Logic and Probability, EPFL, Fall 2011
Lecture Notes 3
Val Tannen

10 Ehrenfeucht-Fraïssé Games and the First Proof of the 0-1 Law

We consider FO models over a signature (vocabulary) σ which has *only relational and constant symbols*. If $\mathcal{A}$ is a σ-structure and $s \in \sigma$ is a constant or relation symbol we denote by $s^{\mathcal{A}}$ the interpretation of s in $\mathcal{A}$.

The Ehrenfeucht-Fraïssé game is as follows. There are two players, called Spoiler and the Duplicator. The board of the game consists of two structures $\mathcal{A}$ and $\mathcal{B}$. The goal of Spoiler is to show that these two structures are different; the goal of Duplicator is to show that they are the same.

In the classical Ehrenfeucht-Fraïssé game, the players play a certain number of rounds. Each round consists of the following steps:

1. Spoiler picks a structure ($\mathcal{A}$ or $\mathcal{B}$) and makes a move by picking an element of that structure: either $a \in \mathcal{A}$ or $b \in \mathcal{B}$.
2. Duplicator responds by picking an element in the other structure.

Suppose that Spoiler and Duplicator play n rounds and let $\overrightarrow{a} = (a_1, \ldots, a_n)$ and $\overrightarrow{b} = (b_1, \ldots, b_n)$ be the (not necessarily distinct!) moves made by the players on $\mathcal{A}$, respectively $\mathcal{B}$. Who has won? To define this, we need a crucial definition: that of a *partial isomorphism.*

Definition 10.1 (Partial isomorphism). *Let $\mathcal{A}$, $\mathcal{B}$ be two σ-structures, and $\overrightarrow{a} = (a_1, \ldots, a_n)$ and $\overrightarrow{b} = (b_1, \ldots, b_n)$ be two tuples of elements from $\mathcal{A}$ and $\mathcal{B}$ respectively. Then $(\overrightarrow{a}, \overrightarrow{b})$ defines a* partial isomorphism *between $\mathcal{A}$ and $\mathcal{B}$ if the following conditions hold:*

- *For every $i, j \leq n$,*
$$a_i = a_j \text{ iff } b_i = b_j.$$
- *For every constant symbol c from σ, and every $i \leq n$,*
$$a_i = c^{\mathcal{A}} \text{ iff } b_i = c^{\mathcal{B}}.$$
- *For every m-ary relation symbol R from σ and every sequence $(i_1, \ldots, i_m)$ of (not necessarily distinct) numbers $1 \leq i_1, \ldots, i_m \leq n$,*
$$(a_{i_1}, \ldots, a_{i_m}) \in R^{\mathcal{A}} \text{ iff } (b_{i_1}, \ldots, b_{i_m}) \in R^{\mathcal{B}}.$$

Definition 10.2 (Who wins) *The game run $(\overrightarrow{a}, \overrightarrow{b})$ has been won by Duplicator if it defines a partial isomorphism. Otherwise, this game run was won by Spoiler.*

Definition 10.3 *We write $\mathcal{A} \sim_n \mathcal{B}$ if Duplicator has a winning strategy for $\mathcal{A}$ and $\mathcal{B}$ that works in any n-round game.*

Observe that $\mathcal{A} \sim_n \mathcal{B}$ implies $\mathcal{A} \sim_k \mathcal{B}$ for every $k \leq n$.

Although this is not at all obvious, $\sim_n$ is an equivalence relation on structures. This can be shown directly, or it can be seen to follow from Theorem 10.1 below.

Definition 10.4 (Quantifier Rank). *The quantifier rank of a formula $qr(\varphi)$ is its depth of quantifier nesting. That is:*

- *If φ is atomic, then $qr(\varphi) = 0$.*
- $qr(\varphi_1 \vee \varphi_2) = qr(\varphi_1 \wedge \varphi_2) = max(qr(\varphi_1), qr(\varphi_2))$.
- $qr(\neg\varphi) = qr(\varphi)$.
- $qr(\exists x\varphi) = qr(\forall x\varphi) = qr(\varphi) + 1$.

Theorem 10.1 (Ehrenfeucht-Fraïssé). *Let $\mathcal{A}$ and $\mathcal{B}$ be two structures in a relational/constants vocabulary σ. Then the following are equivalent:*

1. *For any φ, $qr(\varphi) \leq k$, $\mathcal{A} \models \varphi$ iff $\mathcal{B} \models \varphi$.*
2. $\mathcal{A} \sim_k \mathcal{B}$.

For the proof of this fundamental theorem you can consult Libkin's book (see course bibliography) or Kolaitis's chapter in "Finite Model Theory and Its Applications", Grädel et al., eds., Springer 2007.

Lemma 10.2 *For any k, G_1, G_2 (finite or not!) if G_1, $G_2 \models EA_{r,s}$ for all $r + s \leq k$, then $G_1 \sim_k G_2$. In other words, if G_1, G_2 satisfy enough extension axioms then Duplicator has a winning strategy for any Ehrenfeuch-Fraisse game played on G_1 and G_2.*

Proof. Duplicator's strategy is given quite directly by the extension axioms. Indeed, if at round $n+1$ Spoiler plays (say) $u \in G_1$ then let A_1 be the subset of the vertices already played on G_1 which are adjacent to u and let N_1 be the set of the other vertices already played, those not adjacent to u. Let Y_2 and N_2 be the sets of moves played on G_2 corresponding to Y_1 and N_1. Duplicator has a move to answer $u \in G_1$ by Spoiler because $G_2 \models EA_{r,s}$ where $r = |Y_2|$ and $s = |N_2|$ with $r + s \leq n$. □

Now we prove the 0-1 Law for FOL.

Theorem 10.3 (0-1 Law) *Suppose $p(n) = p \neq 0, 1$ is constant (does not depend on n) i.e., assume the case of the random graph of constant probability. Then, for any φ in FOL $\mu(\varphi)$ exists and is either 0 or 1.*

Proof. Fix an arbitrary φ and let $k = qr(\varphi)$. Consider

$$EA_k = \bigwedge_{r+s \leq k} EA_{r,s}$$

It follows from Lemmas 8.1 and 7.1 that $\mu(EA_k) = 1$.

Now we consider the following two cases:

Case 1 $EA_k \wedge \varphi$ is satisfiable.

Let the graph G_0 be a model for this sentence. Then, for any graph G such that $G \models EA_k$, we have $G_0 \sim_k G$ by Lemma 10.2. Then $G \models \varphi$ by Theorem **??**. Therefore EA_k implies φ and by Lemma 7.1(d) $\mu(\varphi)$ exists and is 1.

Case 2 $EA_k \wedge \varphi$ is unsatisfiable.

Then any graph G that satisfies EA_k cannot also satisfy φ and must therefore satisfy $\neg varphi$. Therefore EA_k implies $\neg\varphi$, and by Lemma 7.1(d) $\mu(\neg\varphi)$ exists and is 1. By Lemma 7.1(a) this means that $\mu(\varphi)$ exists and is 0.

□

In the next lecture we will provide another proof of the 0-1 law, a proof that will allow us to also derive the nice additional fact that given φ it is *decidable* which one of $\mu(\varphi) = 1$ or $\mu(\varphi) = 0$ holds. (But probably not *efficiently* decidable: it was also shown that the problem is in fact PSPACE-complete.)