}, p_{\mathbf{a}\mathfrak{G}}, p_{\mathbf{ba}}, p_{\mathbf{bb}}, p_{\mathbf{b}\mathfrak{G}}, p_{\mathfrak{G}\mathbf{a}}, p_{\mathfrak{G}\mathbf{b}}) = (0.25, 0, 0.25, 0, 0.5, 0, 0.25, 0).$$

A second alignment ν is given by

$$\begin{array}{c|c|c|c|c} x & \mathbf{a} & \mathbf{a} & \mathbf{b} & \mathbf{b} \\ \hline y & \mathbf{a} & \mathbf{b} & \mathbf{a} & \mathbf{b} \end{array}$$

This time we find the following empirical distribution,

$$\vec{p}_\nu(x, y) = (p_{\mathbf{aa}}, p_{\mathbf{ab}}, p_{\mathbf{a}\mathfrak{G}}, p_{\mathbf{ba}}, p_{\mathbf{bb}}, p_{\mathbf{b}\mathfrak{G}}, p_{\mathfrak{G}\mathbf{a}}, p_{\mathfrak{G}\mathbf{b}}) = (0.25, 0.25, 0, 0.25, 0.25, 0, 0, 0).$$

Finally, consider the alignment μ given by

$$\begin{array}{c|c|c|c|c} x & \mathbf{a} & \mathbf{a} & \mathbf{b} & & \mathbf{b} \\ \hline y & \mathbf{a} & \mathbf{b} & & \mathbf{a} & \mathbf{b} \end{array},$$

The empirical distribution is given by

$$\vec{p}_\mu(x, y) = (p_{\mathbf{aa}}, p_{\mathbf{ab}}, p_{\mathbf{a}\mathfrak{G}}, p_{\mathbf{ba}}, p_{\mathbf{bb}}, p_{\mathbf{b}\mathfrak{G}}, p_{\mathfrak{G}\mathbf{a}}, p_{\mathfrak{G}\mathbf{b}}) = (0.25, 0.25, 0, 0, 0.25, 0.25, 0.25, 0).$$

Note that the coefficients of the empirical distribution vector do not usually add up to one. For example, in the case of alignment μ , we get

$$p_{\mathbf{aa}} + p_{\mathbf{ab}} + p_{\mathbf{a}\mathfrak{G}} + p_{\mathbf{ba}} + p_{\mathbf{bb}} + p_{\mathbf{b}\mathfrak{G}} + p_{\mathfrak{G}\mathbf{a}} + p_{\mathfrak{G}\mathbf{b}} = 1.25 > 1,$$

the reason being that we divided by the length n of the strings instead of the number of columns of the alignment. If the alignment were without gaps, the coefficients would add up to one and represent frequencies, but in general this is not the case. However, the coefficients of the empirical distribution

vector are proportional to the actual frequencies and hence may be thought of intuitively as representing these.

Let us next consider two random strings $X = X_1 \dots X_n$ and $Y = Y_1 \dots Y_n$, where X_i and Y_i are i.i.d. random variables taking values in the alphabet $\mathcal{A}$. Consider the set of all empirical distribution vectors that can be obtained from aligning X with Y and inserting the gaps in different places. We denote the convex hull of this set by

$$SET^n := \text{conv} (SET(X, Y)).$$

One of our main results, Theorem 2.1, will establish that

$$\lim_{n \rightarrow \infty} d(SET^n, SET) \stackrel{a.s.}{=} 0,$$

where SET is a unique limiting set that only depends on the distribution of the sequences X and Y , but not their realization, and where $d(\cdot, \cdot)$ denotes the Hausdorff distance between two subsets in $\mathbb{R}^n$ and is defined in (1.12) below.

Let $\mathcal{A}^*$ denote the alphabet $\mathcal{A}$ augmented by the symbol $\mathfrak{G}$, which stands for a gap, and consider functions S from $\mathcal{A}^* \times \mathcal{A}^*$ into the set of real numbers. Such functions will be called *scoring functions*. For a gapped alignment π of $x = x_1 \dots x_n$ and $y = y_1 \dots y_n$, let us define the score $S_\pi(x, y)$ under a scoring function S as the sum of scores of the aligned symbols pairs from $\mathcal{A}^*$. An alignment of x and y is called *optimal under S* if it maximizes $S_\pi(x, y)$ amongst all gapped alignments of x and y .

Another main result of this paper, Theorem 3.2, shows that when the scoring function S is chosen at random, the empirical distribution of any optimal alignment of the random strings $X = X_1 \dots X_n$ and $Y = Y_1 \dots Y_n$ with respect to S almost surely approaches a unique limiting vector $\vec{p}_S$ as n tends to infinity. Apart from the realization of S , this limit vector only depends on the *distribution* of X and Y , but not on their *realizations*.

Let us further illustrate the concept of *optimal alignment* of two strings by means of an example. Consider a scoring function S that takes the value 1 for identically aligned letters, and 0 otherwise. In

the context of the three examples of alignments introduced earlier, we find the following scores,

$$\begin{aligned} S_\pi(x, y) &= S(\mathfrak{a}, \mathfrak{a}) + S(\mathfrak{a}, \mathfrak{G}) + S(\mathfrak{b}, \mathfrak{b}) + S(\mathfrak{G}, \mathfrak{a}) + S(\mathfrak{b}, \mathfrak{b}) = 1 + 0 + 1 + 0 + 1 = 3, \\ S_\mu(x, y) &= S(\mathfrak{a}, \mathfrak{a}) + S(\mathfrak{a}, \mathfrak{b}) + S(\mathfrak{b}, \mathfrak{G}) + S(\mathfrak{G}, \mathfrak{a}) + S(\mathfrak{b}, \mathfrak{b}) = 1 + 0 + 0 + 0 + 1 = 2, \\ S_\nu(x, y) &= S(\mathfrak{a}, \mathfrak{a}) + S(\mathfrak{a}, \mathfrak{b}) + S(\mathfrak{b}, \mathfrak{a}) + S(\mathfrak{b}, \mathfrak{b}) = 1 + 0 + 0 + 1 = 2. \end{aligned}$$

One could check that the score of 3 is not exceeded for any alignment of x and y and hence find that π maximizes the alignment score of x and y under S . We thus say that π is an optimal alignment under S , although note that it may not be the unique alignment with this property.

The alignment score can also be viewed as the value that an appropriately defined linear functional takes on the empirical distribution vector of an alignment: Let $f_S : \mathbb{R}^8 \rightarrow \mathbb{R}$ be defined by

$$\begin{aligned} (1.1) \quad f_S(\vec{p}) &= f_S(p_{\mathfrak{a}\mathfrak{a}}, p_{\mathfrak{a}\mathfrak{b}}, p_{\mathfrak{a}\mathfrak{G}}, p_{\mathfrak{b}\mathfrak{a}}, p_{\mathfrak{b}\mathfrak{b}}, p_{\mathfrak{b}\mathfrak{G}}, p_{\mathfrak{G}\mathfrak{a}}, p_{\mathfrak{G}\mathfrak{b}}) \\ &= S(\mathfrak{a}, \mathfrak{a})p_{\mathfrak{a}\mathfrak{a}} + S(\mathfrak{a}, \mathfrak{b})p_{\mathfrak{a}\mathfrak{b}} + S(\mathfrak{a}, \mathfrak{G})p_{\mathfrak{a}\mathfrak{G}} + S(\mathfrak{b}, \mathfrak{a})p_{\mathfrak{b}\mathfrak{a}} + S(\mathfrak{b}, \mathfrak{b})p_{\mathfrak{b}\mathfrak{b}} \\ &\quad + S(\mathfrak{b}, \mathfrak{G})p_{\mathfrak{b}\mathfrak{G}} + S(\mathfrak{G}, \mathfrak{a})p_{\mathfrak{G}\mathfrak{a}} + S(\mathfrak{G}, \mathfrak{b})p_{\mathfrak{G}\mathfrak{b}}. \end{aligned}$$

It is then the case that

$$(1.2) \quad S_\pi(x, y) = n f_S(\vec{p}_\pi(x, y))$$

holds for any alignment π of x and y .

Recall that $SET(x, y)$ is defined as the set of all empirical distribution vectors of alignments of x and y ,

$$SET(x, y) := \{\vec{p}_\pi(x, y) : \pi \text{ is an alignment of } x \text{ and } y\}.$$

In particular, in our current example $SET(x, y) = SET(\mathfrak{a}\mathfrak{b}\mathfrak{a}\mathfrak{b}, \mathfrak{a}\mathfrak{a}\mathfrak{b}\mathfrak{b})$ contains the vectors

$$\begin{aligned} \vec{p}_\mu(x, y) &= (0.25, 0.25, 0, 0, 0.25, 0.25, 0.25, 0) \\ \vec{p}_\pi(x, y) &= (0.25, 0, 0.25, 0, 0.5, 0, 0.25, 0) \\ \vec{p}_\nu(x, y) &= (0.25, 0.25, 0, 0.25, 0.25, 0, 0, 0). \end{aligned}$$

Further, we write

$$(1.3) \quad L_S(x, y) := \max_{\pi} S_\pi(x, y)$$

for the optimal alignment score of x and y , where the maximum is taken over all gapped alignments π of x with y . Thus, in our current example we have $L_S(\mathbf{abab}, \mathbf{aabb}) = 3$. The rescaled maximum alignment score can also be seen as the maximum value taken by the functional f_S over $SET(x, y)$,

$$\frac{L_S(x, y)}{n} = \max_{\vec{p} \in SET(x, y)} f_S(\vec{p}).$$

1.2. Motivation.

1.2.1. *First passage Percolation.* The problem of understanding the structure of an optimal path in first and last passage percolation has been recognized as being important several decades ago but still remains largely unresolved, see Kesten [12]. Consider the set of edges

$$E := \{(z, w), (z, w + 1)\}, \{(z, w), (z + 1, w)\} : z, w \in \mathbb{Z}\}$$

of the integer lattice. E thus consists of vertical and horizontal edges of unit length incident to points in $\mathbb{R}^2$ with integer coordinates. Consider a setup in which a random weight $w(e)$ is associated with each edge $e \in E$. In the classical setting of First Passage Percolation, these random weights are i.i.d. distributed, and a path of smallest total weight between two points a and b is sought. Any admissible path must consist of consecutive adjacent edges $e_1, e_2, \dots, e_n \in E$, and e_1 and e_n must be incident to a and b respectively. The weights can also be interpreted as the time it takes to cross an edge, with the total weight $w(e_1) + w(e_2) + \dots + w(e_n)$ of the path corresponding to the passage time from a to b via the chosen path, and a minimum weight path corresponding to a fastest link between the two points.

An example of an open problem relating to the microstructure of an optimal path in first passage percolation is the following (see Kesten [12]): Consider the two points $a = (0, 0)$ and $b = (0, n)$. What is the proportion of vertical and horizontal edges in an shortest path from the point $(0, 0)$ to $(0, n)$, and does this proportion converge as n goes to infinity? We shall now argue that these questions are closely related to the central problem of this paper, which is to find the limiting empirical distribution of the aligned letter pairs. Consider the set of oriented edges

$$E' := \{((z, w), (z, w + 1)), ((z, w), (z + 1, w)), ((z, w), (z + 1, w + 1)) : z, w \in \mathbb{Z}\},$$

let $x = x_1 \dots x_n$ and $y = y_1 \dots y_n$ be strings of letters from the alphabet $\mathcal{A}$, let a scoring function $S : \mathcal{A}^* \times \mathcal{A}^* \rightarrow \mathbb{R}$ be given, and let the weight $w(e)$ of an edge of type $e = ((z, w), (z + 1, w + 1))$ (a diagonal edge) be equal to the score obtained by aligning the letter x_{z+1} with y_{w+1} ,

$$w(e) = S(x_{z+1}, y_{w+1}).$$

For edges of type $e = ((z, w), (z + 1, w))$ (horizontal edges), let the weight $w(e)$ be given by the score of aligning a gap with the letter x_{z+1} ,

$$w(e) = S(x_{z+1}, \mathfrak{G}).$$

Likewise, for vertical edges $e = ((z, w), (z, w + 1))$, let

$$w(e) = S(\mathfrak{G}, y_{w+1}).$$

In this manner, the problem of aligning x with y in an optimally according to the scoring function S becomes a Last Passage Percolation problem. The optimal alignment score $S(x, y)$ equals the weight of the maximum weight path going from $(0, 0)$ to (n, n) . An optimal path $e_1 e_2 \dots e_m$, that is, a path of maximum total weight among those that follow oriented edges from E' and link $(0, 0)$ to (n, n) , defines an optimal alignment of x with y in the following fashion: For any diagonal edge $((z, w), (z + 1, w + 1))$ that lies along the path, align the letter x_{z+1} with y_{w+1} . align all other letters with gaps. Now note that when we know the limit of the empirical distribution vector of the aligned letter pairs, we also know the proportion of gaps on the long run. In other words, the limiting distribution of the aligned letter pairs yields the asymptotic proportion of horizontal and vertical edges in the optimal path in E' . This information is the equivalent to knowing the asymptotic proportion of vertical and horizontal edges in last passage percolation in a model where the distribution of edge weights is different, and where diagonal edges are present. The corresponding results for first passage percolation can be obtained by multiplying the edge weights by -1 .

1.2.2. Computational Genomics. In computational genomics the alignment score is a maximum likelihood ratio to decide which alignment is the most likely association of sequences that diverged by evolution. Any gapped alignment of two DNA or RNA strings $x = x_1 \dots x_n$ and $y = y_1 \dots y_n$ represents a hypothesis about the evolutionary history of the two species relative to one another. Assume that the strings x and y are sections of

genetic sequences from two extant species with a common ancestor. Aligning x_i with y_j corresponds to the hypothesis that both letters descended from a particular letter in the corresponding section in the ancestor's genome. A letter aligned with a gap may correspond to a letter in the ancestral genome that disappeared in one of the two descendants or a new letter that appeared through mutation. The value $S(\mathbf{a}, \mathbf{b})$ of the scoring function at $\mathbf{a}, \mathbf{b} \in \mathcal{A}^*$ is equal to the logarithm of the probability that a letter from the ancestral genome evolved into a letter $\mathbf{a}$ in one of the extant species and into a letter $\mathbf{b}$ in the other, assuming that letters mutate independently of their neighbors¹. Naturally, the scoring function depends on how long ago the two species got separated in the evolutionary tree. Given more time since separation, the probability of mutation increases, and as a result the scoring function also will look different.

In practice it can sometimes be difficult to determine whether two given sequences are related or not, since a good choice of the scoring function S may not be known a priori in absence of a good estimate of the time since evolutionary divergence between the two extant species. Sections of DNA-sequences might also look similar because they have a similar distribution of amino acids rather than being related by direct evolution from a particular section of the ancestral genome. An approach to overcoming this problem is to relate sequences by the microscopic path structure of optimal alignments. Hirao, Lember and Matzinger [10] found that optimal alignments of related and non-related sequences have entirely different microscopic structures. Preliminary experiments showed that basing a relatedness test on path structure works at least as well as the widely used test based on the BLASTZ algorithm, which works with a sophisticated scoring function. The results of our paper will help to take this work further.

