Colored Convex Linear Orders and Logical Limit Laws

Matthew Kukla

Abstract

We extend previous work on logical limit laws for several classes of ordered structures to the case of structures equipped with a coloring.

1 Introduction

In [1], first-order logical limit laws were proven for convex linear orders by adapting a Markov chain-style proof of Ehrenfeucht. We present a generalization of this argument to the case of convex linear orders equipped with a coloring (henceforth, "colored convex linear orders" or "CCLOs"). These colorings are expressed by expanding the language of convex linear orders to include a countable number of unary predicates, each indicating the color of a point. Every point is assigned a color, and multiple points may have the same color.

Many of the proofs here will follow similar arguments to those in of [1]. We present this note as one set of examples demonstrating how Markov chain arguments may be extended to show limit laws for broader classes of ordered structures.

2 Preliminaries

The language of t-colored convex linear orders, for $t \in \mathbb{N}$, is given by $\mathcal{L}_t = \{<, E, C_1(x), \dots, C_t(x)\}$, where $<$ is a total order on points, E is an equivalence relation whose classes are $<$-intervals, and $C_1(x), \dots, C_t(x)$ are unary predicates (each corresponding to a "color"). A t-colored convex linear order (t-CCLO) is a finite $\mathcal{L}_t$-structure $\mathfrak{M}$ such that, for each point x in $\mathfrak{M}$, there is exactly one $1 \leq i \leq t$ where $C_i(x)$ holds. Stated formally, we require that each $C_i(x)$ satisfies:

$$C_i(x) \iff \neg \bigvee_{\substack{1 \leq \ell \leq t \\ \ell \neq i}}^{t} C_\ell(x)$$

We say that x is an *i-colored point* when $C_i(x)$ holds.

Definition 2.1. *Let $\bullet_i$ denote the CCLO with one class, containing one i-colored point.*

Definition 2.2. *For CCLOs $\mathfrak{M}$, $\mathfrak{N}$, define $\mathfrak{M} \oplus \mathfrak{N}$ to be the CCLO such that $\mathfrak{N}$ comes after $\mathfrak{M}$ with respect to $<$.*

Definition 2.3. *Let $\mathfrak{M}$ be a CCLO. Define $\widehat{\mathfrak{M}}^i$ to be the CCLO obtained by adding one i-colored point to the $<$-last class of $\mathfrak{M}$.*

We will denote the empty CCLO by $\varnothing$. As this structure contains no classes, $\widehat{\varnothing}^i$ is not defined.

Lemma 2.4. *Any t-CCLO of size n can be constructed uniquely, in n steps, by applying $\widehat{(-)}^i$ and $- \oplus \bullet_i$ to $\varnothing$.*

Proof. We follow an inductive argument in the same spirit as Lemma 2.4 of [1]. Let $\mathfrak{N}$ be a CCLO of size n having t colors. If $n = 1$, $\mathfrak{N}$ contains a single point of some color i; this is equivalent to $\varnothing \oplus \bullet_i$.

Assume now that any CCLO of size $n-1$ can be constructed from the above operations. For some CCLO $\mathfrak{N}$ of size n, let $\mathfrak{M}$ denote $\mathfrak{N}$ minus the $<$-last point. If the last class of $\mathfrak{N}$ contains exactly one i-colored point, $\mathfrak{N} \simeq \mathfrak{M} \oplus \bullet_i$. Otherwise, the last point of $\mathfrak{N}$ is obtained as $\widehat{\mathfrak{M}}^i$. $\square$

Example 2.5. *Suppose we are working in the language of CCLOs with two colors, drawn pointwise as* $\circ$ *and* $\bullet$*. Depicting* E*-classes with square brackets, and reading* $<$ *as left-to-right, we can visualize all 2-CCLOs of size* n *as shown in 2.5 for* $n = 0, 1, 2$*.*

