Lectures discuss one idea, reinvented often under different names, which I call *local weak convergence* (LWC).

Setting: given a random finite structure, which either includes an n-vertex graph or to which we can naturally attach an n-vertex graph, want to study $n \to \infty$ asymptotics. Do this by considering distributions of structure in neighborhood of a random vertex.

One purpose: sometimes one can get limit behavior of "global" statistics out of this "local" limit.

My own main focus has been on

- combinatorial optimization over random data
- in particular, within "mean-field model of distance".

(lectures 2 and 3) but other areas include

- (pure probabilistic combinatorics): fringe subtrees of random trees (lecture 1); random quadrangulations etc

- deterministic approximate counting: of spanning trees (Lyons); of independent sets and colorings (Bandyopadhyay - Gamarnik)

- involution invariance as "stationarity" for countable random graphs; Benjamini-Schramm and Aldous-Lyons

- in modeling complex networks (lecture 3)

Lecture 1 gives set-up and basic examples – very easy, no serious theorems, just "a way of looking at things".

The serious theorems require different ad hoc technical methods; LWC is just a starting place.

Stuff you already know about weak convergence

Consider an abstract space S (complete separable metric space) with a notion of convergence $x_n \to x$. There is a notion of convergence of probability measures on S which respects the topology: all reasonable definitions are equivalent and the most intuitive is

$\mu_n \to \mu_\infty$ iff there exist S-valued random variables X'_n such that

$$\text{dist}(X'_n) = \mu_n; \quad P(X'_n \to X'_\infty) = 1.$$

This is called **weak convergence**; but rather than naming distributions explicitly we typically write $X_n \stackrel{d}{\to} X$ to mean $\text{dist}(X_n) \to \text{dist}(X)$ and call it **convergence in distribution**.

Conceptual point: In our context of a random finite structure, we get to **choose** how to represent it as a random element of some abstract space S of our choice.

Having done that, we don't need to think about what convergence of distributions means.

Digression on definitions: consider Chung's definition of a random variable.

A real, extended-valued random variable is a function X whose domain is a set Δ in $\mathcal{F}$ and whose range is contained in $\mathbb{R}^* = [-\infty, \infty]$ such that for each B in $\mathcal{B}^*$ we have

$$\{\omega : X(w) \in B\} \in \Delta \cap \mathcal{F}$$

where $\Delta \cap \mathcal{F}$ is the trace of $\mathcal{F}$ on Δ.

Easy to poke fun but illustrates a genuine issue – do you want to cover every possible variant in an initial definition?

Stuff you already know about graphs.

A **graph** G has vertices v and (undirected) edges e. degree(v) = number of edges at v.

A **root** is a distinguished vertex (for now, assume other vertices unlabeled).

Distance $d(v, w)$ is (as a default) number of edges on shortest path.

Can define a subgraph Ball($G; r$) on the vertices at distance $\leq r$ from the root. Say G is **locally finite** if each Ball($G; r$) is finite (i.e. finite number of vertices). If G is connected then "locally finite" equivalent to "each v has finite degree".

Stuff you probably haven't thought about

We can define an abstract space

$$S = \{\text{locally finite rooted graphs}\}$$

after identifying isomorphic ones. This has a natural topology:

$G_n \to G_\infty$ means that for each fixed r, for $n > n_0(r)$ there is an isomorphism between $\mathrm{Ball}(G_n; r)$ and $\mathrm{Ball}(G_\infty; r)$.

This space S is nice enough; so we automatically have a notion of convergence in distribution for random locally finite rooted graphs.

Note: typically G_n finite, G_∞ infinite.
Analogous to convergence of infinite sequences of integers.

A complication; often we will deal with a **network**, that is a graph with extra structure, typically marks or numbers attached to edges or vertices. Hard (cf. Chung) to choose a level of generality in which to write down a definition.

Extend previous definition

$G_n \to G_\infty$ means that for each fixed r, for $n > n_0(r)$ there is an isomorphism between Ball$(G_n; r)$ and Ball$(G_\infty; r)$.

by requiring that marks converge too (under isomorphism).