1.3. Monte Carlo Simulation. An appealing approach to the investigation of asymptotic qualities of optimal alignments of random sequences is to use Monte Carlo simulation. However, to derive rigorous bounds one usually needs to know the variance $\text{VAR}(L_n(S))$ of the alignment score as a function of the length n of the aligned random sequences $X_1 \dots X_n$ and $Y_1 \dots Y_n$, and this order of fluctuation is not yet well understood. In the special case of the longest common subsequence problem with sequences consisting of i.i.d.

¹This is of course an inexact approximation of true mutation dynamics.

Ber(p) variables, Chvatal and Sankoff [6] conjectured that $\text{VAR}(L_n(S)) = o(n^{2/3})$ when $p = 0.5$. Steele [19] later proved the bound $\text{VAR}(L_n(S)) \leq 2p(1-p)n$. Waterman [21] asked the question of whether this bound can be improved and found that for $p < 0.5$, simulations suggest that the dependence of $\text{VAR}(L_n(S))$ on n is linear. Boutet de Monvel [4] found that this also applies to the case $p = 0.5$, although the linear growth only sets in for very large n . Lember and Matzinger [14] gave a rigorous proof of the linear order $\text{VAR}(L_n(S)) = \Theta(n)$ in the case where p is very small. Their analysis was based on showing that the manipulation of randomly selecting a letter of specified type from one of the two sequences and changing it into another specified type has a positive biased effect on the optimal alignment score. In Section 4, we significantly extend the applicability of this result to general scoring functions and random sequences whose distributions are not highly asymmetric: Theorem 4.2 yields a sufficient criterion under which the asymptotic order of fluctuation $\text{VAR}(L_n(S)) = \Theta(n)$ holds.

1.4. Summary of Main Results. We now amend the notation introduced in (1.3) and write

$$(1.4) \quad L_n(S) = \max_{\pi} S_{\pi}(X, Y)$$

for the optimal alignment score of the two i.i.d. strings $X = X_1, X_2, \dots, X_n$ and $Y = Y_1, Y_2, \dots, Y_n$ with letters from the alphabet $\mathcal{A}$. Since X and Y are random, the maximum score $L_n(S)$ is a random variable, and we are interested in its dependence on the scoring function $S : \mathcal{A}^* \times \mathcal{A}^* \rightarrow \mathbb{R}$ in the asymptotic regime where the length n of the strings tends to infinity. Let $\lambda_n(S)$ denote the rescaled expected alignment score,

$$(1.5) \quad \lambda_n(S) := \frac{E[L_n(S)]}{n}.$$

A simple subadditivity argument, see Chvátal & Sankoff [6], shows that

$$(1.6) \quad \lambda_n(S) \leq \lambda_m(S), \quad \forall n \leq m \in \mathbb{N},$$

$$(1.7) \quad \lambda(S) := \lim_{n \rightarrow \infty} \frac{E[L_n(S)]}{n} \text{ exists,}$$

$$(1.8) \quad \mathbb{P} \left[\lim_{n \rightarrow \infty} \frac{L_n(S)}{n} = \lambda(S) \right] = 1.$$

Furthermore, Alexander [1] showed that in the case of the longest common subsequence problem, the convergence is slower than the order $\sqrt{\ln n}/\sqrt{n}$,

$$\lambda(S) - \lambda_n(S) \geq C \frac{\sqrt{\ln n}}{\sqrt{n}}.$$

In Lemma 6.2 we show that a lower bound of the same order also exists. The exact value of the Chvátal-Sankoff constant $\lambda(S)$ is unknown even in the simplest cases, and Montecarlo simulations to obtain estimates of even moderate accuracy are quite involved [8, 9]. Note that $\lambda(S)$ depends of course on the distribution of the random strings.

Let $H(S)$ designate the half space

$$(1.9) \quad H(S) := \left\{ \vec{x} \in \mathbb{R}^{|\mathcal{A}^*|^2} : f_S(\vec{x}) \leq \lambda(S) \right\},$$

where f_S is defined by

$$f_S : \mathbb{R}^{|\mathcal{A}^*|^2} \rightarrow \mathbb{R},$$

$$\vec{x} \mapsto \sum_{\mathbf{a} \in \mathcal{A}^*} \sum_{\mathbf{b} \in \mathcal{A}^*} S(\mathbf{a}, \mathbf{b}) x_{\mathbf{a}\mathbf{b}}.$$

We also consider

$$(1.10) \quad SET := \bigcap_S H(S),$$

where the intersection is taken over all scoring functions S . It is immediate from its definition that SET is a closed convex set, and in Lemma 5.2 we will furthermore show that it is compact with nonempty interior. Recall also the notation

$$SET(X, Y) := \{ \vec{p}_\pi(X, Y) : \pi \text{ is an alignment with gaps of } X \text{ and } Y \} \subset \mathbb{R}^{|\mathcal{A}^*|^2}$$

introduced earlier for the set of empirical distribution vectors for gapped alignments of X and Y . Since X and Y are random strings, $SET(X, Y)$ is a random set. We denote the convex hull of this set by

$$(1.11) \quad SET^n := \text{conv}(SET(X, Y)),$$

where we account for the length n of the random strings X and Y notationally because we are interested in the asymptotic behavior when n tends to infinity.

The Hausdorff distance [7] between two sets $A, B \subset \mathbb{R}^n$ is defined as follows, where $\|\cdot\|$ denotes the Euclidean norm,

$$(1.12) \quad \begin{aligned} d(A, B) &= \max \left(\sup_{x \in A} d(x, B), \sup_{y \in B} d(A, y) \right) \\ d(x, B) &= \inf \{ \|x - y\| : y \in B \}, \\ d(A, y) &= \inf \{ \|x - y\| : x \in A \}. \end{aligned}$$

Let us now discuss the main results of this paper:

- (1) Theorem 2.1 will establish that the random set SET^n almost surely converges to the deterministic set SET in terms of Hausdorff distance. As a consequence, if one were to simulate the sequences $X_1 \dots X_n$ and $Y_1 \dots Y_n$ and compute the convex hull of the empirical distribution vectors of all their alignments with gaps, one would find a set that closely resembles SET , provided n is large enough.
- (2) Theorem 3.1 will show that the empirical distributions of all optimal alignments of X and Y almost surely converge to a deterministic distribution as n tends to infinity, on condition that the scoring function S be chosen such that f_S has a unique maximizer in SET . When this condition is met, we denote the unique maximiser by $\vec{p}_S$. The statement of the theorem then says that the probability that there exists an optimal alignment of X and Y with respect to S with empirical distribution further away than $\epsilon > 0$ from $\vec{p}_S$ is negatively exponentially small in n , where ϵ is an arbitrary small constant independent of n .
- (3) The condition of Theorem 3.1 is difficult to verify in practice, but Theorem 5.1 shows that when the scoring function S is chosen randomly, then the condition is almost surely met, that is, f_S has a unique maximizer in SET with probability 1. As a corollary, we obtain Theorem 3.2, which says that for almost all scoring function S the empirical distributions of all optimal alignments of X and Y almost surely converge to a deterministic distribution.
- (4) Theorem 3.2 allows for the derivation of a sufficient criterion to guarantee that the order of fluctuation of the optimal alignment score is linear in the length n of the aligned random strings $X_1 \dots X_n$ and $Y_1 \dots Y_n$. The sufficient criterion depends on the scoring function S and on the unique maximizer $\vec{p}_S$ of f_S over

SET and constitutes a practical tool in the design of a statistical test on the order of fluctuation of the optimal score, see [2].

1.5. A Few Key Ideas. Let $S : \mathcal{A}^* \times \mathcal{A}^* \rightarrow \mathbb{R}$ be a scoring function, x and y two strings of length n with letters from $\mathcal{A}$, and π a gapped alignment of x and y . By (1.2), the optimal alignment score of x and y with respect to S satisfies

$$\max_{\pi} \frac{S_{\pi}(x, y)}{n} = \max_{\vec{p} \in SET(x, y)} f_S(\vec{p}).$$

Applied to the random strings $X = X_1 \dots X_n$ and $Y = Y_1 \dots Y_n$, we find

$$(1.13) \quad \frac{L_n(S)}{n} = \max_{\vec{p} \in SET(X, Y)} f_S(\vec{p}).$$

A crucial observation is now that the maximum of a linear function over a closed set in $\mathbb{R}^n$ equals the maximum of the given function over the convex hull of the given set. Since $SET^n = \text{conv}(SET(X, Y))$, (1.13) implies

$$(1.14) \quad \frac{L_n(S)}{n} = \max_{\vec{p} \in SET^n} f_S(\vec{p}).$$

By (1.8), $L_n(S)/n$ almost surely converges to a deterministic constant $\lambda(S)$ which was also used to define SET , see (1.9) and (1.10). The definition of SET shows further that

$$(1.15) \quad \max_{\vec{p} \in SET} f_S(\vec{p}) \leq \lambda(S),$$

and Lemma 5.2 d), proven below, shows that the inequality in (1.15) holds in fact at equality. Combined with (1.14), this implies

$$(1.16) \quad \max_{\vec{p} \in SET^n} f_S(\vec{p}) \xrightarrow{n \rightarrow \infty} \max_{\vec{p} \in SET} f_S(\vec{p}) \quad \text{almost surely.}$$

At a first pass it is illustrative to consider an approximate proof of Theorem 3.2 that is free of technical details relating to large deviations that will be necessary to render the proof rigorous. Proposition 5.2 and Theorems 5.4 and 5.1, which will be proven in Section 5, provide the crucial insight: by (1.16), the conditions of Theorem 5.4 are approximately met for $C = SET$ and $C^n = SET^n$, $n \in \mathbb{N}$, and hence, it is plausible to argue that $SET^n \rightarrow SET$. We remark that in rendering this argument rigorous later, we will use the fact that $L_n(S)/n$ converges to $\lambda(S)$ at a rate on the order of $\ln(n)/\sqrt{n}$, which follows directly from the Azuma-Hoeffding Inequality, as we shall see in Section 6.

The convergence of SET^n to SET occurs thus at the same rate. Choosing the scoring function S at random is tantamount to choosing f_S randomly, whence Theorem 5.1 shows that the conditions of Proposition 5.2 are satisfied. Theorem 3.2 thus follows.

1.6. Some Key Difficulties. Consider the optimal alignment score $L_n(S)$ for some scoring function S . $L_n(S)$ is a random variable, because it depends on the realization of the i.i.d. random strings $X = X_1 \dots X_n$ and $Y = Y_1 \dots Y_n$. However, it is easy to see that changing the realization of only one of these variables results in a change of $L_n(S)$ by at most the deterministic constant

$$\max_{c,d,e \in \mathcal{A}^*} |S(c,d) - S(c,e)|.$$

See Lemma 6.1 for details. One can therefore apply the Azuma-Hoeffding Inequality to find that, on a scale of $\sqrt{n}$, the tail of $L_n(S)$ decays at least quadratically exponentially fast, see Lemma 6.3. This powerful tool lends itself to an elegant analysis of the asymptotic convergence of the alignment score and its fluctuation.

In contrast, analyzing the convergence of the empirical distribution of letter pairs in optimal alignments is much harder: upon changing the realization of one of the random letters, it has to be assumed a priori that the entire optimal alignment and hence the relative frequencies at which alignments occur have changed. As a consequence, the Azuma-Hoeffding Inequality cannot be applied directly. Luckily, it can be applied indirectly through the optimal alignment scores of different scoring functions at the cost of having to deal with additional technicalities.

A further key difficulty is that for the scoring functions S under consideration, it is required that f_S be maximized in only one point on SET . This condition would be met if SET was known to be strictly convex everywhere, but this seems very difficult to verify in practice, since the exact shape of SET is unknown: SET corresponds to the asymptotic shape of the wet zone in the first/last passage percolation formulation of our problem, and determining the shape of the corresponding zone in standard first passage percolation is a long-standing open problem in the general case. We get around this problem by showing that if the scoring function S is chosen at random, then with probability one there exists a unique maximizer of f_S on SET , see Theorem 5.1.

We remark that convergence of the empirical distribution of aligned letter pairs may not hold when S is not chosen randomly. For a counter-example, construct a scoring function that takes the value 1 for pairs of identical letters, 0 otherwise, and take $\mathcal{A} = \{0, 1\}$ and X_i, Y_j i.i.d. Bernoulli $\text{Ber}(1/2)$ variables. The optimal alignment score then corresponds to the length of the longest common subsequence (LCS), and the asymptotic empirical distribution of optimally aligned letters is not unique: Write out the optimally aligned sequences with the introduced gaps and subdivide them into sections of length 3, e.g.,

$$\begin{array}{c|c|c|c|c} 1 & 0 & 0 & 0 & 1 & 0 & 1 & 0 & 0 & 0 & 0 & \dots \\ \hline 1 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 1 & 0 & 0 & \dots \end{array}$$

One observes empirically that a positive proportion of triplets is of the form

$$\begin{array}{c|c|c} 0 & 1 & 0 \\ \hline 0 & 1 & 0 \end{array}, \begin{array}{c|c|c} 0 & 1 & 0 \\ \hline 0 & 1 & 0 \end{array}, \begin{array}{c|c|c} 0 & 1 & 0 \\ \hline 0 & 1 & 0 \end{array} \text{ or } \begin{array}{c|c|c} 1 & 0 & 0 \\ \hline 1 & 0 & 0 \end{array}.$$

The first two correspond to the pattern 01 in X being aligned with the pattern 10 in Y , and the last two to the inverted situation. Thus, the first triplet can be exchanged for the second, and the third for the fourth without affecting $L_n(S)$ nor the optimality of the alignment. The empirical distribution of optimally aligned letter pairs changes however, as weight is shifted from the pairing $(1, 1)$ to the pairing $(0, 0)$.