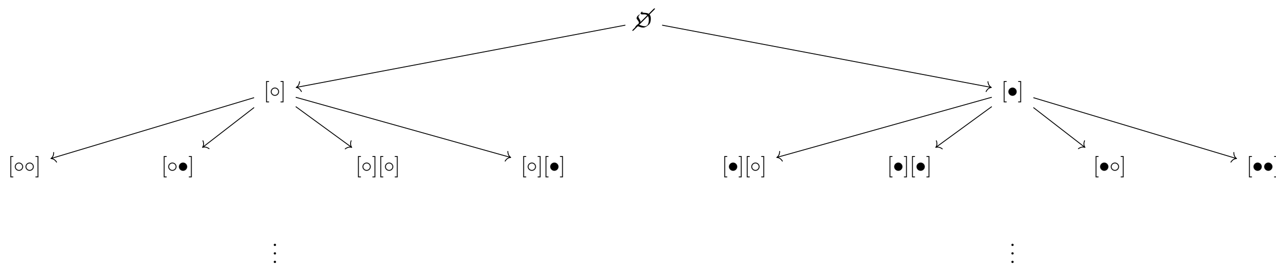

Figure 1: Constructing all 2-CCLOs of size $0, 1, 2$

We write $\mathfrak{M} \equiv_k \mathfrak{N}$ to mean structures $\mathfrak{M}$, $\mathfrak{N}$ agree up to first-order sentences with a maximum quantifier depth of k. This is equivalent to requiring that Duplicator has a winning strategy in a length k Ehrenfeucht–Fraïssé game [2].

Lemma 2.6. *Let* $\mathfrak{M}, \mathfrak{N}, \mathfrak{M}', \mathfrak{N}'$ *be CCLOs with* $\mathfrak{M} \equiv_k \mathfrak{N}$ *and* $\mathfrak{M}' \equiv_k \mathfrak{N}'$*. Then,*

1. $\mathfrak{M} \oplus \mathfrak{M}' \equiv_k \mathfrak{N} \oplus \mathfrak{N}'$
2. *For* $k \in \mathbb{N}$*, there exists* $\ell \in \mathbb{N}$ *such that for all* $s, t > \ell$*,*

$$\bigoplus_s \mathfrak{M} \equiv_k \bigoplus_t \mathfrak{M}$$

Proof. These are Lemmas 2.7 and 2.10 in [1]. $\square$

Lemma 2.7. *Suppose* $\mathfrak{M} \equiv_k \mathfrak{N}$*, then,* $\widehat{\mathfrak{M}}^i \equiv_k \widehat{\mathfrak{N}}^i$*.*

Proof. We construct a winning strategy for Duplicator. If Spoiler plays any point in $\mathfrak{M}$ or $\mathfrak{N}$, Duplicator responds with the corresponding point as they would in a game between $\mathfrak{M}$ and $\mathfrak{N}$ (such a response is guaranteed to exist because $\mathfrak{M} \equiv_k \mathfrak{N}$). In the situation that Spoiler selects the last i-colored point in either structure, Duplicator responds with the corresponding point added in the other structure. $\square$

3 Constructing a Markov chain

Fix a first-order sentence φ in $\mathcal{L}_t$ with quantifier rank k. We associate a Markov chain M_φ to φ in a manner similar to the uncolored case.

For a $\equiv_k$-class $\mathcal{C}$, and any $\mathfrak{M} \in \mathcal{C}$, define

$$\mathcal{C} \oplus \bullet_i := [\mathfrak{M} \oplus \bullet_i]_{\equiv_k}$$

$$\widehat{\mathcal{C}}^i := \left[\widehat{\mathfrak{M}}^i\right]_{\equiv_k}$$

By Lemmas 2.7 and 2.6, any choice of representative $\mathfrak{M}$ will yield a $\equiv_k$-equivalent result. We define M_φ recursively. The starting state is $\varnothing$. There are t possible transitions out of $\varnothing$ to $\bullet_1, \ldots, \bullet_t$, each having probability $1/t$. These initial transitions move only to CCLOs obtained from $- \oplus \bullet_i$ due to the fact that $\widehat{(-)}^i$ is not defined on $\varnothing$. For every $\mathcal{C} \not\simeq \varnothing$, there are $1/2t$

transitions out: one to $\widehat{\mathcal{C}}^i$ and one to $\mathcal{C} \oplus \bullet_i$ (for each $1 \leq i \leq t$). Because any t-CCLO can be constructed uniquely by applying $- \oplus \bullet_i$ and $\widehat{(-)}^i$ to $\varnothing$ repeated n times, this procedure will uniformly randomly sample all t-CCLO structures of size n.