But a special rule comes into play when edge-marks are **lengths**. Then distance $d(v, w)$ is shortest route-length and this distance is used in definition of Ball$(G; r)$ and hence in the meaning of "locally finite" and the topology of convergence of locally finite rooted networks. Call this the "continuum setting" in contrast to "graph setting".

Given n-vertex network G_n (deterministic or random) let U_n be uniform random vertex. Write $G_n[U_n]$ for G_n rooted at U_n.

Definition. If $G_n[U_n] \stackrel{d}{\to}$ some G_∞, call this **local weak convergence (LWC)** of G_n to G_∞ and write $G_n \to_{LWC} G_\infty$.

Formalizes the idea: for large n the local structure of G_n near a typical vertex is approximately the local structure of G_∞ near the root.

Note odd syntax; convergence of finite unrooted networks to an infinite rooted network

Intuition: in models where degree(U_n) is tight as $n \to \infty$ we expect LWC to some limit infinite network. More precisely, in the "plain graph" setting the condition

for each r the size of Ball($G_n[U_n]; r$) is tight

is the condition for compactness, i.e. for some convergent subsequence.

The rest of Lecture 1 is playing around with this definition.

Let's start with a simple deterministic example; consider the discrete d-dimensional cube graph of side-length m:

$$C_m^d = [0, 1, \ldots, m-1]^d; \quad n = m^d.$$

Clearly as $m \to \infty$ we have

$$C_m^d \to_{LWC} \mathbb{Z}^d \text{ rooted at 0.}$$

Now make C_m^d into a random network by attaching IID marks to edges; clearly we have LWC to $\mathbb{Z}^d$ with IID marks on edges.

What happens if the edge-marks are random but not IID? Let's think about the $d = 1$ case and forget graphs for a moment.

For each n let $(X_{n,1}, \dots X_{n,n})$ be arbitrary $\mathbb{R}$-valued. Take $U_n \stackrel{d}{=}$ Uniform$[1, n]$ and suppose

$$(X_{n,U_n},\ n \geq 1) \text{ is tight .}$$

Then there is a subsequence in which

$$\begin{aligned}(Y_{n,i}, -\infty < i < \infty) &:= (X_{n,U_n+i}, -\infty < i < \infty) \\ &\stackrel{d}{\to} (Y_i, -\infty < i < \infty)\end{aligned}$$

where the limit process is stationary; moreover every stationary process arises this way.

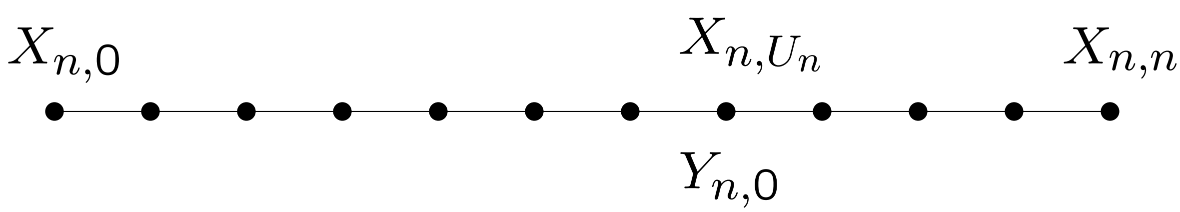

Could view example as random network over line graph.

Conceptual point: In LWC the possible limits (random rooted infinite networks) are the network analogs of **stationary processes**.

Note also that LWC provides a sense in which sampling without replacement converges to sampling with replacement.

Returning to the discrete cube graph

$$C_m^d = [0, 1, \ldots, m-1]^d; \quad n = m^d$$

take IID **edge-lengths** (L_e). Provided L_e strictly positive, we do have LWC in the "continuum setting" to the obvious limit – $\mathbb{Z}^d$ with IID edge-lengths.

