Rory Soiffer

Math 336

18 May 2017

Friendly Frogs, Stable Marriage, and the Magic of Invariance

Game Theory Terminology

The paper talks heavily about game theory, so it's important to introduce some general game theory terminology.

A **game** has a current **position** and one or more **players**. Some positions are **terminal positions**, and are labeled as a victory for one of the players. The game immediately ends upon encountering a terminal position. A game is always ends in a victory for exactly one player. A game has a current player who turn it is to move, and players usually alternate taking turns. The player making a move has a list of moves they are allowed to make, where each move transforms the position of the game in a certain way. All games discussed in this paper are **deterministic** and have **complete information**, meaning that there is no element of random chance, and no information about the game's position is hidden from the players.

A **strategy** is a map from game positions to moves. A **winning strategy** is a strategy that guarantees that the player following it will win the game in a finite number of moves, no matter what the other player does. A game position is a **win** for a player if there exists a winning strategy for that player, a game position is a **loss** for a player if there exists a winning strategy for another player. If a game always ends after a finite number of moves, then every position is either a win or a loss for a given player.

A position is a **P-position** if it is a win for the player that moved previously. A position is an **N-position** if it is a win for the player whose turn it is to move now. In a game with two players who alternate taking turns, every position is either a P-position or an N-position.

Stable Matching Terminology

The paper also pulls in ideas from computer science – specifically the Gale-Shapley algorithm for the stable marriage problem. A full discussion of the importance of the Gale-Shapley algorithm is beyond the scope of this paper, but a few key terms from it are defined below.

A **matching** of a set of points L is a set M of unordered pairs of distinct points in L such that each point of L is included in at most one pair. We say $x \in L$ is **matched** if x is included in some pair, and x is **unmatched** otherwise. The matching is **perfect** if each point in L is matched.

We can talk about additional properties of matchings if we can define some preference function for every point x that represents how much x desires to be matched to every other point. Every point x prefers to be matched with any other point y than with no point. Given a matching M of L, a pair of points $x, y \in L$ is called **unstable** if x prefers y to its current partner (if any) and y prefers x to its current partner (if any). A matching is called **stable** if there are no unstable pairs. In the context of this paper, for $x, y, z \in \mathbb{R}^d$, we choose for the preference function to be such that x prefers y to z if and only if $|x - y| < |x - z|$.

Gale and Shapley introduced the concept of stable matchings in the context of n heterosexual marriages between n woman and n men, each of whom has an arbitrary preference order for each person of the opposite sex. Gale and Shapley proved that a stable matching always exists in this case and gave an efficient algorithm to find one.

Infinite Point Set Terminology

The authors state several theorems about games that involve countably infinite sets of points. As such, I introduce some terminology for dealing with infinite point sets.

Let L be an infinite subset of $\mathbb{R}^d$, and assume that all distances between pairs of points in L are distinct. We call a sequence of points $\{x_i\}$ a **descending chain** if the distances $|x_{i+1} - x_i|$ form a strictly decreasing sequence. L is **discrete** if any bounded subset of L contains finitely many points. Equivalently, L is discrete if and only if every point in L is isolated. It is a theorem (though it is not proved here) that if L has no descending chains, then L is discrete.

To formally define the notion of a random infinite set, we use a **Poisson process**. A Poisson process can be defined as a limit as $\varepsilon \to 0$ of a grid of d-dimensional cubes of volume ε, each of which contains a point with independent probability $\varepsilon\lambda$ for some fixed positive real number λ, called the **intensity**. It is a theorem (though it is not proved here) that in a Poisson process, there is probability 1 that all pairs of points have distinct distances and that there are no descending chains. A Poisson process is **translation-invariant**, which means that the distribution of points is invariant under any translation of $\mathbb{R}^d$.

Friendly Frogs

As the focus of their paper, the authors introduce a new game, which they call **friendly frogs**. Friendly frogs has two players (who we call Alice and Bob), who alternate taking turns (Alice always moves first). The game position is a set of points **L** in Euclidean space $\mathbb{R}^d$, and two frogs that occupy distinct points in that set. The game starts with frogs not occupying any points in the set. As her first move, Alice must choose a point for the first frog to occupy, and then as his first move, Bob must choose a point for the second frog to occupy. On their turn, a player must move either frog from its original point to a new point, such that the distance between the two frogs strictly decreases. If a player has no legal moves, they lose the game (and the other player wins the game).

