

A Category Theoretic Approach to Compositional Approximation in Game Theory

Neil Ghani

MSP Group, University of Strathclyde, Scotland
and Symbolica, London
ng@gcis.strath.ac.uk
neil@symbolica.ai

This paper uses category theory to guide the development of an entirely new approach to approximate game theory. Game theory is the study of how different agents within a multi-agent system take decisions. At its core, game theory asks what an optimal decision is in a given scenario. Thus approximate game theory asks what is an approximately optimal decision in a given scenario. This is important in practice as — just like in much of computing — exact answers maybe too difficult (or even impossible) to compute given inherent uncertainty in input.

We consider first *Selection Functions* which are a simple model of compositional game theory. We develop i) a simple yet robust model of approximate equilibria; ii) the algebraic properties of approximation wrt selection functions; and iii) relate approximation to the compositional structure of selection functions. We then repeat this process for *Open Games* — a more advanced model of game theory featuring the key operation of sequential composition of games. Finally, we use approximation to develop new metrics on game theory and show these yield vital theorems in what one might term "Quantitative Compositionality".

1 Introduction

This paper uses category theory to develop an entirely new approach to compositional approximate game theory. Approximation in computation is ubiquitous e.g. it arises in: i) stochastic systems, where one merely has probability distributions over values induced by inherent/simulated randomness; ii) resource limited environments, where exact computation is prohibitively expensive; iii) systems with imperfect/partial recall, where one only has limited information about what has happened or the intentions/trustworthiness of each agent; and iv) non-exact computation where primitive data (e.g. from sensors) is inexact and supplied with error bars. These scenarios arise in e.g. cyber-physical systems, machine learning, robotics, automotive engineering, aerospace, and energy systems.

On the other hand, game theory is the study of how different agents within a multi-agent system take decisions. The simplest model of game theory we consider are the *selection functions* which are functions $s : (X \rightarrow V) \rightarrow PX$. Here, i) the set X is thought of as a set of actions (or moves, or choices) the agent can consider playing; ii) the set V is thought of as a type of utilities; and iii) a function $X \rightarrow V$ is thought of as assigning, to each possible action, a utility measuring how good that action is deemed to be. A selection function then maps *every* utility function to those actions deemed optimal for that utility function. The canonical selection function is argmax defined by defined by

$$x \in \text{argmax } f \text{ iff } (\forall x' : X) f x \geq f x'$$

Another selection function is ε -argmax defined by

$$x \in \varepsilon\text{-argmax } f \text{ iff } (\forall x' : X) f x \geq f x' - \varepsilon$$

which means for an action x to be ε -optimal it need merely have an associated utility which is within ε of a true optimal value. This approach aligns with the general approach to approximation of uncertainty budgets as we track - and therefore can check - how much utility is lost by moving to an approximately optimal decision. Notice this is NOT saying x is close to an optimal choice of action, rather that the utility of x is close to being the utility of an optimal action. One might study the former and, indeed, this was out first intention. However, the mathematical behaviour of that concept was poorer so we switched to the approach taken in this paper. Nevertheless, this was the first of several moments where key decisions need to be taken to obtain a well behaved mathematical theory showing the combination of approximation and game theory to be a delicate and nuanced endeavour.

A second difficulty arises from the desire for *compositionality* to ensure scalability. Compositional systems are those where larger systems can be defined by assembling smaller subsystems together which means that proving properties of larger systems can be reduced to proving properties of smaller (and hence easier to reason about) subsystems. At its core, compositionality is a story about structure and so the natural mathematical medium in which to develop compositionality is category theory. Recently, Compositional Game Theory [6] was introduced to bring the advantages of compositionality to game theory. However, within compositional game theory there are a number of different notions of morphisms of open games and, as we shall see, certain notions of morphism behave better wrt approximation than others. Finding the right notion of morphism of open game to interact well with approximation was another issue we had to solve.

Concretely, we

- introduce a notion of the ε -approximation of a game by defining for any game G , another game $T_\varepsilon G$ which we think of as containing not just the equilibria of G , but also the ε -equilibria of G — this is an entirely new approach to approximation of equilibria and shows novelty of methodology.
- show that show T_ε has good algebraic structure in that it is i) functorial; and ii) posses graded structure — this shows our methodology is canonical in being aligned with other work on quantitative computation.
- show T_ε interacts well with the compositional structure of game theory, in particular with Nash equilibria and sequential composition. There is no particular reason that the approximation of a large game built from game-constructors should retain that structure — we show this holds for the key cases of parallel (or Nash) product and sequential composition. This is crucial evidence for the utility of our approach, namely that preserves the compositional structure of game theory.
- use our approximate game theory to create what one might call *Quantitative Compositional Game Theory* which i) not only builds new games from old; but ii) ensures if one substitutes a subgame with another which has similar behaviour, the overall system remains similar to the original. The crucial construction permitting *Quantitative Compositional Game Theory* is a metric on games which we build from our approximate game theory.