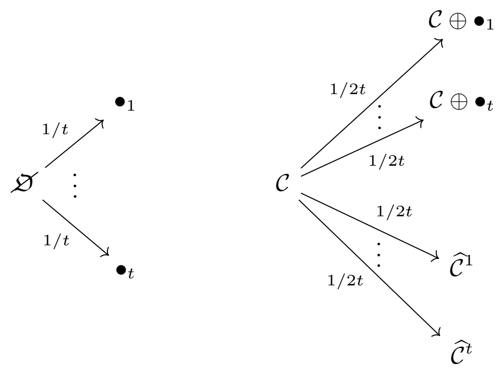

In order for this Markov chain to converge, we require that it is aperiodic in the sense of Definition 2.11 of [1].

Lemma 3.1. *M_φ is aperiodic for all φ.*

Proof. We follow the same argument as Lemma 2.13 of [1]. Suppose M_φ were periodic. Then, there would exist disjoint sets of M_φ-states ($\equiv_k$-classes) $P_0, P_1, \ldots, P_{d-1}$ for some $d > 1$ such that for every state in P_i, M_φ transitions to a state in P_{i+1} with probability 1 (with P_{d-1} transitioning to P_0). Writing $j\bullet_i$ to mean $\bigoplus_j \bullet_i$, we have that for any $\mathcal{C} \in P_0$, $\mathcal{C} \oplus j\bullet_i$ is in P_0 iff $d \mid j$. From Lemmas 2.6 and 2.6, $\mathcal{C} \oplus j\bullet_i \equiv_k \mathcal{C} \oplus (j+1)\bullet_i$ for sufficiently large j, contradicting this. □

Theorem 3.2. *The class of t-CCLOs admits a logical limit law for all $t \in \mathbb{N}$.*

Proof. Consider M_φ for some fixed φ. In any state of M_φ (a $\equiv_k$-class) S, either every structure in S satisfies φ or no structures in S satisfy φ. By the definitions of $- \oplus \bullet_i$ and $\widehat{(-)}^i$ for $\equiv_k$-classes, moving n steps in M_φ (starting from $\varnothing$) is equivalent to uniformly randomly selecting a CCLO of size n and taking its $\equiv_k$-class. Hence, the probability of M_φ being in a state which satisfies φ after n steps is equal to the probability that a randomly selected CCLO of size n satisfies φ. It is sufficient to show that the probability of M_φ being in a satisfactory state after n steps converges as $n \to \infty$; this follows from the fact that M_φ is finite and aperiodic. □

4 Reduction to the uncolored case

We briefly note that limit laws for uncolored convex linear orders can be obtained as a special case of 3.2. An uncolored structure may be equivalently viewed as a colored structure with exactly one color. Hence, the relation $C_1(x)$ holds for every point x, so that there is no distinction in terms of color on the points.

We have two operations for building such structures: $\widehat{(-)}^1$ and $- \oplus \bullet_1$. These are equivalent to the corresponding operators $\widehat{(-)}$ and $- \oplus \bullet$ in Definition 2.2 and Lemma 2.4 respectively of [1]. Because there is only one color, the subscripts are dropped hereafter. Following the procedure in 3, we construct M_φ for first-order sentence φ as:

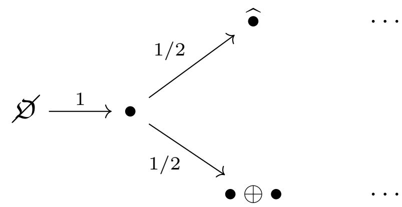

The initial transition has probability 1, as there is only one way to construct $\bullet$ from the empty order. From this diagram, it can be seen that moving n steps in M_φ is equivalent to moving $n-1$ steps in the Markov chain defined by [1], due to the fact that the latter is defined starting at $\bullet$ rather than $\varnothing$. The two Markov chains will converge to the same limiting probability as $n \to \infty$.

References

[1] Samuel Braunfeld and Matthew Kukla. Logical Limit Laws for Layered Permutations and Related Structures. *Enumerative Combinatorics and Applications*, 2(S4PP2), 2021.

[2] Joel Spencer. *The Strange Logic of Random Graphs*, volume 22. Springer Science & Business Media, 2001.