The authors proceed to show that friendly frogs has several interesting properties. They show that in most circumstances, if L is finite, the winning player only depends on the size of L. They also show that if L is infinite and randomly distributed, then Bob wins with probability 1.

Theorems

The following theorems are copied exactly as stated in the source paper. The proofs take their core ideas from the source paper, though the presentation is entirely my own. In particular, I go into much greater detail than the source paper, formalizing all of the proofs. For example, lemmas marked with a * were used in the source paper but were never actually proved.

Theorem 1. *Consider friendly frogs played on a finite set $L \subset \mathbb{R}^d$ of size n in which all pairs of points have distinct distances. The game is a win for Alice if n is odd, and a win for Bob if n is even.*

We let **M** be a matching of points in L, where at most one point remains unmatched (and only if n is odd). We say that a position is in M if and only if the frogs are at the points x and y, where x and y are matched in M.

We construct M by ordering all pairs in L in increasing order of distance. Then for each pair, starting with the pair with the smallest distance, we match the two points to each other if and only if neither point is already matched. This algorithm clearly terminates and matches all points in M (except one leftover point if n is odd).

***Lemma 1.1a.** Assume x is matched to x’ and y is matched to y’, and x is closer to x’ than y is to y’. Then x is closer to x’ than x is to y.

Proof. When M was being constructed, x and x' were matched before y and y' were matched, as the pairs were matched in increasing order of distance. Therefore, immediately before x was matched with x', all four of x, x', y, and y' were unmatched. This means that x was chosen to match with x' before x was considered for matching with y, so x must be closer to x' than to y.

***Lemma 1.1b.** Assume x is matched to x' and y is not matched to any point. Then x is closer to x' than x is to y.

Proof. When M was being constructed, immediately before x was matched with x', all three of x, x', and y were unmatched. This means that x was chosen to match with x' before x was chosen to match with y, so x must be closer to x' than to y.

Lemma 1.2. If the current position is not in M, then the current player can move so that the next position is in M.

Proof. Let the frogs be at the points x and y, where x and y are not matched in M. If both x and y are matched, then let x be matched to x' and y be matched to y', and without loss of generality, assume x is closer to x' than y is to y'. By Lemma 1.1a, x is closer to x' than x is to y. If either x or y is unmatched, then without loss of generality, assume y is unmatched. By Lemma 1.1b, x is closer to x' than x is to y. Therefore, in either case, moving the frog from y to x' decreases the distance between it and the frog at x, so the current player can move the frog from y to x', and the frogs are made to be at matched points.

Lemma 1.3. If the current position is in M, then any move by the current player results in the next position not being in M.

Proof. This lemma is trivial, as every point is matched to only one other point, so moving the location of either frog makes it so that the frogs are no longer at matching points.

***Lemma 1.4.** If the current position is not in M, then the current position is an N-position. Moreover, the current player has a winning strategy as follows: always move so that the next position is in M.

Proof. We call the current player A, and we call the previous player B. Lemma 1.2 guarantees that if the current position is in M, A can always follow the given strategy. After A moves, the position is in M and it's B's turn to move, so Lemma 1.3 guarantees that the next position is not in M. Then, the position is not in M and it's again A's turn to move, so A can continue to follow the given strategy. Note that Lemma 1.2 guarantees that A will always have a legal move, so A cannot lose the game. B is not guaranteed to always have a move, so B will eventually lose the game.

***Lemma 1.5.** If the current position is in M, then the current position is a P-position.

Proof. We call the current player A, and we call the previous player B. By Lemma 1.3, A must move so that the next position is not in M. Then, by Lemma 1.4, the current position is not in M, so it is a win for B.

We must now consider the starting two moves of the game. If n is even, then we show that the game is a win for Bob. As her first move, Alice must choose a point x at which to place the first frog. Bob can then choose to place the second frog at the point matched to x. Then the current position is in M, so by Lemma 1.5, it is a loss for the current player (Alice).

If n is odd, then we show that the game is a win for Alice. As her first move, Alice should choose to place her frog at the single unmatched point x. Bob must then choose to place the second frog at a point not matched to x. Then the current position is not in M, so by Lemma 1.4, it is a win for the current player (Alice).