2. SET CONVERGENCE OF EMPIRICAL DISTRIBUTIONS

Theorem 2.1. *Let SET and SET^n be the sets defined in (1.10) and (1.11), and let d denote the Hausdorff distance defined in (1.12). Then*

$$(2.1) \quad \mathbb{P} \left[d(SET^n, SET) \xrightarrow{n \rightarrow \infty} 0 \right] = 1.$$

Proof. By the definition of d in (1.12), we need to prove the two identities

$$(2.2) \quad \mathbb{P} \left[\max_{\vec{x} \in SET^n} d(\vec{x}, SET) \xrightarrow{n \rightarrow \infty} 0 \right] = 1,$$

$$(2.3) \quad \mathbb{P} \left[\max_{\vec{x} \in SET} d(\vec{x}, SET^n) \xrightarrow{n \rightarrow \infty} 0 \right] = 1.$$

To prove Equation (2.2), we use Lemma 5.1 which establishes that, given $\epsilon > 0$, there exist finitely many scoring functions $S_1, \dots, S_k$ such that

$$\max_{\vec{x} \in \bigcap_{i=1}^k H(S_i)} d(\vec{x}, SET) \leq \epsilon,$$

where the half-spaces

$$H(S_i) := \left\{ \vec{x} \in \mathbb{R}^{|\mathcal{A}^*|^2} : f_{S_i}(\vec{x}) \leq \lambda(S_i) \right\},$$

are defined as in Equation (1.9). Let $H_n^+(S_i)$ denote the shifted half-space

$$H_n^+(S_i) = \left\{ \vec{x} : f_{S_i}(\vec{x}) \leq \lambda_n(S_i) + \frac{\ln(n)}{\sqrt{n}} \right\},$$

and let us define the event

$$\mathcal{A}_n(S_i) = \left\{ \omega \in \Omega : SET^n(\omega) \subset H_n^+(S_i) \right\},$$

where Ω is the probability space over which the random sequences X and Y are defined. Corollary 6.1 and its proof show that the Azuma-Hoeffding Inequality implies

$$\mathbb{P}[\mathcal{A}_n(S_i)^c] \leq n^{-c_{S_i} \ln n},$$

where $c_{S_i} > 0$ is a constant that does not depend on n , see Theorem 6.1. It follows that

$$\mathbb{P} \left[SET^n \subset \bigcap_{i=1}^k H_n^+(S_i) \right] \geq 1 - \sum_{i=1}^k n^{-c_{S_i} \ln n} \geq 1 - n^{-c \ln n},$$

where $c > 0$ is a constant independent of n . The series $\sum_n n^{-c \ln n}$ being convergent, the Borel-Cantelli Lemma implies that almost surely there exists $n_0 \in \mathbb{N}$ such that

$$SET^n \subset \bigcap_{i=1}^k H_n^+(S_i), \quad \forall n \geq n_0.$$

By the definition of the S_i , this implies that for all $n \geq n_0$, we have

$$(2.4) \quad \max_{\vec{x} \in SET^n} d(\vec{x}, SET) \leq C \times \frac{\ln n}{\sqrt{n}} + \epsilon$$

where $C > 0$ is a constant independent of n . This implies that

$$\mathbb{P} \left[\limsup_{n \rightarrow \infty} \max_{\vec{x} \in SET^n} d(\vec{x}, SET) \leq \epsilon \right] = 1, \quad \forall \epsilon > 0.$$

Finally, since this is true for all ϵ rational, Equation (2.2) follows.

To prove Equation (2.3), we employ Theorem 5.3 that establishes that for any given $\epsilon > 0$, there exist points $\vec{x}_1, \dots, \vec{x}_k \in SET_{SE}$, a certain subset² of the set of extreme points of SET , and chosen such that

$$d(\vec{x}, \text{conv}(\vec{x}_1, \dots, \vec{x}_k)) \leq \epsilon, \quad \forall \vec{x} \in SET.$$

²See Section 5 for the relevant theory.

We now claim that for each $\vec{x}_i$ there almost surely exists a sequence of points $\vec{x}_{1,i}, \vec{x}_{2,i}, \vec{x}_{3,i}, \dots$ such that $\vec{x}_{j,i} \in SET^j$ for all $j \in \mathbb{N}$ and

$$\limsup_{j \rightarrow \infty} \|\vec{x}_{j,i} - \vec{x}_i\| \leq \epsilon.$$

By the triangular inequality, our claim implies that almost surely it is the case that

$$\limsup_{n \rightarrow \infty} \max_{\vec{x} \in SET} d(\vec{x}, SET^n) \leq 2\epsilon, \quad \forall \epsilon > 0,$$

and since this is true for all ϵ rational, Equation (2.3) follows.

It remains to prove our claim. Proposition 5.1 shows that there exist scoring functions $S_{\vec{x}_i}$ and $S_1, \dots, S_\ell$ such that

$$\vec{x}_i \in C_i := \left\{ \vec{x} : f_{S_{\vec{x}_i}}(\vec{x}) \geq \lambda(S_{\vec{x}_i}), f_{S_j}(\vec{x}) \leq \lambda(S_j), (j = 1, \dots, \ell) \right\} \subset B_\epsilon(\vec{x}_i)$$

and $C_{n,i}$ compact for all $n \in \mathbb{N}$, where

$$C_{n,i} := \left\{ \vec{x} : f_{S_{\vec{x}_i}}(\vec{x}) \geq \lambda_n(S_{\vec{x}_i}) - \frac{\ln n}{\sqrt{n}}, f_{S_j}(\vec{x}) \leq \lambda(S_j) + \frac{\ln n}{\sqrt{n}}, (j = 1, \dots, \ell) \right\}.$$

Further, by (1.6), the sets $C_{n,i}$ are nested, and by compactness and (1.7), we have

$$(2.5) \quad \limsup_{n \rightarrow \infty} d(\vec{x}_i, C_{n,i}) \leq \epsilon.$$

We will now show that with high probability $C_{n,i}$ has a nonempty intersection with SET^n . Consider the events

$$\begin{aligned} \mathcal{B}_{n,j} &:= \left\{ \omega \in \Omega : \frac{L_n(S_j)}{n} \leq \lambda(S_j) + \frac{\ln n}{\sqrt{n}} \right\}, \\ \mathcal{C}_{n,i} &:= \left\{ \omega \in \Omega : \exists \vec{x} \in SET^n \text{ s.t. } f_{S_{\vec{x}_i}}(\vec{x}) \geq \lambda_n(S_{\vec{x}_i}) - \frac{\ln n}{\sqrt{n}} \right\}. \end{aligned}$$

By Theorem 6.1, we have

$$\mathbb{P}[\mathcal{B}_{n,j}^c] \leq n^{-K_j \ln n}$$

where $K_j > 0$ is a constant that does not depend on n . Note also that Equation (1.14) implies

$$\mathcal{B}_{n,j} = \left\{ \omega \in \Omega : f_{S_j}(\vec{x}) \leq \lambda(S_j) + \frac{\ln n}{\sqrt{n}}, \forall \vec{x} \in SET^n \right\}.$$

In conjunction with (1.6), Theorem 6.1 further implies

$$\mathbb{P} [\mathcal{C}_{n,i}^c] \leq n^{-\ln(n)}.$$

But note that when the events $\mathcal{C}_{n,i}$ and $\mathcal{B}_{n,1}, \dots, \mathcal{B}_{n,\ell}$ occur jointly, then $SET^n \cap C_{n,i} \neq \emptyset$ holds. The probability that the intersection is empty is thus bounded from above by

$$\mathbb{P} [\mathcal{C}_{n,i}^c] + \sum_{i=1}^{\ell} \mathbb{P} [\mathcal{B}_{n,i}^c] \leq (\ell + 1)n^{-K \ln n},$$

where $K > 0$ is a constant that does not depend on n .

In view of the fact that the series

$$\sum_{n=1}^{\infty} (\ell + 1)n^{-K \ln n}$$

converges, the Borel-Cantelli Lemma now implies that, almost surely, for all but a finite number of $n \in \mathbb{N}$ there exists $x_{n,i} \in SET^n \cap C_{n,i}$. In the finitely many cases where $SET^n \cap C_{n,i} = \emptyset$ we can pick an arbitrary point $x_{n,i} \in SET^n$ to complete the sequence. In view of (2.5), we thus find that almost surely it is possible to construct a sequence $(x_{n,i})_{n \in \mathbb{N}}$ with the claimed properties. Hence, this settles the theorem. $\square$

3. POINT CONVERGENCE

So far we established that the empirical distributions of optimal alignments of random sequences under any scoring function asymptotically lie in SET . We will now show that for a fixed, randomly chosen scoring function S , the empirical distributions of all optimal alignments of X and Y under S converge to a unique point in SET . Recall the notation $\vec{p}_\pi(x, y)$ introduced in Section 1.1, and let us write

$$SET^*(X, Y) = \{\vec{p}_\pi(X, Y) : \pi \text{ is an optimal alignment of } X \text{ and } Y\}$$

for the set of empirical distributions corresponding to optimal alignments of $X = X_1 \dots X_n$ and $Y = Y_1 \dots Y_n$. Consider the event

$$\mathcal{D}_n(\vec{p}, \epsilon) := \{\omega \in \Omega : SET^*(X(\omega), Y(\omega)) \setminus \overline{B}_\epsilon(\vec{p}) \neq \emptyset\}$$

that there exists an optimal alignment π of $x = X(\omega)$ and $y = Y(\omega)$ under the scoring function S such that $\|\vec{p}_\pi(x, y) - \vec{p}\| > \epsilon$.

Theorem 3.1. *Let S be a scoring function such that the hyperplane*

$$(3.1) \quad \{\vec{x} : f_S(\vec{x}) = \lambda(S)\}$$

intersects SET in a unique point $\vec{p}_S$, and let $\epsilon > 0$ be given. Then there exists a constant K_ϵ such that for all $n \in \mathbb{N}$ it is true that

$$\mathbb{P}[\mathcal{D}_n(\vec{p}_S, \epsilon)] \leq e^{-K_\epsilon n}.$$

Furthermore, $SET^(X, Y) \rightarrow \{\vec{p}_S\}$ almost surely as n tends to infinity.*

Proof. By Lemma 5.2, SET is a compact convex set with nonempty intersection with the hyperplane (3.1), and by (1.15) all such intersection points are maximizers of the optimization problem $\max_{\vec{y} \in SET} \langle \vec{s}, \vec{y} \rangle$, where $\vec{s}$ is the normalization of the vector representation of the linear functional f_S defined by the scoring function. It follows that $\vec{p}_S$ satisfies Definition 5.1 of a point of strict curvature of SET . Proposition 5.1 therefore implies that there exist finitely many scoring functions $S_1, S_2, \dots, S_k$ and thresholds $\epsilon_0, \dots, \epsilon_k > 0$ such that

$$(3.2) \quad \{\vec{x} : f_S(\vec{x}) \geq \lambda(S) - \epsilon_0\} \cap \bigcap_{i=1}^k \{\vec{x} : f_{S_i}(\vec{x}) \leq \lambda(S_i) + \epsilon_i\} \subset B_\epsilon(\vec{p}_S),$$

Consider now the events

$$\mathcal{E}_{n,i} := \{\omega \in \Omega : SET^n \subset \{\vec{x} : f_{S_i}(\vec{x}) \leq \lambda(S_i) + \epsilon_i\}\}.$$

By (1.14) this is equivalent to requiring that the rescaled optimal alignment score $L_n(S_i)/n$ satisfy $L_n(S_i)/n \leq \lambda(S_i) + \epsilon_i$. By Theorem 6.1 there exists $K_i > 0$ such that

$$(3.3) \quad \mathbb{P}[\mathcal{E}_{n,i}] \geq 1 - e^{-K_i n} \quad \forall n.$$

Let us further define the event

$$\mathcal{E}_{n,0} := \{\omega \in \Omega : SET^n \cap \{\vec{x} : f_S(\vec{x}) \geq \lambda(S) - \epsilon_0\} \neq \emptyset\},$$

which is the same as requiring that $L_n(S)/n$ exceed the value $\lambda(S) - \epsilon_0$. Corollary 6.1 once again shows that there exists $K_0 > 0$ such that

$$(3.4) \quad \mathbb{P}[\mathcal{E}_{n,0}] \geq 1 - e^{-K_0 n} \quad \forall n.$$

Combining all of the above, we now find $\mathcal{D}_n^c \subseteq \bigcup_{i=0}^k \mathcal{E}_{n,i}$, so that

$$\mathbb{P}[\mathcal{D}_n] \leq \sum_{i=0}^k \mathbb{P}[\mathcal{E}_{n,i}^c] \leq \sum_{i=0}^k e^{-K_i n} \leq e^{-K_\epsilon n}$$

for some constant $K_\epsilon > 0$, as claimed.

The last statement follows from the Borel-Cantelli Lemma in a similar construction as in the proof of Theorem 2.1. $\square$

The above theorem shows that if $\vec{p}_S$ is the only solution to $f_S(\vec{p}) = \lambda_S$, then, denoting any optimal alignment of $X_1 \dots X_n$ and $Y_1 \dots Y_n$ with respect to S by π_n , it is true a.s. that $\vec{p}_{\pi_n}(X_1 \dots X_n, Y_1 \dots Y_n) \rightarrow \vec{p}_S$. Note however that the convergence rate was not specified. Our convergence argument, which is based on the Azuma-Hoeffding Inequality – see Theorem 6.1 – could be made quantitative if a bound on the curvature of SET at $\vec{p}_S$ were known.

Our second and main result of this section shows that the above theorem generically applies. For this purpose we consider a scoring function S that is chosen randomly in such a way that if $\vec{S}$ denotes the normalization of the vector representation of the linear functional f_S , then $\vec{S}$ has an absolutely continuous distribution with respect to the Hausdorff measure (or uniform measure) on the sphere. In this case we say that S has absolutely continuous distribution.

Theorem 3.2. *Let the scoring function S be chosen randomly from an absolutely continuous distribution, and let π_n denote any optimal alignment of $X_1 \dots X_n$ with $Y_1 \dots Y_n$. Then almost surely $\vec{p}_{\pi_n}(X, Y)$ converges to a unique empirical distribution.*

Proof. By Theorem 5.1, the conditions of Theorem 3.1 apply with probability 1. $\square$

4. FLUCTUATION OF THE OPTIMAL ALIGNMENT SCORE

Let $X = X_1 \dots X_n$ and $Y = Y_1 \dots Y_n$ be the random strings introduced earlier, let $\mathfrak{a}$ and $\mathfrak{b}$ be two distinct letters from the alphabet $\mathcal{A}$, and let us define a new random string $\tilde{X} = \tilde{X}_1 \dots \tilde{X}_n$ via the following compound procedure:

- (1) sample a realization $x = x_1 \dots x_n$ of X ,
- (2) if $\mathcal{J} := \{i : x_i = \mathfrak{a}\} \neq \emptyset$,
 - (a) let J be a random index defined on some probability space $(\tilde{\Omega}, \tilde{\mathbb{P}})$ and taking values with uniform distribution on $\mathcal{J}$,
 - (b) select a sample $j = J(\tilde{\omega})$,
 - (c) set $\tilde{x}_j = \mathfrak{b}$ and $\tilde{x}_i = x_i$ for $i \neq j$,
- (3) else set $\tilde{x} = x$.

Note that the distribution of $\tilde{X}$ generally differs from the distribution of X , and that, while $X = X_1 \dots X_n$ consists of the first n letters of a random sequence $(X_i)_{i \in \mathbb{N}}$, the same cannot be said about $\tilde{X}$: we only ever sample (at most) one entry of X realized in the form of an $\mathfrak{a}$, independently of n , so that the probability of any given index to be chosen diminishes as n grows:

The following result was proven by Lember and Matzinger [14], where we use the notation

$$\tilde{L}_n(S) := \max_{\pi} S_{\pi}(\tilde{X}, Y),$$

in analogy to the earlier introduced random variable $L_n(S) = \max_{\pi} S_{\pi}(X, Y)$, and where we write $f(n) = \Theta(n)$ if there exist constants $0 < c_1 < c_2$ such that $c_1 n \leq f(n) \leq c_2 n$ for all $n \in \mathbb{N}$.