Consider a point process model: intuitively, n random (independent, uniform) points in square of **area** n "converges" to Poisson point process of rate 1 on $\mathbb{R}^2$.

One way to state this: from the n points $(\xi_{n,i}, 1 \leq i \leq n)$ pick a random point ξ_{n,U_n} and look at displacements

$$(\xi_{n,i} - \xi_{n,U_n}, i \leq i \leq n).$$

These converge (usual sense of weak convergence of point processes) to Poisson point process on $\mathbb{R}^2$ with a point planted at 0.

Equivalently if we take the complete graph on the n points $(\xi_{n,i}, i \leq i \leq n)$ with edge-lengths = Euclidean lengths, then this random network G_n converges (LWC in continuum setting) to the complete graph on the Poisson point process on $\mathbb{R}^2$.

This example indicates why the continuum setting is useful; can be applied even when limit graph has infinite degrees.

Similarly, fix $c > 0$ and consider the **random geometric graph** $G_{n,c}$ on the n points $(\xi_{n,i}, 1 \leq i \leq n)$ where [definition] $G_{n,c}$ contains only edges of length $\leq c$. Then

$$G_{n,c} \to_{LWC} G_{\infty,c}$$

the limit being the random geometric graph on the limit Poisson process.

In practice thinking in terms of LWC doesn't add anything to these examples, but . . . . . .

Conceptual point: convergence of $\mathbb{Z}^d$-indexed processes and convergence of point processes on $\mathbb{R}^d$ can often be viewed as special cases of LWC.

Note the word **local** in LWC is intended to contrast with **global** weak convergence, exemplified by convergence of random walk to Brownian motion, in which the entire finite random structure is rescaled.

Roughly speaking, every bounded-mean-degree sequence of deterministic or random graphs has some local weak limit. Here are 2 more deterministic examples (do on blackboard).

Balanced finite binary tree.

de Bruijn graph.

The (implicitly) classic example concerns the sparse Erdos-Renyi random graph $\mathcal{G}(n, c/n)$.

Recall this has n vertices, and each possible edge is present independently with chance c/n, for fixed $0 < c < \infty$. We have

$$\mathcal{G}(n, c/n) \to_{LWC} \mathsf{PGW}(c)$$

where the limit is the Galton-Watson branching process (viewed as a rooted tree) with Poisson(c) offspring.

(PGW stands for Poisson-Galton-Watson).

Why is this true? – outline on blackboard.

Digression; every freshman probability/statistics course should include an example like "family size distributions":

if $(p(i), i \geq 0)$ is distribution of "number i of children per family", then the distribution $(\tilde{p}(j), j \geq 0)$ of "number j of siblings of random child" is

$$\tilde{p}(j) \propto (j+1)p(j+1).$$

. . . a basic instance of size-biasing; cf. internet search session lengths.

There are a variety of models $G(n, \mathbf{q})$ on n vertices which formalize the notion "random subject to degree distribution approximately the prescribed distribution $\mathbf{q} = (q(i), i \geq 0)$.

All such models have the property

$$G(n, \mathbf{q}) \to_{LWC} \mathsf{GW}(\mathbf{q}, \tilde{\mathbf{q}})$$

where the limit is the Galton-Watson branching process with $\mathbf{q}$ offspring in the first generation and $\tilde{\mathbf{q}}$ offspring in subsequent generations.

In particular, for a random r-regular graph the LWC limit is the (deterministic) infinite r-regular rooted tree.

An example to be used in Lecture 2.