Theorem 2. *Suppose $L \subset \mathbb{R}^d$ has all pairwise distances distinct and has no infinite descending chains. Then there exists a unique stable matching of L.*

Define a pair of points (x, y) to be **mutually closest** in a set S if for all z in S – {x, y}, |x – z| > |x – y| and |y – z| > |y – x|. Let {$\mathbf{M_i}$} be an infinite sequence of matchings of points in L. Define M_0 = {}. For each i > 0, define M_i as follows: let S be the set of points that are in L but that are not matched in M_{i-1}. Then, find the set of all mutually closest pairs of points in S, and let M_i be the union of M_{i-1} and those pairs. Let **M** be the union of all M_i in the sequence.

***Lemma 2.1.** M is a subset of every stable matching of L.

Proof. We show that every M_i is a subset of every stable matching. We use proof by induction on M_i. Our base case is M_0: the empty set is a subset of every set, so the statement is clearly true. As our induction step, we assume that for any stable matching P, M_i is a subset of P, and we show that the same holds for M_{i+1}. Let (x, y) be a pair of points that are matched in M_{i+1} but not in M_i. Assume by way of contradiction that x is not matched with y in P. Let x be matched with x' and y be matched with y' in P. If x' is matched to some element z in M_i, then (x', z) is an element of P by the assumption of the induction step, which is a contradiction, as x' is not matched to z in P. Therefore, x' is unmatched in M_i, and by the same argument, so is y'. Since x and y are matched in M_{i+1}, x and y are a mutually closest pair of points among the remaining set of points, so x prefers y to x' and y prefers x to y'. Therefore (x, y) is an unstable pair of P, a contradiction.

***Lemma 2.2.** Every subset S of L that contains at least two points must contain at least one pair of points that are mutually closest.

Proof. Assume by way of contradiction that no pair of points in S is mutually closest. For any point x of S, let C(x) denote the closest point in S to x. C is well-defined because S is discrete, so any bounded subset of S contains finitely many points, and since pairwise distances are discrete, we can identify a unique smallest distance from x to another point. Now, take any point x_0 in S, and define $x_i = C(x_{i-1})$ for $i > 0$. Since x_i and x_{i+1} cannot form a mutually closest pair, we know $x_i \neq x_{i+2}$, so $|x_i - x_{i+1}| > |x_{i+1} - x_{i+2}|$. Therefore $\{|x_i - x_{i+1}|\}$ is a strictly decreasing sequence, so $\{x_i\}$ is an infinite descending chain, a contradiction.

Lemma 2.3. At most one point in L is left unmatched in M.

Proof. Let S be the set of points in L that are unmatched in M, and assume by way of contradiction that S contains at least two points. By Lemma 2.2, S contains a pair of two mutually closest points, which we call x and y. Since L is discrete, the set of points $T = \{z \in (L - S): |x - z| < |x - y| \text{ or } |y - z| < |x - y|\}$ contains finitely many points z. Each of these points is matched in some M_i, so we can find some large k such that every point in T is matched in M_k. Then we consider the points that are matched in M_{k+1}. x prefers y to every point not in T, and y prefers x to every point not in T, and every point in T is already matched, so x and y are a mutually closest pair of points in $L - M_k$. Therefore x and y should have been matched in M_{k+1}.

Now, we only need to show that M is stable. By way of contradiction, assume that (x, y) is an unstable pair of M. By Lemma 2.3, at least one of x and y is matched, so we consider the smallest i such that x or y is matched in M_i. Without loss of generality, this gives us that x and y are both unmatched in M_{i-1}, and that x is matched to x' in M_i. However, since (x, y) is an unstable pair and since y is still unmatched in M_{i-1}, then (x, x') could not have been a mutually closest pair, as x is closer to y than to x'.

Finally, by Lemma 2.1, M is a subset of any stable matching P of L. By Lemma 2.3, M is a stable matching of all points in L (except possibly one point), so it is not possible to add another pair to M, so there is no stable matching that contains any pairs that are not in M. Therefore P is equal to M, so M is the unique stable matching of L.

Theorem 3. *Suppose $L \subset \mathbb{R}^d$ has all pairwise distances distinct and has no infinite descending chains. Let M be the stable matching of L and consider friendly frogs on L. The position with the two frogs at x and y is a P-position if and only if x is matched to y in M.*

We say that a position is in M if and only if the frogs are at the points x and y, where x and y are matched in M.