This third item is particularly surprising and thus our strongest contribution — as mentioned there in not inherent reason why, for example, the approximations of a Nash equilibria should consist of strategies

that form a Nash equilibria. Finally, we argue in the conclusions that our methodology is also general in that it can apply to other forms of game theory, eg Markov Decision Processes, and beyond game theory into the huge area of machine learning. This latter application area particularly exciting as formally quantifying distance from optimality is a way of expressing how much has been learned and thus contributing to the XAI agenda.

This paper is structured as follows. In Section 2, we cover preliminaries pertaining to metric spaces, lenses and selection functions. Section 3 develops approximation for selection functions, discusses the algebraic properties of approximation and show approximation is compositional. Section 4 extends these results to open games [6], while Section 5 uses approximation to develop Metric Game Theory. Section 6 concludes the paper with conclusions and directions for future research.

2 Preliminaries: Metric Spaces and Lenses

This paper should be accessible to those with a basic understanding of game theory and category theory. Introductory texts are [14, 10, 12], for game theory and [11, 9] for category theory.

2.1 Metric Spaces

Let \mathbb{R} be the positive reals extended with an infinity ∞

Definition 2.1. *Define:*

1. An extended metric space is a set X with a function $d : X \times X \rightarrow \mathbb{R}$ satisfying symmetry, the triangle inequality and $d(x, y) = 0$ iff $x = y$.
2. EMet is the category of extended metric spaces and short maps, i.e maps f satisfying $d(fx, fy) \leq d(x, y)$.

We write $|X|$ for the underlying set of a metric space X .

Definition 2.2. *If (X, d) and (Y, d') are extended metric spaces, then $(X, d) \otimes (Y, d')$ is the extended metric space with underlying set $X \times Y$, where*

$$d_{X \otimes Y}((x, y), (x', y')) = \max\{d(x, x'), d(y, y')\}$$

Other options are possible, eg sum, average etc. The following explains the use of ∞ as a possible distance — it is needed to ensure the metric on function spaces is well defined.

Definition 2.3. *Let X, Y be a metric spaces. Then a non-expansive (or short) map $f : X \rightarrow Y$ is a function $f : X \rightarrow Y$ such that*

$$d(fx, fy) \leq d(x, y)$$

The set of non-expansive functions $[X, Y]$ is a metric space via the sup-metric:

$$d(f, g) = \sup_{x \in X} d(fx, gx).$$

2.2 Lenses

While lenses can be built over categories with appropriate structure, to keep our presentation simple, we work with lenses over sets.

Definition 2.4. A lens is given by a pair of sets (X, V) . A map of lenses $f : (X, V) \rightarrow (Y, Z)$ is given by a pair (f_0, f_1) consisting of functions $f_0 : X \rightarrow Y$ and $f_1 : X \times Z \rightarrow V$.

The following is well known:

Lemma 2.5. Lenses from a symmetric monoidal category Lens with tensor given by $(X, V) \otimes (X', V') = (X \times X', V \times V')$ and with unit given by $I = (1, 1)$.

Note i) an element of X is a lens-map $I \rightarrow (X, V)$; and ii) a function $X \rightarrow V$ is equivalently a lens map $(X, V) \rightarrow I$. These two observations are useful, eg we can often represent complex functions $X \rightarrow V$ more simply as composites of lens maps.

2.3 Selection Functions

A basic form of game theory is given by selection functions [2] which, given sets X, Y , are defined by $\text{Sel}(X, Y) = (X \rightarrow Y) \rightarrow PX$ where PX is the set of powersets of X . Intuitively, a selection function maps a utility function (which assigns to each potential action $x \in X$ a resulting utility $y \in Y$) to a set of actions deemed to be optimal for that specific utility function. Give two selection functions $s, t : (X \rightarrow Y) \rightarrow PX$, we write $s \subseteq t$ iff $(\forall f : X \rightarrow Y) sf \subseteq tf$. The canonical selection function is argmax defined as follows:

Example 2.6. Define the argmax selection function $\text{argmax} : (X \rightarrow Y) \rightarrow PX$ by

$$\text{argmax}f = \{ x \in X \mid (\forall x' \in X) fx \geq fx' \}$$

Note that the selection function model of game theory deals with not merely one utility function but, rather, all utility functions. This is also the case of open games and is a hallmark — indeed key distinguishing feature — of compositional models of game theory. When it comes to maps of selection functions, we choose the following for reasons that will become clearer later.