Cayley's formula says that the number of trees on n labeled vertices equals n^{n-2} (unrooted) or n^{n-1} (rooted). For a uniform random such tree G_n we can write down many explicit formulas "just by counting". For instance, the chance we see

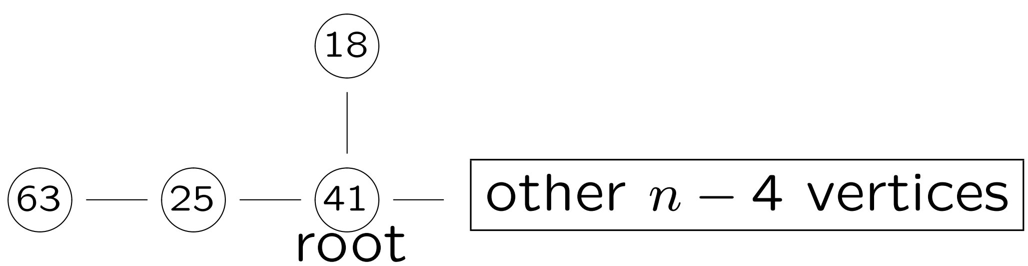

equals $\frac{(n-4)^{(n-4)-1}}{n^{n-1}}$. Removing labels, the chance we see

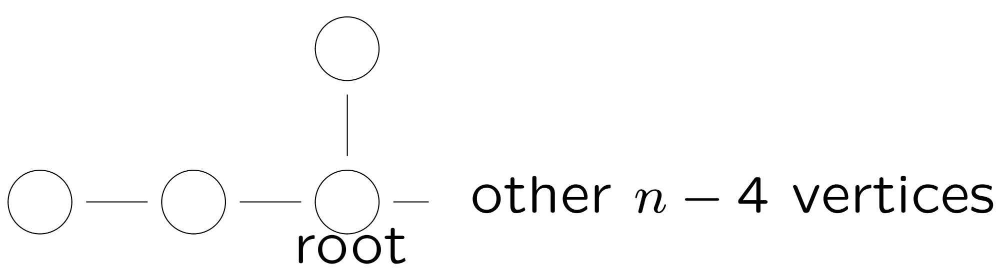

equals $\frac{(n-4)^{(n-4)-1}}{n^{n-1}} \times (n)_4 \sim e^{-4}$.

Now the chance the PGW(1), drawn in an unusual way, is

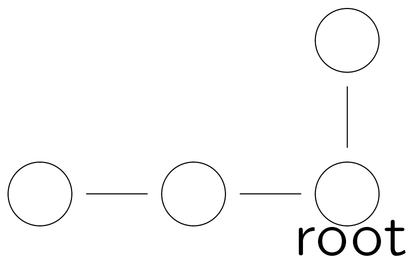

equals $e^{-1}/2 \times e^{-1} \times 1 \cdot e^{-1} \times e^{-1} \times 2 = e^{-4}$, the final $\times 2$ because either first-generation offspring could have the child.

Repeating the argument with an arbitrary finite rooted tree shows the following. Let $G_n =$ uniform random tree on n labeled vertices (then delete labels). Let U_n be uniform random vertex. Delete largest component attached to U_n, and write $G_n^{small}[U_n]$ for remaining rooted tree. Then

$$G_n^{small}[U_n] \xrightarrow{d} \text{PGW}(1).$$

This idea goes back to Grimmett (1980).

This isn't quite LWC, but by continuing these combinatorial arguments, or by a general result mentioned below,

one can show

$$G_n \to_{LWC} \mathrm{PGW}^{\infty}(1)$$

where the limit is as follows. Take one-sided infinite path from root; attach i.i.d. "bushes" which are independent PGW(1).