***Lemma 3.1.** Either the game terminates in finitely many moves, or each frog moves infinitely many times.

Proof. Assume by way of contradiction that there exists an infinite sequence of legal moves from some starting position, and that one frog (we call this frog F, and the other frog G) is moved only finitely many times. Then after some finite move when F is at the point x and G is at the point y, F stays at x for the rest of the game while G is moved closer and closer to x. Then G is moved to infinitely many distinct points, all of which must in the ball around x of radius $|x - y|$, as the distance between F and G cannot increase above $|x - y|$. Therefore a bounded subset of L contains infinitely many points, so L is not discrete, which is a contradiction, as every set with no infinite descending chains is discrete.

***Lemma 3.2.** The game is guaranteed to terminate in finitely many moves.

Proof. Assume by way of contradiction that there exists an infinite sequence of legal moves from some starting position. Then Lemma 3.1 tells us that each frog is moved infinitely many times. Let x_0 and x_1 be the initial positions of the frogs. Define x_2 as the first point that the first frog is moved to. Then for each $n > 2$, if x_{n-1} was chosen as a point that one frog was moved to on move M, define x_n as the first point that the other frog is moved to on any move after M. Note that since the distance between the frogs is strictly decreasing on every move, $|x_n - x_{n-1}| < |x_{n-1} - x_{n-2}|$. Therefore the sequence $\{x_n\}$ is an infinite descending chain, a contradiction.

Lemma 3.3. If the current position is not in M, then the current player can move so that the next position is in M.

Proof. Let the frogs be at the points x and y, where x and y are not matched in M. If both x and y are matched, then let x be matched to x’ and y be matched to y’, and without loss of generality, assume $|x - x'| < |y - y'|$. If $|x - y| < |x - x'|$, then $|x - y| < |y - y'|$, so (x, y) is an unstable pair of M, which is a contradiction as M is stable, so $|x - y| > |x - x'|$. If either x or y is unmatched, then without loss of generality, assume y is unmatched. If $|x - y| < |x - x'|$, then (x, y) is an unstable pair of M, which is a contradiction as M is stable, so $|x - y| > |x - x'|$. Therefore, in either case, moving the frog from y to x’ decreases the distance between it and the frog at x, so the current player can move the frog from y to x’, and the frogs are made to be at matched points.

Lemma 3.4. If the current position is in M, then any move by the current player results in the next position not being in M.

Proof. Let the frogs be at the points x and y, where x and y are matched in M. Without loss of generality, assume the current player moves the frog at y to z. Each point is only matched to one other point, so x is not matched to z, so the next position is not in M.

We copy the rest of the proof verbatim from Lemmas 1.4 and 1.5 and the proof of Theorem 1. Lemma 3.3 replaces Lemma 1.2 and Lemma 3.4 replaces Lemma 1.3, but the proof is otherwise completely unchanged. Lemma 3.2 guarantees that the game eventually terminates. This gives us that if the current position is not in M, then the current position is an N-position, and that if the current position is in M, then the current position is a P-position, so we are done proving Theorem 3. The proof of Theorem 1 gives one additional result: that Alice wins if M has an unmatched point, and that Bob wins if M is a perfect matching. We call this result Corollary 3.5.

Corollary 3.5. Suppose $L \subset \mathbb{R}^d$ has all pairwise distances distinct and has no infinite descending chains. Let M be the stable matching of L and consider friendly frogs on L. The game is a win for Alice if M is not perfect, and a win for Bob if M is perfect.

Theorem 4. *Let L be a Poisson process on $\mathbb{R}^d$. With probability 1, friendly frogs on L is a win for Bob.*

As stated in the Infinite Point Set Terminology section, it is a theorem (though it is not proved here) that in a Poisson process, there is probability 1 that all pairs of points have distinct distances and that there are no descending chains. Therefore, by Theorem 2, there is probability 1 that there is a unique stable matching M of L. By Corollary 3.5, to show that the game is a win for Bob with probability 1, we only need to show that M is perfect with probability 1. Equivalently, we will show that there is an unmatched point with probability 0.

Let X be the expected value of the number of unmatched points. By Lemma 2.3, there is at most one unmatched point, so $X \leq 1$. Let Y be the expected value of the number of unmatched points in a single unit cube. By the translation-invariance of the Poisson process, Y does not depend on the location of the unit cube, so Y is simply a constant. We can partition $\mathbb{R}^d$ into an infinite set S of disjoint unit cubes, such as by partitioning $\mathbb{R}^d$ by a d-dimensional unit grid. Then $X = \Sigma_S Y$, so if $Y > 0$, then the sum doesn't converge and X is infinite, which is a contradiction as $X \leq 1$. Therefore Y must be equal to 0, so X is an infinite sum of zeroes, so $X = 0$.