Definition 2.7. Let $s : \text{Sel}(X, V)$ and $t : \text{Sel}(X', V')$. A map from s to t consists of a map of lenses $\alpha : (X, V) \rightarrow (X', V')$ such that

$$(\forall k : X' \rightarrow V')(\forall x : X) \alpha x \in tk \Rightarrow x \in s(k\alpha)$$

Note in the above we write the map $X \rightarrow V$ as the composite of lens maps $k\alpha$ as this is much cleaner than a direct definition. This is an one example of the power of lens composition to simplify definitions as we remarked earlier. We continue to use lens composition via juxtaposition through the rest of the paper without remarking on it. The above definition has also possesses a contravariance akin to how maps of containers (aka dependent lenses) are contravariant in their position maps. An alternative covariant definition for a map of selection functions would have been to require

$$(\forall k : X' \rightarrow V') x \in s(k\alpha) \Rightarrow \alpha x \in tk$$

but, as we shall see later, this does not give a functorial theory of approximation. Henceforth we write Sel for the category of all selection functions and the maps between them. Selection functions are a compositional model of game theory as they support a Nash product. This arises by extending the monoidal product on lenses to one on selection functions:

Definition 2.8. *The category Sel of selection functions and maps of selection functions is monoidal with monoidal product*

$$\otimes : \text{Sel}(X, Y) \times \text{Sel}(X', Y') \rightarrow \text{Sel}(X \times X', Y \times Y')$$

defined by

$$(x, x') \in (s \otimes t)f \quad \text{iff} \quad \begin{array}{l} x \in s(\pi_0 f(-, x')) \\ \wedge \\ x' \in t(\pi_1 f(x, -)) \end{array}$$

Note as an example, we have the beautiful equation showing Nash equilibria are not a primitive but rather a derived concept.

$$\text{Nash} = \text{argmax} \otimes \text{argmax}$$

This equation also highlights how our tensor of selection functions can be seen concretely as the two games played in parallel with Nash as a specific example of two argmaxes run in parallel. Not only will we develop approximate game theory, we will ensure it is compositional in that approximation interacts well with compositional operators such as the Nash product.

3 Approximation and Selection Functions

Given selection functions form a compositional model of game theory we can ask two natural questions: i) can we account for approximate equilibria in selection functions; and ii) is approximation compositional in that — for example — if we approximate a Nash product of two games, do we get the Nash product of two approximations? We answer both questions positively — the second to our surprise as there is no particular reason to believe that approximation of a structured game preserves that structure. The fact it does suggests the compositional approach to game theory and the approach to approximation advocated here are well designed concepts capturing fundamental logical structure which other approaches have missed. We begin by defining approximation in selection functions via a function

$$\top : (\varepsilon : \mathbb{R}) \rightarrow \text{Sel}(X, V) \rightarrow \text{Sel}(X, V)$$

by

$$x \in (\top_\varepsilon G)k \quad \text{iff} \quad (\exists k' \sim_\varepsilon k) \quad x \in Gk'$$

where we write $k' \sim_\varepsilon k$ for $d(k', k) \leq \varepsilon$. What is nice about this definition is that not only is it extremely simple and thus elegant, but it also formalises the intuition that an ε -optimal choice for a utility function k is an actual optimal choice, but for a utility function k' differing no more than ε from k . Getting this definition right, and getting the right notion of morphism between selection functions, were the fundamental keys required to achieving the smooth mathematical theory we have been able to develop in this paper. The first thing to do is to relate our notion of approximation to the traditional notion of approximation that one finds in the literature.

Lemma 3.1. *Let $\varepsilon \geq 0$. Then $\top_\varepsilon(\text{argmax}) \subseteq 2\varepsilon\text{-argmax}$*

Proof. Let $x \in \top_\varepsilon(\text{argmax})k$ for some utility function $k : X \rightarrow V$. Then there is a $k' \sim_\varepsilon k$ such that $x \in \text{argmax}(k')$, ie for all $x' \in X$, we have $k'x \geq k'x'$. But since $k' \sim_\varepsilon k$, we thus have $kx \geq kx' - 2\varepsilon$. Thus $x \in 2\varepsilon\text{-argmax } k$ □

Assuming X has a decidable equality, which includes the key case where X is finite, we can go further