Theorem 4.1. *Let the scoring function S and the distribution of X and Y be chosen so that there exist parameters $\beta, \varepsilon > 0$ for which*

$$\mathbb{P} \left[\mathbb{E}_{\tilde{\mathbb{P}}} \left[\tilde{L}_n(S) - L_n(S) \mid X, Y \right] \geq \varepsilon \right] \geq 1 - e^{-\beta n}, \quad \forall n \in \mathbb{N}.$$

Then the order of fluctuation of the optimal alignment score is given by

$$\text{VAR} [L_n(S)] = \Theta(n).$$

Up until now, the criterion of Theorem 4.1 could only be verified in a few special cases. We will next see that when the scoring function satisfies the conditions of Theorem 3.1, then the criterion of Theorem 4.1 can be reduced to a condition that solely depends on $\vec{p}_S$ and that can be verified by Montecarlo simulation to high confidence:

Theorem 4.2. *Let S be such that $\{\vec{x} : f_S(\vec{x}) = \lambda(S)\}$ intersects S in a unique point $\vec{p}_S = (p_{\mathfrak{c}\mathfrak{d}})$ and such that there exist $\mathfrak{a}, \mathfrak{b} \in \mathcal{A}$ for which it is the case that*

$$\sum_{\mathfrak{c} \in \mathcal{A}^*} p_{\mathfrak{ac}} (S_{\mathfrak{bc}} - S_{\mathfrak{ac}}) > 0.$$

Then the order of fluctuation of the optimal alignment score is given by

$$\text{VAR}[L_n(S)] = \Theta(n).$$

Proof. Let $\mathcal{J} = \{i \in \{1, \dots, n\} : X_i = \mathfrak{a}\}$, $q_{\mathfrak{a}} = \mathbb{P}[X_1 = \mathfrak{a}]$, and let us define the event

$$\mathcal{F}_n := \left\{ \omega \in \Omega : \frac{n}{|\mathcal{J}|} \geq \frac{1}{2q_{\mathfrak{a}}} \right\}.$$

Since X has i.i.d. entries, McDiarmid's Inequality – see Lemma 6.3 below – implies that for all $n \in \mathbb{N}$,

$$(4.1) \quad \mathbb{P}[\mathcal{F}_n] \geq 1 - e^{-n \frac{q_{\mathfrak{a}}^2}{2}}.$$

Next, let

$$\varepsilon := \frac{1}{4q_{\mathfrak{a}}} \langle \vec{p}_S, (S_{\mathfrak{bc}} - S_{\mathfrak{ac}})_{\mathfrak{c}} \rangle := \frac{1}{4q_{\mathfrak{a}}} \sum_{\mathfrak{c} \in \mathcal{A}^*} p_{\mathfrak{ac}} (S_{\mathfrak{bc}} - S_{\mathfrak{ac}}),$$

where $q_{\mathfrak{a}} = \mathbb{P}[X_1 = \mathfrak{a}]$. By continuity of inner products, there exists a $\delta > 0$ so that for any $\vec{p} \in B_{\delta}(\vec{p}_S)$, we have

$$(4.2) \quad \frac{1}{2q_{\mathfrak{a}}} \langle \vec{p}, (S_{\mathfrak{bc}} - S_{\mathfrak{ac}})_{\mathfrak{c}} \rangle \geq \varepsilon.$$

Recall now the notations $SET^*(X, Y)$ and $\mathcal{D}_n(\vec{p}, \epsilon)$ introduced in Section 3. Theorem 3.1 shows that there exists $K_{\delta} > 0$ such that the probability that all optimal alignments of X and Y have empirical distributions that lie within a distance δ of $\vec{p}_S$ equals

$$(4.3) \quad \mathbb{P}[SET^*(X, Y) \subseteq B_{\delta}(\vec{p}_S)] = \mathbb{P}[\mathcal{D}_n^c(\vec{p}_S, \delta)] \geq 1 - e^{-K_{\delta} n}, \quad \forall n \in \mathbb{N}.$$

But when $\mathcal{D}_n^c(\vec{p}_S, \delta)$ occurs, then for any optimal alignment π_n^* of X and Y , (4.2) holds with $\vec{p} = \vec{p}_{\pi_n^*}(X, Y)$. Denoting the components of $\vec{p}_{\pi_n^*}(X, Y)$ by $p_{\mathfrak{c}\mathfrak{d}}^*$, where $(\mathfrak{c}, \mathfrak{d})$ are pairs

of letters from $\mathcal{A}^*$, we have

$$\begin{aligned}
& \mathbb{P} \left[\mathbb{E}_{\tilde{\mathcal{P}}} \left[\tilde{L}_n(S) - L_n(S) \parallel X, Y \right] \geq \varepsilon \parallel \mathcal{D}_n^c(\vec{p}_S, \delta), \mathcal{F}_n \right] \\
& \geq \mathbb{P} \left[\mathbb{E}_{\tilde{\mathcal{P}}} \left[S_{\pi_n^*}(\tilde{X}, Y) - S_{\pi_n^*}(X, Y) \parallel X, Y \right] \geq \varepsilon \parallel \mathcal{D}_n^c(\vec{p}_S, \delta), \mathcal{F}_n \right] \\
& = \mathbb{P} \left[\frac{n}{|\mathcal{J}|} \sum_{c \in \mathcal{A}^*} p_{ac}^* (S_{bc} - S_{ac}) \geq \varepsilon \parallel \mathcal{D}_n^c(\vec{p}_S, \delta), \mathcal{F}_n \right] \\
(4.4) \quad & \stackrel{(4.2)}{\geq} \mathbb{P} \left[\frac{2nq_a}{|\mathcal{J}|} \varepsilon \geq \varepsilon \parallel \mathcal{D}_n^c(\vec{p}_S, \delta), \mathcal{F}_n \right] = 1
\end{aligned}$$

Therefore,

$$\begin{aligned}
& \mathbb{P} \left[\mathbb{E}_{\tilde{\mathcal{P}}} \left[\tilde{L}_n(S) - L_n(S) \parallel X, Y \right] \geq \varepsilon \right] \\
& \geq \mathbb{P} \left[\mathbb{E}_{\tilde{\mathcal{P}}} \left[\tilde{L}_n(S) - L_n(S) \parallel X, Y \right] \geq \varepsilon \parallel \mathcal{D}_n^c(\vec{p}_S, \delta), \mathcal{F}_n \right] \times \mathbb{P} \left[\mathcal{D}_n^c(\vec{p}_S, \delta), \mathcal{F}_n \right] \\
& \stackrel{(4.1), (4.3), (4.4)}{\geq} 1 - e^{-n \frac{q_a^2}{2}} - e^{-K_\delta n}, \quad \forall n \in \mathbb{N}.
\end{aligned}$$

Thus, the conditions of Theorem 4.1 are met for $\beta > 0$ small enough, and the claimed order of fluctuation holds. $\square$

5. APPENDIX: CONVEX GEOMETRY

We will now present geometric results required in the analysis of earlier parts of this paper. S^{n-1} will denote the unit sphere in $\mathbb{R}^n$, $B_\rho(\vec{x})$ the Euclidean ball of radius ρ around $\vec{x} \in \mathbb{R}^n$, d the Hausdorff distance, $\text{conv}(\cdot)$ the convex hull and $\text{cl}(\cdot)$ the closure of a set in the canonical subspace topology inherited from $\mathbb{R}^n$. We say that a convex set $C \subset \mathbb{R}^n$ has dimension k if its affine hull $\text{aff}(C) \subset \mathbb{R}^n$ has dimension k .

Theorem 5.1. *Let $C \subset \mathbb{R}^n$ be nonempty compact convex, and let $\vec{S} : \Omega \rightarrow S^{n-1}$ be a random vector that takes values in the unit sphere with uniform distribution, defined on some probability space $(\Omega, \mathcal{A}, \mathbb{P})$. Then for almost all $\omega \in \Omega$, the optimization problem $\arg \max_{\vec{y} \in C} \langle \vec{S}(\omega), \vec{y} \rangle$ has a unique solution.*

Proof. Let us first consider the case where C has nonempty interior. Upon a shift of C we may assume without loss of generality that $\vec{0}$ lies in the interior of C . Then the polar




FIGURE 1. The geometry of the Lipschitz estimate.

of C ,

$$C^\circ = \{\vec{w} \in \mathbb{R}^n : \langle \vec{w}, \vec{y} \rangle \leq 1, \forall \vec{y} \in C\},$$

is also compact convex with nonempty interior. Seen as the claim of the theorem is invariant under positive scaling, we may further assume without loss of generality that

$$B_3(\vec{0}) \subset C^\circ \subset B_\varrho(\vec{0}).$$

Next, let $\vec{s} \in S^{n-1}$ be a given point on the unit sphere and consider the function

$$\begin{aligned} \tau_{\vec{s}} : T_{\vec{s}} S^{n-1} &\rightarrow \mathbb{R}, \\ \vec{w} &\mapsto \max \{\tau > 0 : \tau \vec{w} \in C^\circ\} \end{aligned}$$


defined on the tangent space at $\vec{s}$. We claim that $\tau_{\vec{s}}$ is Lipschitz continuous on a sufficiently small neighbourhood $\mathcal{V}_{\vec{s}}$ of $\vec{s}$ in $T_{\vec{s}} S^{n-1} \cap B_2(\vec{0})$. Let $\vec{w}_1, \vec{w}_2 \in T_{\vec{s}} S^{n-1} \cap B_2(\vec{0})$ and $W = \text{span}\{\vec{w}_1, \vec{w}_2\}$. For $(i = 1, 2)$ we then have

$$(5.1) \quad 1 \leq \|\vec{w}_i\| < 2,$$

$$(5.2) \quad \tau_{\vec{s}}(\vec{w}_i) = \max \{\tau > 0 : \tau \vec{w}_i \in C^\circ \cap W\},$$

$$(5.3) \quad 1 < \tau_{\vec{s}}(\vec{w}_i) \|\vec{w}_i\| \leq \varrho.$$

By (5.2), we may assume without loss of generality that $\mathbb{R}^n = W$ for the purposes of proving $|\tau_{\vec{s}}(\vec{w}_1) - \tau_{\vec{s}}(\vec{w}_2)| \leq L \|\vec{w}_1 - \vec{w}_2\|$. We refer the reader to Figure 1 for an illustration of the geometric setup. The lines a and b are the tangents from $\tau_{\vec{s}}(\vec{w}_1)\vec{w}_1$ to the unit sphere S^1 in W . Denote the angle between the line $\overline{\vec{w}_1 \vec{w}_2}$ and the horizontal at $\vec{w}_1$ by θ , the angle

FIGURE 2. Bounding $\tau_{\vec{s}}(\vec{w}_2)$ by ratios.

between the horizontal at $\tau_{\vec{s}}(\vec{w}_1)\vec{w}_1$ and the tangents a, b by α , and the angle between the horizontal at $\vec{w}_1$ and the two tangents from $\vec{w}_1$ to S^1 by β . Since the affine hull $\text{aff}(\vec{w}_1, \vec{w}_2)$ cannot enter $B_1(\vec{0})$, it must lie wedged between the latter two tangents. In combination with (5.1), this implies

$$(5.4) \quad |\theta| \leq \beta = \frac{\pi}{2} - \arcsin \frac{1}{\|\vec{w}_1\|} \leq \frac{\pi}{2} - \arcsin \frac{1}{\varrho}.$$

Further, (5.3) implies

$$(5.5) \quad \alpha = \frac{\pi}{2} - \arcsin \frac{1}{\tau_{\vec{s}}(\vec{w}_1) \|\vec{w}_1\|} \leq \frac{\pi}{2} - \arcsin \frac{1}{\varrho}.$$

Observe that, by convexity, the line segment between the point of tangency of a at S^1 and $\tau_{\vec{s}}(\vec{w}_1)\vec{w}_1$ lies in C° , and further that the definition of $\tau_{\vec{s}}(\vec{w}_1)$ implies $\tau_{\vec{s}}(\vec{w}_1)\vec{w}_1 \in \partial C^\circ$. Therefore, the segment of a above $\tau_{\vec{s}}(\vec{w}_1)\vec{w}_1$ lies outside C° , and it follows that

$$(5.6) \quad \frac{\|B\|}{\|C\|} \leq \tau_{\vec{s}}(\vec{w}_2) \leq \frac{\|A\|}{\|D\|},$$

see Figure 2. Let φ be the angle between $\vec{w}_1$ and $\vec{w}_2$, and let us assume $\varphi < (\pi - 2\theta)/2$, so that the intersection points A, B, C, D exist. This assumption is equivalent to limiting our analysis to a sufficiently small neighbourhood of $\vec{s}$ in $T_{\vec{s}}S^{n-1}$, as assumed earlier. We

can now express the inequalities (5.6) in terms of the angles we introduced,

$$\tau_{\vec{s}}(\vec{w}_1) \frac{1 - \tan \varphi \tan \beta}{1 + \tan \varphi \tan \alpha} \leq \tau_{\vec{s}}(\vec{w}_2) \leq \tau_{\vec{s}}(\vec{w}_1) \frac{1 + \tan \varphi \tan \beta}{1 - \tan \varphi \tan \alpha}.$$

This can be simplified by Taylor expansion,

$$(5.7) \quad |\tau_{\vec{s}}(\vec{w}_2) - \tau_{\vec{s}}(\vec{w}_1)| \leq \tau_{\vec{s}}(\vec{w}_1) \tan \varphi (\tan \alpha + \tan \beta) \stackrel{(5.1),(5.3),(5.4),(5.5)}{\leq} \varrho \tan \varphi \left(\varrho \sqrt{1 - \varrho^{-2}} + \sqrt{3} \right),$$

and since $\|\vec{w}_1 - \vec{w}_2\| \geq \|\vec{w}_1\| \tan \varphi$, Equations (5.1) and (5.7) imply

$$|\tau_{\vec{s}}(\vec{w}_1) - \tau_{\vec{s}}(\vec{w}_2)| \leq L \|\vec{w}_1 - \vec{w}_2\|,$$

with $L = \varrho(\varrho \sqrt{1 - \varrho^{-2}} + \sqrt{3})$.