root 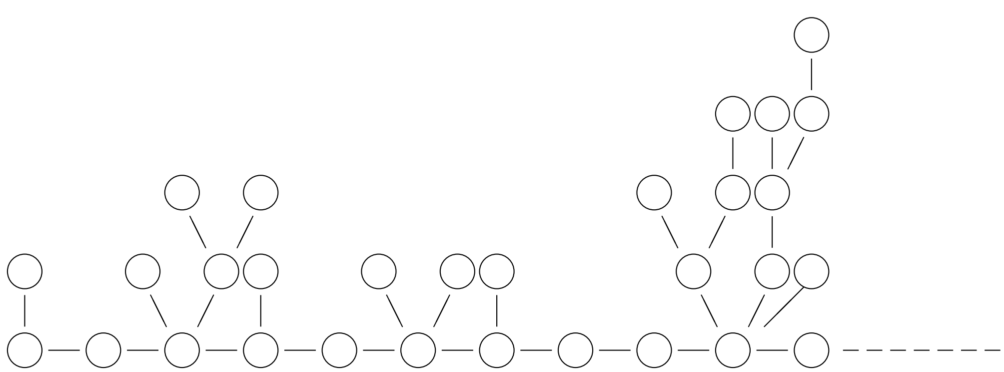

Though **not** the most interesting setting, it turns out (Aldous 1991) there's a "general theory" of LWC when the given graphs G_n are (rooted) trees, and this parallels the example above. Each vertex v of a finite rooted tree defines a subtree rooted at v. Take some model of random n-vertex rooted tree G_n; pick $v = U_n$ uniformly at random to get the <u>random fringe subtree</u> $\mathcal{F}_n$.

Empirical Observation. For most "natural" families of random trees,

$$\mathcal{F}_n \xrightarrow{d} \mathcal{F} \text{ (say), as } n \to \infty$$

where the limit is a finite random rooted tree but with infinite mean number of vertices.

Note (blackboard) the path and the star are extreme "bad" examples. To illustrate what's going on, consider the Crump-Mode-Jagers general continuous-time branching process, i.e.

The Pessimist's View of Life

1. You're born; you have a random number of children at random times; you die.

2. Your children behave in the same way, independently of you.

<u>Model.</u> Continuous-time BP where each individual has C children ($EC > 1$) at times $(\xi_1, \xi_2, \ldots, \xi_C)$ (arbitrary distribution) after own birth.

Standard facts. Under minor technical assumptions, conditional on non-extinction:

1. (Number born before t) $:= N(t) \sim Ze^{\theta t}$ for a certain constant θ.

2. Pick individual at random from those born before deterministic time T; look at individual and descendants born before T. As $T \to \infty$ this "random family tree" $\mathcal{F}_T$ has the following limit $\mathcal{F}$:

Start the BP with 1 individual and watch for an Exponential(θ) time.

Note the "infinite mean" size of $\mathcal{F}$ arises as $\int EN(t) \cdot \theta e^{-\theta t} \, dt$.

Point. Many of the tractable models of combinatorial random trees are tractable precisely because they are similar to critical or supercritical branching process models.

Example: greedy undirected tree.

n vertices, one distinguished (root). Start with no edges. Repeat $n-1$ times

add edge, chosen at random (uniformly) from set of all edges whose addition would not create a cycle

to get random tree T_n.

Fact. random fringe subtrees $\mathcal{F}_n \stackrel{d}{\rightarrow} \mathcal{F}$ where $\mathcal{F}$ is family tree of the following multitype BP.

Type space $(0, \infty)$. Type s individual has:
Poisson($\lambda(s)$) offspring of type s
Poisson (rate $\rho(s^*)$) process of offspring of types $s^* < s$
Progenitor type has density $\nu(s)$.

(Explicit formulas for ν, λ, ρ omitted).

So asymptotic proportion of leaves in T_n equals chance progenitor in $\mathcal{F}$ has no offspring

$$\int_0^\infty \exp\left(-\lambda(s) - \int_0^s \rho(s^*) ds^*\right) \nu(ds) \approx 0.408.$$

Recall setting: underlying network G_n is itself a random tree, so can define random fringe subtree $\mathcal{F}_n$.

A (perhaps) surprising **Theorem** is that convergence of random fringe subtrees to some limit $\mathcal{F}_n \stackrel{d}{\rightarrow} \mathcal{F}_\infty$ implies the (*a priori* stronger) LWC of G_n to a limit $\mathcal{T}_\infty$ determined by $\mathcal{F}_\infty$. The limit always has the same qualitative structure as in previous two examples:

semi-infinite path with finite bushes attached to baseline. Bush at root is $\mathcal{F}_\infty$.