Lemma 3.2. *Let $\varepsilon \geq 0$. If X has a decidable equality, then*

$$\varepsilon\text{-argmax} \subseteq T_\varepsilon(\text{argmax})$$

Proof. Let $x \in \varepsilon\text{-argmax}(k)$ for a utility function $k : X \rightarrow V$. Then for all $y \in X$, $kx \geq ky - \varepsilon$. Define $k' : X \rightarrow V$ by $k'x = kx + \varepsilon$ and $k'y = ky$ for all $y \neq x$. Clearly $k \sim_\varepsilon k'$. Further, for any $y \neq x$, we have $k'x \geq k'y$. Thus $x \in \text{argmax}(k')$ and so $x \in T_\varepsilon(\text{argmax})(k)$ \square

Moving to a categorical analysis, the first thing we would like to establish is the functoriality of T_ε . This is non-trivial as the needs of approximation put constraints on how lens maps back-propagate values. In essence, we must restrict to *short* lens maps which we define now

Definition 3.3. *A short lens map is a lens map $\alpha : (X, V) \rightarrow (X', V')$ where, for every $x \in X$, $\alpha(x, -) : V' \rightarrow V$ is a short map, ie non-expanding.*

Short maps are a natural condition — the condition reflects desire that coutilty (or backpropagation) doesn't inflate differences in the values to be backpropagated. Were we to allow such inflation, system behaviour could well become chaotic as small amounts of sub-optimality would be magnified into large system reconfiguratation. Note the identity is a short map and the composite of short maps is short and hence the category of lenses restricts to a category of lenses and short maps denoted Lens_s . Similarly, we define the category of selection functions and maps whose underlying lens map is short by Sel_s . And, finally, the tensor product on lenses and selection functions restricts to a monoidal structure on Sel_s

Lemma 3.4. *Let $\alpha : (X, V) \rightarrow (X', V')$ be a short lens map. If $k, k' : X' \rightarrow V'$ is such that $k \sim_\varepsilon k'$, then $k\alpha \sim_\varepsilon k'\alpha$*

Proof. Direct from the definition of what a short lens map is. \square

We can now prove functoriality:

Lemma 3.5. *Let $\alpha : s \rightarrow t$ be a map of selection functions over the short lens map $\alpha : (X, V) \rightarrow (X', V')$ of lenses. Then α also defines a map of selection functions $T_\varepsilon s \rightarrow T_\varepsilon t$*

Proof. Consider $x \in X$ and $k : X' \rightarrow V'$. We need to show if $\alpha x \in (T_\varepsilon t)k$, then $x \in (T_\varepsilon s)(k\alpha)$. The assumption implies there is a $k_1 : X' \rightarrow V'$ such that $k \sim_\varepsilon k_1$ and $\alpha x \in tk_1$. As $\alpha : s \rightarrow t$, this means $x \in s(k_1\alpha)$. By the previous lemma, $k\alpha \sim_\varepsilon k_1\alpha$ and so $x \in T_\varepsilon s(k\alpha)$ as required. \square

Next we consider the role/structure of the parameter ε in approximation. The following graded structure is what one expects in many quantitative situations.

Lemma 3.6. *The following are easily provable*

- $T_0 G = G$
- If $\varepsilon \leq \varepsilon'$, then $T_\varepsilon G \subseteq T_{\varepsilon'} G$
- $T_\varepsilon(T_\delta G) \subseteq T_{\varepsilon+\delta} G$

Proof. Straightforward \square

3.1 Compositional Approximation

In this subsection we discuss the relationship between the compositional structure of selection functions and approximation of selection functions. To do this we first prove a small lemma

Lemma 3.7. *Let $k, k' : X \times X' \rightarrow V \times V'$ be functions such that $k \sim_\varepsilon k'$. Then for every x and x' ,*

$$\pi_0 k(-, x') \sim_\varepsilon \pi_0 k'(-, x') \quad \text{and} \quad \pi_1 k(x, -) \sim_\varepsilon \pi_1 k'(x, -)$$

Proof. Straightforward calculation □

Now, imagine we take a Nash equilibrium of selection functions G and H defined by the monoidal product of selection function as $G \otimes H$. What do its approximate equilibria look like? Is there any structure here or does it just consist of a fuzzy ball of extra equilibria enveloping G and H . The following shows that structure is preserved, ie that an approximate equilibrium of a Nash equilibrium is still a Nash equilibrium.