Next, having shown that $\tau_{\vec{s}}$ is Lipschitz continuous on a sufficiently small open neighbourhood $\mathcal{V}_{\vec{s}} \subset T_{\vec{s}}S^{n-1}$ of $\vec{s}$, Rademacher's Theorem [18] implies that $\tau_{\vec{s}}$ is Fréchet-differentiable everywhere on $\mathcal{V}_{\vec{s}}$ except on a null-set $\mathcal{B}_{\vec{s}} \subset \mathcal{V}_{\vec{s}}$. We now claim that if the optimization problem

$$(5.8) \quad \vec{x}(\vec{s}) = \arg \max_{\vec{y} \in C} \langle \vec{s}, \vec{y} \rangle$$

has multiple solutions, then $\tau_{\vec{s}}$ is Gâteaux-nondifferentiable at $\vec{s}$. Since $\tau_{\vec{s}}$ is then also Fréchet nondifferentiable at $\vec{s}$, it must be the case that $\vec{s} \in \mathcal{B}_{\vec{s}}$. Let us thus suppose that (5.8) has two different solutions, $\vec{x}_0 \neq \vec{x}_1$. Then $\langle \vec{s}, \vec{x}_1 - \vec{x}_0 \rangle = 0$, so that we have $c_1 := \langle \vec{s}, \vec{x}_1 \rangle = \langle \vec{s}, \vec{x}_0 \rangle$. Furthermore, writing $c_2 := \langle \vec{x}_1 - \vec{x}_0, \vec{x}_0 \rangle$ and $c_3 := \langle \vec{x}_1 - \vec{x}_0, \vec{x}_1 \rangle$, our assumption that $\vec{x}_0 \neq \vec{x}_1$ implies $c_2 \neq c_3$. Without loss of generality we may assume that $c_2 < c_3$. For all $\xi \in \mathbb{R}$ let us define $\vec{w}_{\xi} := \vec{s} + \xi(\vec{x}_1 - \vec{x}_0)$ and consider the restriction $\tau_{\vec{s}}|_{\vec{s} + \text{span}(\vec{x}_1 - \vec{x}_0)}$ which we shall denote by $\tau(\xi) := \tau_{\vec{s}}(\vec{w}_{\xi})$. Clearly, if $\tau(\xi)$ is nondifferentiable at $\xi = 0$, then $\tau_{\vec{s}}(\vec{w})$ is Gâteaux-nondifferentiable at $\vec{w} = \vec{s}$. The definition of $\tau(\xi)$ implies $\langle \tau(\xi) \vec{w}_{\xi}, \vec{x}_j \rangle \leq 1$ for $(j = 0, 1)$, so that

$$\begin{aligned} \tau(\xi) &\leq \min \left(\frac{1}{c_1 + c_2 \xi}, \frac{1}{c_1 + c_3 \xi} \right) \\ &= \frac{1}{c_1} \min \left(1 - \frac{c_2}{c_1} \xi + \mathcal{O}(\xi^2), 1 - \frac{c_3}{c_1} \xi + \mathcal{O}(\xi^2) \right). \end{aligned}$$

Furthermore, we have $\tau(0) = 1/c_1$. Therefore,

$$\frac{d}{d\xi_+}\tau(0) = \lim_{\xi \rightarrow 0^+} \frac{\tau(\xi) - \frac{1}{c_1}}{\xi} \leq -\frac{c_3}{c_1^2} < -\frac{c_2}{c_1^2} \leq \lim_{\xi \rightarrow 0^-} \frac{\tau(\xi) - \frac{1}{c_1}}{\xi} = \frac{d}{d\xi_-}\tau(0),$$

showing that $\tau(\xi)$ is nondifferentiable at $\xi = 0$, as claimed.

Next, observe that $\tau_{\vec{s}}$ is Fréchet differentiable at $\vec{w} \in \mathcal{V}_{\vec{s}}$ if and only if the map

$$\begin{aligned} \hat{\tau} : S^{n-1} &\rightarrow \mathbb{R}, \\ \vec{z} &\mapsto \max \{ \tau > 0 : \tau \vec{z} \in C^\circ \} \end{aligned}$$

is differentiable at $\hat{w} := \vec{w}/\|\vec{w}\|$ and if and only if $\tau_{\hat{w}}$ is differentiable at $\hat{w}$. Denoting the spherical projections of $\mathcal{V}_{\vec{s}}$ and $\mathcal{B}_{\vec{s}}$ by $\hat{\mathcal{V}}_{\vec{s}}$ and $\hat{\mathcal{B}}_{\vec{s}}$, the compactness of S^{n-1} implies the existence of finitely many points $\vec{s}_1, \dots, \vec{s}_k \in S^{n-1}$ such that $\cup_{i=1}^k \mathcal{V}_{\vec{s}_i} = S^{n-1}$. Consequently,

$$\mathcal{B} = \cup_{i=1}^k \mathcal{B}_{\vec{s}_i}$$

is a null-set with the property that if Problem (5.8) has multiple solutions for a given $\vec{s} \in S^{n-1}$, then $\vec{s} \in \mathcal{B}$. This proves the claim of the theorem in the case where C has nonempty interior.

Let us now consider the general case. When C consists of a singleton, the claim of the theorem is trivial. We may thus assume that $\dim(C) \geq 1$. Upon a shift we may assume without loss of generality that $\vec{0} \in C$. Let $W = \text{span}(C)$ be the subspace spanned by C , and $W^\perp$ its orthogonal complement under the Euclidean inner product of $\mathbb{R}^n$. We denote the orthogonal projections onto these spaces by π_W and $\pi_{W^\perp}$ respectively. Finally, let $S_W = S^{n-1} \cap W$ be the unit sphere in W , and

$$\begin{aligned} \pi_S : S^{n-1} &\rightarrow S_W, \\ \vec{s} &\mapsto \frac{\pi_W(\vec{s})}{\|\pi_W(\vec{s})\|} \end{aligned}$$

the rescaled projection of S^{n-1} into W .

The condition $\dim(C) \geq 1$ implies $\dim(W^\perp) \leq n-1$, and $\mathcal{B}_{W^\perp} = \{\omega \in \Omega : \vec{S}(\omega) \in W^\perp\}$ is a null-set. Hence, $\pi_S(\vec{s})$ is defined for almost all $\vec{s} \in S^{n-1}$. Further, by isotropy of the uniform distribution on S^{n-1} , the random vector

$$\pi_S(\vec{S}) : \Omega \setminus \mathcal{B}_{W^\perp} \rightarrow S_W$$

is uniformly distributed on S_W . Since C has nonempty interior in the subspace topology of W , the case we already settled above applies and implies that

$$\mathcal{B}_W = \left\{ \omega \in \Omega \setminus \mathcal{B}_{W^\perp} : \arg \max_{\vec{y} \in C} \langle \pi_S(\vec{S}(\omega)), \vec{y} \rangle \text{ is nonunique} \right\}$$

is a null-set. Observing that for $\vec{s} \in S^{n-1} \setminus W^\perp$ it is the case that

$$\arg \max_{\vec{y} \in C} \langle \vec{s}, \vec{y} \rangle = \arg \max_{\vec{y} \in C} \langle \pi_S(\vec{S}(\omega)), \vec{y} \rangle,$$

we find that $\arg \max_{\vec{y} \in C} \langle \vec{S}(\omega), \vec{y} \rangle$ has a unique solution if and only if ω is not in the null-set $\mathcal{B} = \mathcal{B}_{W^\perp} \cup \mathcal{B}_W$. $\square$

The following notion will play a key role in the sequel.

Definition 5.1. *Let $C \subset \mathbb{R}^n$ be convex compact. We say that a boundary point $\vec{x} \in \partial C$ is a point of strict curvature if there exists $\vec{s} \in S^{n-1}$ such that the optimization problem $\max_{\vec{y} \in C} \langle \vec{s}, \vec{y} \rangle$ has $\vec{x}$ as unique maximizer. We denote the set of points of strict curvature by C_{SE} .*

Note that if C has a differentiable boundary, then any point where all principal curvatures are nonzero is a point of strict curvature. However, the set of points of strict curvature may be larger. For example, the epigraph of the curve $x \mapsto |x|^3$ has zero curvature at $x = 0$, but this is a point of strict curvature nonetheless under our definition. Furthermore, Definition 5.1 also applies to points where ∂C is nondifferentiable and principal curvatures are not defined. For example, vertices of polytopes are points of strict curvature, while points on edges (1-faces) are not. Proposition 5.1 below also provides further intuition.

For the purposes of the next result, let us recall that the normal cone of C at $\vec{x} \in C$ is defined as follows,

$$N_{\vec{x}} C = \{ \vec{s} \in \mathbb{R}^n : \langle \vec{s}, \vec{x} - \vec{w} \rangle \geq 0, \forall \vec{w} \in C \},$$

or equivalently,

$$\begin{aligned} N_{\vec{x}} C &= \{ \vec{s} \in \mathbb{R}^n : \vec{x} = \arg \max_{\vec{w} \in C} \langle \vec{s}, \vec{w} \rangle \} \\ (5.9) \quad &= \{ \tau \vec{s} : \tau \geq 0, \vec{s} \in S^{n-1}, \vec{x} = \arg \max_{\vec{w} \in C} \langle \vec{s}, \vec{w} \rangle \}. \end{aligned}$$

By the dual description of C , it is the case that

$$(5.10) \quad N_{\vec{x}} C \cap S^{n-1} \cap \text{span}(C) \neq \emptyset$$

if and only if $\vec{x} \in \partial C$, see also Lemma 5.1 c). We denote the set of extreme points³ of C by C_E .

Proposition 5.1. *For any $C \subset \mathbb{R}^n$ nonempty convex compact, the following hold true:*

- a) $\vec{x} \in C_{SE}$ if and only if there exists $\vec{s} \in N_{\vec{x}} C \cap S^{n-1}$ and sequences $(\delta_k)_{k \in \mathbb{N}}, (\epsilon_k)_{k \in \mathbb{N}} \subset \mathbb{R}_+$ such that $\epsilon_k, \delta_k \rightarrow 0$ as k tends to infinity and

$$\{\vec{y} \in C : N_{\vec{y}} C \cap S^{n-1} \cap B_{\delta_k}(\vec{s}) \neq \emptyset\} \subset B_{\epsilon_k}(\vec{x}), \quad \forall k \in \mathbb{N}.$$

- b) $C_{SE} \subseteq C_E \subseteq \text{cl}(C_{SE})$.

- c) $\{\vec{x} \in \partial C : N_{\vec{x}} C \cap N_{\vec{v}} C = \text{span}(C)^\perp, \forall \vec{v} \in C \setminus \{\vec{x}\}\} \subset C_{SE}$.

- d) Let $\vec{x}_0 \in C_{SE}$ and $\vec{s}_0 \in N_{\vec{x}_0} C \cap S^{n-1}$ be chosen such that $\vec{x}_0$ is the unique maximizer of $\max_{\vec{y} \in C} \langle \vec{s}_0, \vec{y} \rangle$, and let $\epsilon > 0$ be given. Then there exist finitely many points $\vec{x}_i \in C$ and normal vectors $\vec{s}_i \in N_{\vec{x}_i} C \cap S^{n-1}$, $(i = 1, \dots, k)$, such that

$$C(\xi_0, \dots, \xi_k) := \{\vec{x} \in \mathbb{R}^n : \langle \vec{s}_0, \vec{x} - \vec{x}_0 \rangle \geq \xi_0\} \cap \bigcap_{i=1}^k \{\vec{x} \in \mathbb{R}^n : \langle \vec{s}_i, \vec{x} - \vec{x}_i \rangle \leq \xi_i\}$$

is compact for all $(\xi_0, \dots, \xi_k) \in \mathbb{R}^{k+1}$, and $C(0, \dots, 0) \subset B_\epsilon(\vec{x}_0)$.

Proof. a) Let $\vec{x} \in C_{SE}$, and let $\vec{s} \in N_{\vec{x}} C \cap S^{n-1}$ such that $\vec{x}$ is the unique maximizer of

$$(5.11) \quad \max_{\vec{y} \in C} \langle \vec{s}, \vec{y} \rangle.$$

Let $(\delta_k)_{k \in \mathbb{N}} \subset \mathbb{R}_+$ be a sequence such that $\delta_k \rightarrow 0$. We claim that there exists a sequence $(\epsilon_k)_{k \in \mathbb{N}}$ with the required properties. Supposing the claim to be wrong, there exists an $\epsilon > 0$ and sequences $(\vec{x}_k)_{k \in \mathbb{N}}, (\vec{s}_k)_{k \in \mathbb{N}}$ such that

$$\begin{aligned} \vec{x}_k &\in C \setminus B_\epsilon(\vec{x}), \\ \vec{s}_k &\in N_{\vec{x}_k} C \cap S^{n-1} \cap B_{\delta_k}(\vec{s}). \end{aligned}$$

³A point $\vec{x} \in C$ is an extreme point of C if it cannot be written as a convex combination of two points $\vec{y}, \vec{z} \in C \setminus \{\vec{x}\}$.

Since $\delta_k \rightarrow 0$, we have $\vec{s}_k \rightarrow \vec{s}$, and $C \setminus B_\epsilon(\vec{x})$ being compact, we may assume without loss of generality that $\vec{x}_k \rightarrow \vec{x}_*$ for some $\vec{x}_* \in C \setminus B_\epsilon(\vec{x})$. By (5.9), this shows that $\vec{x}$ is not the unique maximizer of (5.11) and contradicts the membership of $\vec{x}$ in C_{SE} . Our claim is thus correct, and this establishes the “only if” part.

To prove the converse, let us assume that $(\delta_k)_{k \in \mathbb{N}}$ and $(\epsilon_k)_{k \in \mathbb{N}}$ with the required properties exist. By (5.9), all maximizers $\vec{x}_*$ of (5.11) must satisfy $\vec{s} \in N_{\vec{x}_*} C$, and since $\vec{s} \in S^{n-1} \cap B_{\delta_k}(\vec{s})$, this implies $\vec{x}_* \in B_{\epsilon_k}(\vec{x})$ for all $k \in \mathbb{N}$, which can only be true if $\vec{x}_* = \vec{x}$. This shows that $\vec{x}$ is the unique maximizer of (5.11), and hence $\vec{x} \in C_{SE}$.

b) For any point $\vec{x} \in C \setminus C_E$ there exist two other points $\vec{v}, \vec{w} \in C$ of which $\vec{x}$ is a convex combination, $\vec{x} = \xi \vec{v} + (1 - \xi) \vec{w}$. Let $\vec{s} \in N_{\vec{x}} C \cap S^{n-1}$, so that $\vec{x}$ is a maximizer of (5.11). The existence of such an $\vec{s}$ is guaranteed by (5.9). By convexity, we have $\langle \vec{s}, \vec{x} \rangle \leq \max(\langle \vec{s}, \vec{v} \rangle, \langle \vec{s}, \vec{w} \rangle)$. Without loss of generality, we may assume that the maximum is achieved at $\langle \vec{s}, \vec{w} \rangle$, and therefore, $\vec{w}$ is also a maximizer of (5.11). Since this construction works for any $\vec{s} \in N_{\vec{x}} C \cap S^{n-1}$, this shows that $\vec{x} \notin C_{SE}$, and hence, $C_{SE} \subseteq C_E$.

Next, we claim that $K_1 := \text{cl}(\text{conv}(C_{SE})) = C$. Assuming our claim to be wrong, there exists $\vec{x} \in C \setminus K_1$, and since $K_1 \subset C$ is compact, the Hahn-Banach Separation Theorem then implies that there exists $\vec{s} \in S^{n-1}$ such that

$$(5.12) \quad \eta := \max_{\vec{y} \in K_1} \langle \vec{s}, \vec{y} \rangle < \xi := \langle \vec{s}, \vec{x} \rangle.$$

By Theorem 5.1, there exist sequences $(\vec{s}_k)_{k \in \mathbb{N}} \subset S^{n-1}$ and $(\vec{y}_k)_{k \in \mathbb{N}} \subset C_{SE}$ such that

$$\begin{aligned} \vec{s}_k &\xrightarrow{k \rightarrow \infty} \vec{s}, \\ \vec{y}_k &= \arg \max_{\vec{y} \in C} \langle \vec{s}, \vec{y} \rangle, \end{aligned}$$

and since K_1 is compact, we may further assume without loss of generality that $\vec{y}_k \rightarrow \vec{y}_* \in K_1$. Using the continuity of the function $\vec{s} \mapsto \max_{\vec{y} \in C} \langle \vec{s}, \vec{y} \rangle$, we thus find that

$$\xi = \lim_{k \rightarrow \infty} \langle \vec{s}_k, \vec{y}_k \rangle = \lim_{k \rightarrow \infty} \langle \vec{s}_k, \vec{y}_* \rangle \leq \eta.$$

This contradicts (5.12) and proves our claim to be true.