Extensions

We consider Theorem 1, relax the restriction that all pairs of points have distinct distances, and allow L to be a subset of an arbitrary metric space. I give a method to construct a simple winning strategy in this more general case. This is an extension beyond what is discussed in the source paper, and the authors only make slight references to these as areas that need further exploration. Instead of rigorously proving the following assertions, as the proofs would be almost identical to the above proofs but with slightly different language, I just show the method and explain why generalizing the above proofs doesn't lead to problems.

Extending Euclidean space to a general metric space is natural in this context, as we never use properties specific to the Euclidean metric. We use distances between points and infinite descending chains, which are defined in all metric spaces. We also use the notion of bounded and discrete sets, which are also both defined in all metric spaces. This gives us that Theorems 1, 2, and 3 are all valid if L is a subset of any metric space. Theorem 4 does not easily extends to arbitrary metric spaces, as the Poisson process is normally only defined in Euclidean space, and our argument also makes use of cubes, which don't necessarily exist in any metric spaces.

Removing that pairs of points have distinct distances is a much more interesting extension, as it changes the fundamental ideas of the proofs. To see this, imagine friendly frogs on a finite set L that consists of the vertices of an equilateral triangle. Theorem 1 tells us that the game is a win for Alice, as the size of L is odd. However, no matter which points Alice and Bob start at, Alice is unable to move, as she cannot strictly decrease the distance between the two frogs.

To overcome this difficulty, we extend the concept of matchings to allow for one point to be matched to multiple other points. Formally, a **multi-matching** of a set of points L is a set M of unordered pairs of distinct points in L such that if x is matched to both y and z in M, then $|x - y| = |x - z|$. We also define an unstable pair of a multi-matching to be any pair (x, y) such that x likes y at least as much as its current partner (if any) and y likes x at least as much as its current partner (if any). We replace all references to matchings in the proofs above with the equivalent statements about multi-matchings.

We now change Theorem 1 to state that Alice wins if there exists an unmatched point, while Bob wins if the multi-matching is perfect. We need to do this because perfect multi-matchings don't necessarily contain an even number of elements, so if the size of L is odd, an unmatched point might not exist. Then, in Theorem 1, instead of adding the pair with the smallest distance to M, we add the set of all pairs with the smallest distance to M. Similarly, in Theorem 2, instead of saying x and y are mutually closest if for all z in $S - \{x, y\}$, $|x - z| > |x - y|$ and $|y - z| > |y - x|$, we say x and y are mutually closest if for all z in $S - \{x, y\}$, $|x - z| \geq |x - y|$ and $|y - z| \geq |y - x|$. These changes are important because they allow a point to enter a multi-matching M_i as part of multiple pairs if the point is equally close to multiple other points.

One other small but important change we must make to the construction of M in Theorem 2 is, if x is unmatched in M_i and y is matched to z in M_i, we add the pair (x, y) to M_{i+1} if (x, y) is mutually closest in $S + \{y, z\}$. This means that we can add pairs containing y to M_{i+1} even if y is already matched in M_i. (This avoids some edge cases by ensuring that if x is matched to x' and y is matched to y', and if $|x - x'| < |y - y'|$, then $|x - x'| < |x - y|$, which is required by the new definition of unstable pairs above).

If we make all of these changes, then all of our theorems hold true in this more general case. A unique stable matching of a set of points L in any metric space exists as long as L has no infinite descending chains. Whether L is finite or infinite, Alice wins if there exists an unmatched point by placing her frog at that point, and Bob wins if there does not exist an unmatched point. I don't formally prove these results only because the proofs are almost the exact same as the proofs presented above, so the new proofs don't give any useful insights beyond the core idea of a multi-matching.

Source Paper

This paper was based heavily off of a source paper published in the *American Mathematical Monthly*, called "Friendly Frogs, Stable Marriage, and the Magic of Invariance." All credit for the core ideas of the proofs presented in this paper go to the authors of the source paper, Maria Deijfen, Alexander E. Holroyd and James B. Martin. Only the Extensions section and the proofs of the lemmas marked with a * are completely original work.

Deijfen, Maria, et al. "Friendly Frogs, Stable Marriage, and the Magic of Invariance." *The American Mathematical Monthly*, vol. 124, no. 5, 2017, pp. 387–402. *JSTOR*, www.jstor.org/stable/10.4169/amer.math.monthly.124.5.387.