Next, let $\vec{x} \in C_E$. By what we know so far, there exists a sequence of points $(\vec{x}_k)_{k \in \mathbb{N}} \subset \text{conv}(C_{SE})$ that converges to $\vec{x}$. The Carathéodory-Steinitz Theorem [20] then implies

that each $\vec{x}_k$ can be written as a convex combination

$$\vec{x}_k = \sum_{i=1}^{\dim(C)+1} \theta_i^k \vec{z}_{i,k}$$

of at most $\dim(C) + 1$ points $\vec{z}_{i,k} \in C_{SE}$, ($i = 1, \dots, \dim(C) + 1$). By compactness of

$$\Delta_{\dim(C)+1} \otimes \left(\bigotimes_{i=1}^{\dim(C)+1} \text{cl}(C_{SE}) \right),$$

where $\Delta_{\dim(C)+1}$ denotes the $\dim(C)$ -dimensional simplex, we may assume without loss of generality that

$$\begin{aligned} \theta^k &\xrightarrow{k \rightarrow \infty} \theta^* \in \Delta_{\dim(C)+1}, \\ \vec{z}_{i,k} &\xrightarrow{k \rightarrow \infty} \vec{z}_{i,*} \in \text{cl}(C_{SE}), \end{aligned}$$

so that

$$\vec{x} = \lim_{k \rightarrow \infty} \vec{x}_k = \sum_{i=1}^{\dim(C)+1} \theta_i^* \vec{z}_{i,*} \in \text{conv}(\text{cl}(C_{SE})) \subset C = \text{cl}(\text{conv}(C_{SE})).$$

Seen as $\vec{x}$ is an extreme point of C , it must also be an extreme point of the subset $\text{conv}(\text{cl}(C_{SE}))$. Hence, all $\vec{z}_{i,*}$ must be identical. This shows that $\vec{x} = \vec{z}_{1,*} \in \text{cl}(C_{SE})$ and proves the inclusion $C_E \subseteq \text{cl}(C_{SE})$.

c) This follows directly from Equations (5.9) and (5.10).

d) By the dual description of C , we have

$$\bigcap_{\vec{x} \in C} \bigcap_{\vec{s} \in \mathbb{N}_{\vec{x}} C \cap \mathbb{S}^{n-1}} \{\vec{y} \in \mathbb{R}^n : \langle \vec{s}, \vec{y} - \vec{x} \rangle \leq 0\} = C,$$

and since it suffices to take this intersection over a dense subset of $\vec{x} \in C$, we have

$$(5.13) \quad \bigcap_{\vec{x} \in \partial C \setminus \{\vec{x}_0\}} \bigcap_{\vec{s} \in \mathbb{N}_{\vec{x}} C \cap \mathbb{S}^{n-1}} \{\vec{y} \in \mathbb{R}^n : \langle \vec{s}, \vec{y} - \vec{x} \rangle \leq 0\} = C.$$

By compactness of C there exists $\rho > 0$ such that $C \subset \text{cl}(B_\rho(\vec{0}))$. Consider the compact set $K_2 = \text{cl}(B_\rho(\vec{0})) \setminus B_\varepsilon(\vec{x}_0)$. By the assumed properties of $\vec{x}_0$ and $\vec{s}_0$, it is further true that

$$(5.14) \quad \{\vec{y} \in \mathbb{R}^n : \langle \vec{s}_0, \vec{y} - \vec{x}_0 \rangle \geq 0\} \cap C = \{\vec{x}_0\}.$$

Equations (5.13) and (5.14) now show that

$$K_2 \subset \{\vec{y} \in \mathbb{R}^n : \langle \vec{s}_0, \vec{y} - \vec{x}_0 \rangle < 0\} \cup \bigcup_{\vec{x} \in C \setminus \{\vec{x}_0\}} \bigcup_{\vec{s} \in N_{\vec{x}} C \cap S^{n-1}} \{\vec{y} \in \mathbb{R}^n : \langle \vec{s}, \vec{y} - \vec{x} \rangle > 0\}.$$

By compactness of K_2 , there exist finitely many $\vec{x}_i \in C \setminus \{\vec{x}_0\}$ and $\vec{s}_i \in N_{\vec{x}_i} C \cap S^{n-1}$, ($i = 1, \dots, k$), such that

$$K_2 \subset \{\vec{y} \in \mathbb{R}^n : \langle \vec{s}_0, \vec{y} - \vec{x}_0 \rangle < 0\} \cup \bigcup_{i=1}^k \{\vec{y} \in \mathbb{R}^n : \langle \vec{s}_i, \vec{y} - \vec{x}_i \rangle > 0\}.$$

Let $\vec{e}_j$ be the j -th coordinate vector in $\mathbb{R}^n$. Let us write $\vec{s}_{j,\ell} = (-1)^\ell \vec{e}_j$ for $\ell \in \{0, 1\}$, and let

$$\vec{x}_{j,\ell} \in \arg \max_{\vec{w} \in C} \langle \vec{s}_{j,\ell}, \vec{w} \rangle.$$

Then $\vec{s}_{j,\ell} \in N_{\vec{x}_{j,\ell}} C$, and

$$Q := \bigcap_{\ell=0}^1 \bigcap_{j=1}^n \{\vec{y} \in \mathbb{R}^n : \langle \vec{s}_{j,\ell}, \vec{y} - \vec{x}_{j,\ell} \rangle \leq 0\}$$

is a cuboid. By choosing ρ large enough for $Q \subset \text{cl}(B_\rho(\vec{0}))$ to hold, and by including the $(\vec{x}_{j,\ell}, \vec{s}_{j,\ell})$ among our list of points $\{(\vec{x}_i, \vec{s}_i) : i = 1, \dots, k\}$ if necessary, we can guarantee that

$$\{\vec{y} \in \mathbb{R}^n : \langle \vec{s}_0, \vec{y} - \vec{x}_0 \rangle \geq 0\} \cap \bigcap_{i=1}^k \{\vec{y} \in \mathbb{R}^n : \langle \vec{s}_i, \vec{y} - \vec{x}_i \rangle \leq 0\} \subset Q \setminus K_2 \subset B_\varepsilon(\vec{x}_0),$$

and that $C(\xi_0, \dots, \xi_k)$ is compact (although possibly empty) for all $(\xi_0, \dots, \xi_k)$. $\square$

Proposition 5.2. *Let $C \subset \mathbb{R}^n$ be nonempty convex compact and $C^1, C^2, \dots$ compact subsets of $\mathbb{R}^n$ such that $d(C^n, C) \rightarrow 0$, where d denotes the Hausdorff distance. Let $\vec{s} \in S^{n-1}$ be such that the optimization problem*

$$(5.15) \quad \vec{x}_* = \arg \max_{\vec{x} \in C} \langle \vec{s}, \vec{x} \rangle$$

has a unique solution. For all $n \in \mathbb{N}$ let $\vec{x}_n$ be a solution of

$$\vec{x}_n = \arg \max_{\vec{x} \in C^n} \langle \vec{s}, \vec{x} \rangle.$$

Then $\vec{x}_n \rightarrow \vec{x}_$ as n tends to infinity.*

Proof. Let $f(\vec{x}) = \langle \vec{s}, \vec{x} \rangle$ be the linear functional defined by $\vec{s}$. We note that

$$\liminf_{n \rightarrow \infty} \max_{\vec{y} \in C^n} f(\vec{y}) \geq f(\vec{x}_*),$$

since for any $\epsilon > 0$ there exists n_ϵ such that for all $n \geq n_\epsilon$ there exists $\vec{y}_n \in A^n \cap B_\epsilon(\vec{x}_*)$, and we have

$$(5.16) \quad f(\vec{x}_*) = f(\vec{y}_n) + \langle \vec{s}, \vec{x}_* - \vec{y}_n \rangle < f(\vec{y}_n) + \epsilon.$$

Next, let $W = \text{span}(C)$ and $W^\perp$ be its orthogonal complement in $\mathbb{R}^n$. Upon shifting and rescaling, we may assume without loss of generality that $B_\rho(\vec{0}) \cap W \subset C \subset B_1(\vec{0}) \cap W$ for some $\rho > 0$, so that $f(\vec{x}_*) \leq 1$. The condition $d(C^n, C) \rightarrow 0$ then implies that for any $\epsilon > 0$ we may take n_ϵ to be large enough so that for $n \geq n_\epsilon$,

$$C^n \subset \left(1 + \frac{\epsilon}{\rho}\right) C \times \left(B_\epsilon(\vec{0}) \cap W^\perp\right).$$

Convexity of C and the uniqueness of $\vec{x}_*$ as a maximizer of (5.15) imply that there exists $\delta(\epsilon) > 0$ such that

$$(5.17) \quad f(\vec{x}) < f(\vec{x}_*) - \left(2 + \frac{f(\vec{x}_*)}{\rho}\right) \epsilon$$

for all $\vec{x} \in C \setminus B_{\delta(\epsilon)}(\vec{x}_*)$, and that $\delta(\epsilon) \rightarrow 0$ when $\epsilon \rightarrow 0$. Furthermore, we have

$$(5.18) \quad \sup_{\vec{y} \in B_\epsilon(\vec{0}) \cap W^\perp} f(\vec{y}) < \epsilon.$$

(5.17) and (5.18) imply that for all $\vec{x} \in (1 + \epsilon/\rho) (C \setminus B_{\delta(\epsilon)}(\vec{x}_*)) \times (B_\epsilon(\vec{0}) \cap W^\perp)$,

$$\begin{aligned} f(\vec{x}) &< \left(1 + \frac{\epsilon}{\rho}\right) \left(f(\vec{x}_*) - \left(2 + \frac{f(\vec{x}_*)}{\rho}\right) \epsilon\right) + \epsilon \\ &< f(\vec{x}_*) - \epsilon. \end{aligned}$$

In particular, this applies to all $\vec{x} \in C^n \setminus (1 + \epsilon/\rho) B_{\delta(\epsilon)}(\vec{x}_*)$, and (5.16) shows that for $n \geq n_\epsilon$ we have $\vec{x}_n \in (1 + \epsilon/\rho) B_{\delta(\epsilon)}(\vec{x}_*)$. We conclude that

$$\begin{aligned} \|\vec{x}_n - \vec{x}_*\| &\leq \left\| \vec{x}_n - \left(1 + \frac{\epsilon}{\rho}\right) \vec{x}_* \right\| + \left\| \left(1 + \frac{\epsilon}{\rho}\right) \vec{x}_* - \vec{x}_* \right\| \\ &\leq \left(1 + \frac{\epsilon}{\rho}\right) \delta(\epsilon) + \frac{\epsilon}{\rho}. \end{aligned}$$

Since we may choose $\epsilon, \delta \rightarrow 0$ when n is allowed to go to infinity, the result follows. $\square$

Next, we shall investigate the approximability of compact convex sets by polyhedra and polytopes. Results on outer approximations by polyhedra and algorithms to achieve this in practice are widespread in the literature on the cutting plane approach in numerical optimization, see e.g. Bertsekas [3]. Similar results for inner approximations by polytopes play a key role in Markov chain Monte Carlo methods for the estimation of the volume of high dimensional convex bodies, see e.g. Jerrum [11]. The literature in both areas is focused on algorithms and relies on separation or membership oracles. As a result, the constructions use outer approximations by cutting planes that do not necessarily touch the boundary of the convex body to be approximated, and likewise, inner approximations use generators that generally do not lie on the boundary either.

In contrast, the approximations required by our analysis have a crucial interplay with the boundary. For outer approximations, we would like cutting hyperplanes to be supported at points of strict curvature. Likewise, we would like inner approximations to be generated as the convex hull of points of strict curvature. Since we are not aware of such results appearing in the literature, we derive them from first principles.

Lemma 5.1. *Let $C \subset \mathbb{R}^n$ be a convex set with dual description*

$$(5.19) \quad C = \bigcap_{\vec{s} \in S^{n-1}} H_{\vec{s}},$$

where $H_{\vec{s}} = \{\vec{x} : \langle \vec{s}, \vec{x} \rangle \leq \lambda(\vec{s})\}$ for some continuous function $\vec{s} \mapsto \lambda(\vec{s}) \in \mathbb{R}$. Then the following hold true:

- a) C is compact.
- b) For any given $\epsilon > 0$, there exists a finite collection of points $\vec{s}_1, \dots, \vec{s}_k \in S^{n-1}$ for which

$$(5.20) \quad \max_{\vec{x} \in \bigcap_{i=1}^k H(\vec{s}_i)} d(\vec{x}, C) \leq \epsilon.$$

- c) For every point $\vec{x} \in \partial C$, there exists $s \in S^{n-1}$ such that

$$\langle \vec{s}, \vec{x} \rangle = \max_{\vec{y} \in C} \langle \vec{s}, \vec{y} \rangle = \lambda(\vec{s}).$$

Proof. a) Since C is a closed subset of the compact set

$$\bigcap_{i=1}^n H_{\vec{e}_i} \cap H_{-\vec{e}_i},$$

it is itself compact.

b) The continuity of λ and of the Hausdorff distance imply that the function $s \mapsto d(\vec{x}, H(s))$ for any fixed $\vec{x} \in \mathbb{R}^n$ is continuous. The dual description (5.19) can therefore be amended by taking the intersection over $\vec{s}$ in a dense subset $\mathcal{S} \subset S^{n-1}$ only, and since the unit sphere is a separable space, we may take $\mathcal{S}$ to be a countable set $\{\vec{s}_i : i \in \mathbb{N}\}$. Thus, we have

$$(5.21) \quad C = \bigcap_{i=1}^{\infty} H(\vec{s}_i).$$

Consider the nested sets

$$(5.22) \quad G_j := \bigcap_{i=1}^j H(\vec{s}_i).$$

Since C is compact, we may furthermore order the vectors $\vec{s}_i$ so that G_j is compact for all $j \geq 2n$. Our claim is clearly true if we can establish that

$$(5.23) \quad \max_{\vec{x} \in G_j} d(\vec{x}, C) \xrightarrow{j \rightarrow \infty} 0.$$

Assuming the contrary, there exists $\epsilon > 0$, a subsequence $(G_{j_i})_{i \in \mathbb{N}}$ and points $\vec{y}_{j_i} \in G_{j_i}$ such that

$$(5.24) \quad d(\vec{y}_{j_i}, C) \geq \epsilon, \quad \forall i \in \mathbb{N}.$$

Since all but finitely many terms of the sequence $(\vec{y}_{j_i})_{i \in \mathbb{N}}$ are contained in the compact set G_{2n} , there exists a convergent subsequence $(\vec{z}_k)_{k \in \mathbb{N}} \subseteq (\vec{y}_{j_i})_{i \in \mathbb{N}}$ with limit $\vec{z}_*$. By continuity of the distance function and by virtue of Equation (5.24), we have

$$d(\vec{z}_*, C) = \lim_{k \rightarrow \infty} d(\vec{z}_k, C) \geq \epsilon.$$

Since this is in direct contradiction with $\vec{z}_* \in \bigcap_{i \in \mathbb{N}} G_{j_i} = C$, implied by (5.21), our claim is true.

c) If $C = \emptyset$, there is nothing to prove. Hence, we may assume that C is nonempty. Without loss of generality, we may furthermore assume that C is shifted so that it contains

the origin and $\lambda(\vec{s}) \geq 0$ for all $\vec{s} \in S^{n-1}$. Let $\vec{x} \in \partial C$. Then, by the dual description (5.19),

$$\langle \vec{s}, \vec{x} \rangle \leq \max_{\vec{y} \in C} \langle \vec{s}, \vec{y} \rangle \leq \lambda(\vec{s})$$

for all $\vec{s} \in S^{n-1}$. Thus, it suffices to prove the existence of $\vec{s}$ such that $\langle \vec{s}, \vec{x} \rangle \geq \lambda(\vec{s})$. If $\vec{x} = \vec{0}$, then we may assume without loss of generality that C has empty interior, as any interior point could otherwise be shifted to the origin. By convexity, C then lies in a lower-dimensional subspace, and it suffices to take $\vec{s} \in \text{span}(C)^\perp$. If $\vec{x} \neq \vec{0}$, then $(1 + 1/k)\vec{x} \notin C$ for all $k \in \mathbb{N}$. By virtue of the Hahn-Banach Theorem there exist unit vectors $\vec{s}_k \in S^{n-1}$ such that

$$(5.25) \quad \langle \vec{s}_k, \vec{x} \rangle \geq \frac{\lambda(\vec{s}_k)}{1 + \frac{1}{k}}.$$

S^{n-1} being compact, we can extract a convergent subsequence and assume without loss of generality that $\vec{s}_k \rightarrow \vec{s}_* \in S^{n-1}$, as $k \rightarrow \infty$. By continuity of λ , (5.25) implies

$$\langle \vec{s}_*, \vec{x} \rangle \geq \lambda(\vec{s}_*).$$

□

Theorem 5.2. *Let C be as in Lemma 5.1 and nonempty. Then the points $\vec{s}_i$ in part b) of Lemma 5.1 can be chosen so that $\vec{x}_i = \arg \max_{\vec{y} \in C} \langle \vec{s}_i, \vec{y} \rangle$ is unique for all i , that is, $\vec{x}_i$ are points of strict curvature.*

Proof. By Theorem 5.1, the set

$$\mathcal{S} = \left\{ \vec{s} \in S^{n-1} : \arg \max_{\vec{y} \in C} \langle \vec{s}, \vec{y} \rangle \in C_{SE} \right\}$$

has full measure and is therefore everywhere dense in S^{n-1} . Since $\text{cl}(\mathcal{S}) = S^{n-1}$ is separable, there exists furthermore a countable subset $\{\vec{s}_i : i \in \mathbb{N}\} \subset \mathcal{S}$, also everywhere dense in S^{n-1} , that can be used in the construction. □

Theorem 5.3. *Let $C \subset \mathbb{R}^n$ be nonempty compact convex. Then for all $\epsilon > 0$ there exist finitely many points of strict curvature $\vec{x}_1, \dots, \vec{x}_k \in C_{SE}$ such that*

$$\max_{\vec{x} \in C} d(\vec{x}, \text{conv}(\vec{x}_1, \dots, \vec{x}_k)) \leq \epsilon.$$

Proof. Let $\{\vec{x}_1, \dots, \vec{x}_k\} \subset C_{SE}$ be an ϵ -net on the set C_E of extreme points of C , that is, $\vec{x}_i$ are chosen so that

$$\min_i \|\vec{z} - \vec{x}_i\| \leq \epsilon$$

for all $\vec{z} \in C_E$. The existence of such an ϵ -net is established as follows: C being compact, $\text{cl}(C_{SE})$ is a compact set too, and by the Heine-Borel Theorem, we can extract a finite covering by Euclidean balls of radius $\epsilon/2$ around points $\vec{y}_i \in \text{cl}(C_{SE})$, which by Proposition 5.1 b) is also a covering of C_E ,

$$\bigcup_{i=1}^k B_{\epsilon/2}(\vec{y}_i) \supset C_E.$$

Next, for all i choose $\vec{x}_i \in C_{SE}$ within distance $\epsilon/2$ of $\vec{y}_i$. It then follows from the triangular inequality that $\{\vec{x}_1, \dots, \vec{x}_k\}$ is the required ϵ -net.

By the Theorems of Minkowski⁴ [17] and Carathéodory⁵ [5], any point $\vec{x} \in C$ can be written as a convex combination $\vec{x} = \xi_1 \vec{z}_1 + \dots + \xi_m \vec{z}_m$ of $m \leq n + 1$ extreme points $z_j \in C_E$, and by construction of the ϵ -net, it is then possible to choose $1 \leq i_j \leq k$ such that $\|\vec{z}_j - \vec{x}_{i_j}\| \leq \epsilon$ for all j . Using the triangular inequality once again, we find that

$$\begin{aligned} d(\vec{x}, \text{conv}(\vec{x}_1, \dots, \vec{x}_k)) &\leq d(\vec{x}, \text{conv}(\vec{x}_{i_1}, \dots, \vec{x}_{i_m})) \\ &\leq d(\xi_1 \vec{z}_1 + \dots + \xi_m \vec{z}_m, \xi_1 \vec{x}_{i_1} + \dots + \xi_m \vec{x}_{i_m}) \leq \epsilon, \end{aligned}$$

as claimed. □

The next result can be seen as an intuitive template for Theorem 2.1 free of large deviations complications.

Theorem 5.4. *Let C be a nonempty convex compact subset of $\mathbb{R}^n$ with dual description (5.1), and let $C^1, C^2, \dots$ be convex compact subsets of $\mathbb{R}^n$ such that for all linear functionals $f : \mathbb{R}^n \rightarrow \mathbb{R}$ it is true that*

$$(5.26) \quad \max_{\vec{p} \in C^n} f(\vec{p}) \xrightarrow{n \rightarrow \infty} \max_{\vec{p} \in C} f(\vec{p}).$$

⁴This theorem says that a convex compact set in $\mathbb{R}^n$ is equal to the convex hull of its extreme points. The generalization of this result to arbitrary topological vector spaces is the Krein-Milman Theorem [13, 16].

⁵This result says that if $K = \text{conv}(X)$ for some $X \subset \mathbb{R}^n$, then every point in K is a convex combination of at most $n + 1$ elements of X .

Then $d(C^n, C) \rightarrow 0$.

Proof. Upon translation and rescaling we may assume without loss of generality that $\{0\} \in C \subset B_1(0)$. Let $\epsilon > 0$ be a given small number. By the dual description of C and Lemma 5.1 there exist finitely many vectors $\vec{s}_1, \dots, \vec{s}_k \in S^{n-1}$ such that

$$(5.27) \quad \sup_{\vec{x} \in C_\epsilon} d(\vec{x}, C) < \epsilon,$$

where $C_\epsilon = \cap_{i=1}^k \{\vec{x} : \langle \vec{s}_i, \vec{x} \rangle \leq \lambda_{s_i}\}$ and $\lambda_s = \arg \max_{\vec{x} \in C} \langle \vec{s}, \vec{x} \rangle \geq 0$. Because of Assumption (5.26), there exists $n_0 \in \mathbb{N}$ such that for all $n \geq n_0$,

$$\max_{\vec{x} \in C^n} \langle \vec{s}_i, \vec{x} \rangle \leq (1 + \epsilon) \lambda_{s_i}, \quad (i = 1, \dots, k),$$

and hence, $C^n \subseteq (1 + \epsilon)C_\epsilon$. Therefore, for all $n \geq n_0$,

$$\begin{aligned} \sup_{\vec{x} \in C^n} d(\vec{x}, C) &\leq \sup_{\vec{x} \in (1+\epsilon)C_\epsilon} d(\vec{x}, C) \\ &= (1 + \epsilon) \sup_{\vec{x} \in C_\epsilon} d\left(\vec{x}, \frac{1}{1 + \epsilon}C\right) \\ &\leq (1 + \epsilon) \left(\sup_{\vec{x} \in C_\epsilon} d(\vec{x}, C) + \sup_{\vec{x} \in C} d\left(\vec{x}, \frac{1}{1 + \epsilon}C\right) \right) \\ (5.28) \quad &\leq 2(1 + \epsilon)\epsilon \end{aligned}$$

where we used (5.27) and $d(C, 1/(1 + \epsilon)C_\epsilon) < \epsilon$ to arrive at the last inequality.

Next, let us choose points $\vec{x}_1, \dots, \vec{x}_k$ as in Theorem 5.3. Proposition 5.1 allows us to choose scoring functions $S_{\vec{x}_i}$ and $S_1, \dots, S_\ell$ such that $\vec{x}_i \in C_i \subset B_\epsilon(\vec{x}_i)$ and C_i^j is compact for all $j \in \mathbb{N}$, where

$$\begin{aligned} C_i &:= \left\{ \vec{x} : f_{S_{\vec{x}_i}}(\vec{x}) \geq \max_{\vec{x} \in C} f_{S_{\vec{x}_i}}(\vec{x}), f_{S_j}(\vec{x}) \leq \max_{\vec{x} \in C} f_{S_j}(\vec{x}), (j = 1, \dots, \ell) \right\}, \\ C_i^j &:= \left\{ \vec{x} : f_{S_{\vec{x}_i}}(\vec{x}) \geq \max_{\vec{x} \in C^j} f_{S_{\vec{x}_i}}(\vec{x}), f_{S_j}(\vec{x}) \leq \max_{\vec{x} \in C^j} f_{S_j}(\vec{x}), (j = 1, \dots, \ell) \right\}. \end{aligned}$$

Taking $\vec{x}_{ij}$ to be a maximizer of $\arg \max_{\vec{x} \in C^j} f_{S_{\vec{x}}}$, we find $\vec{x}_{ij} \in C^j \cap C_i^j \neq \emptyset$, and by (5.26) applied to $f_{S_{\vec{x}_i}}$ and f_{S_j} ($j = 1, \dots, k$) and the continuity of the Hausdorff distance, we have

$$\limsup_{n \rightarrow \infty} \|\vec{x}_i - \vec{x}_{ij}\| \leq \epsilon.$$

By the triangular inequality, this implies

$$\limsup_{n \rightarrow \infty} \max_{\vec{x} \in C} d(\vec{x}, C^n) \leq 2\epsilon, \quad \forall \epsilon > 0,$$

and since this is true for all ϵ rational, the Theorem follows. $\square$

The final result shows among other things that all results presented above are all applicable to SET , defined as in (1.10). Recall the notations $\lambda(S)$ for the Chvátal-Sankoff limit (1.7) of a scoring function S , and f_S , the linear functional associated with S as defined in (1.1).

Lemma 5.2.

- a) *The function $S \mapsto \lambda(S)$ is continuous.*
- b) *SET is nonempty convex compact.*
- c) *For every $\vec{x} \in \partial SET$ there exists $S_{\vec{x}} \neq 0$ such that $\{\vec{y} : f_{S_{\vec{x}}}(\vec{y}) = \lambda(S_{\vec{x}})\}$ is a tangent plane to SET supported at $\vec{x}$.*
- d) *$\max_{\vec{x} \in SET} f_S(\vec{x}) = \lambda(S)$ holds true for all scoring functions S .*

Proof. a) It suffices to show that for any two scoring functions S and T the following inequality holds true,

$$\lambda_n(S) - \lambda_n(T) \leq 2 \|T - S\|_{\infty}.$$

For any pair of letters $\mathfrak{c}, \mathfrak{d} \in \mathcal{A}^*$ we have

$$|S(\mathfrak{c}, \mathfrak{d}) - T(\mathfrak{c}, \mathfrak{d})| \leq \|S - T\|_{\infty} = \max_{\mathfrak{a}, \mathfrak{b} \in \mathcal{A}^*} |S(\mathfrak{a}, \mathfrak{b}) - T(\mathfrak{a}, \mathfrak{b})|.$$

Further, since no optimal alignment of two strings of length n contains any aligned pair of gaps, there are at most $2n$ aligned letter pairs, so that the triangular inequality implies

$$|L_n(S) - L_n(T)| \leq 2n \|S - T\|_{\infty}.$$

The claim now follows by dividing by n and taking expectations. Lemma 5.1 now applies because SET has the dual description (1.10), where S can be restricted to the case where f_S is a unit vector, since for all $\tau > 0$ and pairs of strings (x, y) , we have $(\tau S)_{\pi}(x, y) = \tau S_{\pi}(x, y)$, whence $\lambda(\tau S) = \tau \lambda(S)$ and $H(\tau S) = H(S)$.

b) By Part a), Lemma 5.1 is applicable to $C = SET$. This shows that SET is convex compact. It remains to prove that it contains at least one point. Consider the random sequences X, Y introduced in Section 1.1, and for all $n \in \mathbb{N}$ let the first n letters be aligned by π_n , defined as follows,

$$\begin{array}{c|c|c|c|c|c|c|c|c|c} X & X_1 & \mathfrak{G} & X_2 & \mathfrak{G} & X_3 & \dots & X_n & \mathfrak{G} \\ \hline Y & \mathfrak{G} & Y_1 & \mathfrak{G} & Y_2 & \mathfrak{G} & \dots & \mathfrak{G} & Y_n \end{array},$$

where $\mathfrak{G}$ denotes a gap. For all $n \in \mathbb{N}$, the expected empirical distribution of aligned letters is given by the vector $\vec{x}$ with entries defined as follows,

$$x_{\mathfrak{a}, \mathfrak{b}} = \begin{cases} \mathbb{P}[X_1 = \mathfrak{a}], & \text{if } \mathfrak{a} \neq \mathfrak{G}, \mathfrak{b} = \mathfrak{G}, \\ \mathbb{P}[Y_1 = \mathfrak{b}] & \text{if } \mathfrak{a} = \mathfrak{G}, \mathfrak{b} \neq \mathfrak{G}, \\ 0 & \text{otherwise.} \end{cases}$$

For every scoring function S this is a legitimate, albeit suboptimal, alignment. Therefore, we have

$$f_S(\vec{x}) = \frac{1}{n} \mathbb{E}[S_{\pi_n}(X, Y)] \stackrel{(1.4)}{\leq} \frac{1}{n} \mathbb{E}[L_n(S)] \stackrel{(1.6)}{\leq} \lambda(S).$$

By (1.10), this shows that $\vec{x} \in SET$.

c) This follows from Lemma 5.1, which is applicable by Part a).

d) Note that it follows from (1.9) and (1.10) that $f(\vec{x}) \leq \lambda(S)$ for all $\vec{x}$. It suffices thus to show that the hyperplane

$$\mathcal{T} := \{\vec{x} : f_S(\vec{x}) = \lambda(S)\}$$

has nonempty intersection with SET . Assuming the contrary, and using the construction of (5.21) and (5.22), there exists $j \in \mathbb{N}$ such that $\mathcal{T} \cap G_j = \emptyset$. By continuity, there then exists $\delta > 0$ such that

$$(5.29) \quad \mathcal{T}^\delta \cap G_j^\delta = \emptyset,$$

where

$$\begin{aligned}\mathcal{T}^\delta &:= \{\vec{x} : f_S(\vec{x}) \geq \lambda(S) - \delta\}, \\ G_j^\delta &:= \bigcap_{i=1}^j H^\delta(S_i), \\ H^\delta(S_i) &:= \{\vec{x} : f_{S_i}(\vec{x}) \leq \lambda(S_i) + \delta\}.\end{aligned}$$

By (1.14) and the almost sure convergence of $L_n(S_i)/n$ to $\lambda(S_i)$, there exists almost surely $n_0 \in \mathbb{N}$ such that for all $n \geq n_0$,

$$\max_{\vec{x} \in SET^n} f_{S_i}(\vec{x}) \leq \lambda(S_i) + \delta, \quad (i = 1, \dots, j),$$

so that $SET^n \subset G_j^\delta$, and further,

$$\max_{\vec{x} \in SET^n} f_S(\vec{x}) \geq \lambda(S) - \delta,$$

so that $\emptyset \neq \mathcal{T}^\delta \cap SET^n \subset \mathcal{T}^\delta \cap G_j^\delta$. Since this is in direct contradiction with (5.29), the claim holds true. $\square$

6. APPENDIX: LARGE DEVIATIONS

Recall the notations $L_S(x_1 \dots x_i, y_1 \dots y_j)$, $L_n(S) = L_S(X_1 \dots X_n, Y_1 \dots Y_n)$ and $\lambda_n(S) = E[L_n(S)]/n$ introduced in Section 1.1, and the fact that $\lambda_n(S) \rightarrow \lambda(S)$ mentioned in (1.7). In this appendix we will show a stronger result that quantifies the convergence rate as being of order $\mathcal{O}(\sqrt{\ln n/n})$. For this purpose, we introduce the following notation,

$$\begin{aligned}\|S\|_\delta &= \max_{\mathbf{c}, \mathbf{d}, \mathbf{e} \in \mathcal{A}^*} |S(\mathbf{c}, \mathbf{d}) - S(\mathbf{c}, \mathbf{e})|, \\ \|S\|_\infty &= \max_{\mathbf{c}, \mathbf{d} \in \mathcal{A}^*} |S(\mathbf{c}, \mathbf{d})|,\end{aligned}$$

Lemma 6.1. *Let $x = x_1 \dots x_m$ and $y = y_1 \dots y_n$ be two given strings with letters from the alphabet $\mathcal{A}$, and let S be a given scoring function. Let further $\hat{x} \in \mathcal{A}$, and consider two amendments of string x , $x^{[i]} = x_1 \dots x_{i-1} \hat{x} x_{i+1} \dots x_m$, obtained by replacing an arbitrary letter x_i by $\hat{x}$, and $x^{[+]} = x_1 \dots x_m \hat{x}$, obtained by extending x by a letter $\hat{x}$. Then the*

following hold true,

$$(6.1) \quad |L_S(x^{[i]}, y) - L_S(x, y)| \leq \|S\|_\delta,$$

$$(6.2) \quad |L_S(x^{[+]}, y) - L_S(x, y)| \leq \|S\|_\infty.$$

Proof. Let π be an optimal alignment of x and y , so that $S_\pi(x, y) = L_S(x, y)$, and denote the letter with which x_i is aligned under π by $\mathfrak{a} \in \mathcal{A}^*$. Then

$$L_S(x^{[i]}, y) \geq S_\pi(x^{[i]}, y) = S_\pi(x, y) - S(x_i, \mathfrak{a}) + S(\hat{x}, \mathfrak{a}) \geq L_S(x, y) - \|S\|_\delta.$$

Applying the identical argument to an optimal alignment of $x^{[i]}$ and y , we obtain the analogous inequality

$$L_S(x, y) \geq L_S(x^{[i]}, y) - \|S\|_\delta,$$

so that (6.1) follows.

For the second claim, let us use an optimal alignment π of x and y to construct an alignment $\pi^{[+]}$ of $x^{[+]}$ and y by appending an aligned pair of letters $(\hat{x}, \mathfrak{G})$, where $\mathfrak{G}$ denotes a gap. Then we have

$$L_S(x^{[+]}, y) \geq S_{\pi^{[+]}}(x^{[+]}, y) = S_\pi(x, y) + S(\hat{x}, \mathfrak{G}) \geq L_S(x, y) - \|S\|_\infty.$$

Conversely, we can amend an optimal alignment $\tilde{\pi}^{[+]}$ of $x^{[+]}$ and y to become a valid alignment $\tilde{\pi}$ of x and y by cropping the last pair of aligned letters, $(\hat{x}, \mathfrak{a})$. We then have

$$L_S(x, y) \geq S_{\tilde{\pi}}(x, y) = S_{\tilde{\pi}^{[+]}}(x^{[+]}, y) - S(\hat{x}, \mathfrak{a}) \geq L_S(x^{[+]}, y) - \|S\|_\infty,$$

thus establishing (6.2). $\square$

Lemma 6.2. *The convergence of $\lambda_n(S)$ to $\lambda(S)$ is governed by the inequality*

$$(6.3) \quad \lambda_n(S) \leq \lambda(S) \leq \lambda_n(S) + c_n \|S\|_\delta \frac{\sqrt{\ln n}}{\sqrt{n}} + \frac{2\|S\|_\infty}{n}, \quad \forall n \in \mathbb{N},$$

where

$$c_n := \sqrt{\frac{2 \ln 3 + 2 \ln(n+2)}{\ln(n)}}.$$

Note that c_n tends to $\sqrt{2}$ when $n \rightarrow \infty$, so that it effectively acts as a constant.

Proof. See [2]. $\square$

Lemma 6.3 (McDiarmid's Inequality [15]). *Let $Z_1, Z_1, \dots, Z_m$ be i.i.d. random variables that take values in a set D , and let $g : D^m \rightarrow \mathbb{R}$ be a function of m variables with the property that*

$$\max_{i=1, \dots, m} \sup_{z \in D^m, \hat{z}_i \in D} |g(z_1, \dots, z_m) - g(z_1, \dots, \hat{z}_i, \dots, z_m)| \leq C.$$

Thus, changing a single argument of g changes its image by less than a constant C . Then the following bounds hold,

$$\begin{aligned} \mathbb{P}[g(Z_1, \dots, Z_m) - \mathbb{E}[g(Z_1, \dots, Z_m)] \geq \epsilon \times m] &\leq \exp \left\{ -\frac{2\epsilon^2 m}{C^2} \right\}, \\ \mathbb{P}[\mathbb{E}[g(Z_1, \dots, Z_m)] - g(Z_1, \dots, Z_m) \geq \epsilon \times m] &\leq \exp \left\{ -\frac{2\epsilon^2 m}{C^2} \right\}. \end{aligned}$$

Proof. A consequence of the Azuma-Hoeffding Inequality, see [15]. $\square$

Theorem 6.1. *For fixed $\epsilon > 0$ and scoring function S there exists $K > 0$ and $n_\epsilon \in \mathbb{N}$ such that*

$$(6.4) \quad \mathbb{P} \left[\frac{L_n(S)}{n} \geq \lambda(S) + \epsilon \right] \leq e^{-Kn}, \quad \forall n \in \mathbb{N},$$

$$(6.5) \quad \mathbb{P} \left[\frac{L_n(S)}{n} \leq \lambda_n(S) - \epsilon \right] \leq e^{-Kn}, \quad \forall n \in \mathbb{N},$$

$$(6.6) \quad \mathbb{P} \left[\frac{L_n(S)}{n} \leq \lambda(S) - \epsilon \right] \leq e^{-Kn}, \quad \forall n \gg n_\epsilon.$$

Proof. We know from Lemma 6.1 that

$$g(X_1, \dots, X_n, Y_1, \dots, Y_n) = S(X_1 \dots X_n, Y_1 \dots Y_n) = L_n(S)$$

satisfies the assumptions of Lemma 6.3 with $m = 2n$ and $C = \|S\|_\delta$. McDiarmid's Inequality therefore shows

$$\begin{aligned} \mathbb{P} \left[\frac{L_n(S)}{n} \geq \lambda_n(S) + \epsilon \right] &= \mathbb{P} \left[L_n(S) \geq \mathbb{E}[L_n(S)] + \frac{\epsilon}{2} \times 2n \right] \\ (6.7) \quad &\leq \exp \left\{ -\frac{\epsilon^2}{\|S\|_\delta^2} \times n \right\}, \end{aligned}$$

and similarly,

$$(6.8) \quad \mathbb{P} \left[\frac{L_n(S)}{n} \leq \lambda_n(S) - \epsilon \right] \leq \exp \left\{ -\frac{\epsilon^2}{\|S\|_\delta^2} \times n \right\}.$$

Claim (6.5) therefore holds with $K = \frac{\epsilon^2}{4\|S\|_\delta^2}$.

Furthermore, Lemma 6.2 established that

$$(6.9) \quad \lambda_n(S) \leq \lambda(S) \leq \lambda_n(S) + c_n \|S\|_\delta \frac{\sqrt{\ln n}}{\sqrt{n}} + \frac{2\|S\|_\infty}{n}, \quad \forall n \in \mathbb{N},$$

holds, where $c_n = \sqrt{2 \ln 3 + 2 \ln(n+2)} / \sqrt{\ln(n)}$. Using the first inequality from (6.9) in conjunction with (6.7), we find

$$\mathbb{P} \left[\frac{L_n(S)}{n} \geq \lambda(S) + \epsilon \right] \leq \mathbb{P} \left[\frac{L_n(S)}{n} \geq \lambda_n(S) + \epsilon \right] \leq \exp \left\{ -\frac{\epsilon^2}{\|S\|_\delta^2} \times n \right\},$$

which shows that Claim (6.4) holds with $K = \frac{\epsilon^2}{4\|S\|_\delta^2}$.

Using now the second inequality from (6.9) in conjunction with (6.8), we find

$$\begin{aligned} \mathbb{P} \left[\frac{L_n(S)}{n} \leq \lambda(S) - \epsilon \right] &\leq \mathbb{P} \left[\frac{L_n(S)}{n} \leq \lambda_n(S) - \left(\epsilon - c_n \|S\|_\delta \frac{\sqrt{\ln n}}{\sqrt{n}} - \frac{2\|S\|_\infty}{n} \right) \right] \\ &\leq \exp \left\{ -\frac{\left(\epsilon - c_n \|S\|_\delta \frac{\sqrt{\ln n}}{\sqrt{n}} - \frac{2\|S\|_\infty}{n} \right)^2}{\|S\|_\delta^2} \times n \right\} \\ &\leq \exp \left\{ -\frac{\epsilon^2}{4\|S\|_\delta^2} \times n \right\}, \quad \forall n \geq n_\epsilon, \end{aligned}$$

where $n_\epsilon \in \mathbb{N}$ is chosen large enough to satisfy

$$\epsilon - c_n \|S\|_\delta \frac{\sqrt{\ln n}}{\sqrt{n}} - \frac{2\|S\|_\infty}{n} > \frac{\epsilon}{2}, \quad \forall n \geq n_\epsilon.$$

This shows that (6.6) holds for $K = \frac{\epsilon^2}{4\|S\|_\delta^2}$. □

7. ACKNOWLEDGEMENTS

The authors wish to thank the Engineering and Physical Sciences Research Council (EPSRC) and the Institute of Mathematics and its Applications (IMA) and Pembroke College Oxford for generous financial support, Endre Süli for supporting their small grant application to the IMA, Pembroke College for granting a College Associateship to Heinrich Matzinger, and the Oxford Mathematical Institute for hosting him during his Oxford visit.

REFERENCES

- [1] Kenneth S. Alexander. The rate of convergence of the mean length of the longest common subsequence. *Ann. Appl. Probab.*, 4(4):1074–1082, 1994.
- [2] S. Amsalu, R. Hauser, and H. Matzinger. A Monte Carlo approach to the fluctuation problem in optimal alignments of random strings. *University of Oxford, Numerical Analysis Group Research Report NA-12-16*.
- [3] D. Bertsekas and H. Yu. A unifying polyhedral approximation framework for convex optimization. *SIAM J. Optim.*, 21.
- [4] J. Boutet de Monvel. Extensive simulations for longest common subsequences. *Eur. Phys. J. B*, 7:293–308, 1999.
- [5] C. Carathéodory. Über den Variabilitätsbereich der Fourierschen Konstanten von positiven harmonischen Funktionen. *Rendiconti del Circolo Matematico di Palermo*, 32:193–217, 1911.
- [6] V. Chvatal and D. Sankoff. Longest common subsequences of two random sequences. *J. Appl. Probability*, 12:306–315, 1975.
- [7] Felix Hausdorff. *Grundzüge der Mengenlehre*. Veit & Co., 1914.
- [8] R. Hauser, S. Martinez, and H. Matzinger. Large deviation based upper bounds for the lcs-problem. *Advances in Applied Probability*, 38:827–852, 2006.
- [9] R. Hauser, H. Matzinger, and C. Durringer. Approximation to the mean curve in the lcs problem. *Stochastic Processes and their Applications*, 118(1):629–648, 2008.
- [10] E. Hirno, J. Lember, and H. Matzinger. Detecting the homology of DNA-sequences based on the variety of optimal alignments: a case study. *arXiv:1210.3771 [stat.AP]*, 2012.
- [11] M. Jerrum. Lectures in Mathematics ETH Zürich. In *Counting, Sampling and Integrating: Algorithm and Complexity*, pages 62–92. ETHZ, Zürich, 2003.
- [12] H. Kesten, editor. volume 110 of *Encyclopaedia of Mathematical Sciences*, Berlin, 2004. Springer-Verlag.
- [13] M. Krein and D. Milman. On extreme points of regular convex sets. *Studia Mathematica*, 9:133–138, 1940.
- [14] Jyri Lember and Heinrich Matzinger. Standard deviation of the longest common subsequence. *Annals of Probability*, 37(3):1192–1235, 2009.
- [15] C. McDiarmid. On the method of bounded differences. *Surveys in Combinatorics*, 141:148–188, 1989.
- [16] D. Milman. Characteristics of extremal points of regular convex sets (Russian). *Doklady Akademii Nauk SSSR*, 57:119–122, 1947.
- [17] Hermann Minkowski. *Geometrie der Zahlen*. Teubner, 1910.
- [18] H. Rademacher. Über partielle und totale Differenzierbarkeit von Funktionen mehrerer Variablen und über die Transformation der Doppelintegrale. *Math. Ann.*, 79:340–359, 1919.
- [19] Michael J. Steele. An Efron-Stein inequality for non-symmetric statistics. *Annals of Statistics*, 14:753–758, 1986.

- [20] E. Steinitz. Bedingt konvergente Reihen und konvexe Systeme. *J. Reine Angew. Math.*, 143:128–175, 1913.
- [21] Michael S. Waterman. Estimating statistical significance of sequence alignments. *Phil. Trans. R. Soc. Lond. B*, 344:383–390, 1994.

RAPHAEL HAUSER, MATHEMATICAL INSTITUTE, UNIVERSITY OF OXFORD, 24-29 ST GILES', OXFORD OX1 3LB, UNITED KINGDOM

E-mail address: hauser@maths.ox.ac.uk

HEINRICH MATZINGER, SCHOOL OF MATHEMATICS, GEORGIA INSTITUTE OF TECHNOLOGY, 686 CHERRY STREET, ATLANTA, GA 30332-0160 USA

E-mail address: matzi@math.gatech.edu