Lemma 3.8. *Let $G : \text{Sel}(X, V)$ and $H : \text{Sel}(X', Y')$ be selection functions. Then*

$$T_\varepsilon(G \otimes H) \subseteq T_\varepsilon G \otimes T_\varepsilon H$$

Proof. Let $(x, x') \in X \times X'$ and $k : X \times X' \rightarrow V \times V'$. Further assume $(x, x') \in T_\varepsilon(G \otimes H) k$. Then there is a $k' : X \times X' \rightarrow V \times V'$ such that $k \sim_\varepsilon k'$ and $(x, x') \in (G \otimes H) k'$. By definition of $G \otimes H$, this means

$$x \in G(\pi_0 k'(-, x')) \quad \text{and} \quad x' \in H(\pi_1 k'(x, -))$$

By the previous lemma

$$x \in T_\varepsilon(G)(\pi_0 k(-, x')) \quad \text{and} \quad x' \in T_\varepsilon(H)(\pi_1 k(x, -))$$

and hence $(x, x') \in (T_\varepsilon G \otimes T_\varepsilon H) k$. Thus $T_\varepsilon(G \otimes H) \subseteq T_\varepsilon G \otimes T_\varepsilon H$ as required. □

The above lemma states that the approximation of a Nash equilibrium is itself a Nash equilibrium. Noting the contravariance in our notion of morphism of selection functions, we can summarise the above as:

Lemma 3.9. *The functor $T_\varepsilon : \text{Sel}_s \times \text{Sel}_s \rightarrow \text{Sel}_s$ is lax monoidal functor*

However, it would be nice if the reverse were true as then we could *find* and approximate equilibrium for a Nash equilibria by merely approximating the components. Not suprisingly, the ability to take this extra step depends on the existence of sufficient utility functions as we saw in lemma 3.2 which in turn rests on the decidability of equality on actions. The key lemma is the following

Lemma 3.10. *Let X, X' have decidable equality. Further, let $k : X \times X' \rightarrow V \times V'$, $x \in X$ and $x' \in X'$. Even further, let $g \sim_\varepsilon \pi_0 k(-, x')$ and $h \sim_\varepsilon \pi_1 k(x, -)$. Then there is a $k' : X \times X' \rightarrow V \times V'$ such that*

- $k' \sim_\varepsilon k$
- $g = \pi_0 k'(-, x')$; and
- $h = \pi_1 k'(x, -)$.

Proof. Define $k' : X \times X' \rightarrow V \times V'$ as follows

$$\begin{aligned} k'(z, z') &= k(z, z') && \text{iff } z \neq x, z' \neq x' \\ k'(x, z') &= (\pi_0 k(x, z'), h z') && \text{iff } z' \neq x' \\ k'(z, x') &= (g z, \pi_1 k(z, x')) && \text{iff } z \neq x \\ k'(x, x') &= (g x, h x') && \text{otw} \end{aligned}$$

The required properties of k' are now easy to verify. □

Now we can prove the reverse containment

Lemma 3.11. *Let X, X' have decidable equality. Let $G : \text{Sel}(X, V)$ and $H : \text{Sel}(X', Y')$ be selection functions. Then*

$$\top_\varepsilon G \otimes T_\varepsilon H \subseteq \top_\varepsilon (G \otimes H)$$

Proof. Let $(x, x') \in X \times X'$ and $k : X \times X' \rightarrow V \times V'$ be such that $(x, x') \in (\top_\varepsilon G \otimes T_\varepsilon H)k$. Then

$$x \in (\top_\varepsilon G)(\pi_0 k(-, x')) \quad \text{and} \quad x' \in (\top_\varepsilon H)(\pi_1 k(x, -))$$

Thus there is a function $g : X \rightarrow V$ and one $h : X' \rightarrow V'$ such that

$$x \in Gg \quad \text{and} \quad x' \in Hh \quad \text{and} \quad g \sim_\varepsilon \pi_0 k(-, x') \quad \text{and} \quad h \sim_\varepsilon \pi_1 k(x, -)$$

By lemma 3.10, there is a map $k' : X \times X' \rightarrow V \times V'$ such that

$$x \in G(\pi_0 k'(-, x')) \quad \text{and} \quad x' \in H(\pi_1 k'(x, -)) \quad \text{and} \quad k' \sim_\varepsilon k$$

Thus $(x, x') \in (G \otimes H)k'$ and hence $(x, x') \in \top_\varepsilon (G \otimes H)k$ □

This is a very strong — and surprising result — that for decidable action sets, approximation is not merely lax monoidal, but actually monoidal. This we did not expect and, indeed, when it comes to sequential composition of open games, we can only show one containment. Sequential composition of games leads us from selection functions to Open Games.

4 Approximation and Open Games

Selection functions form a simple model of compositional game theory via the Nash product/parallel composition of selection functions. However, what about the sequential composition of selection functions — afterall parallel and sequential composition are intimately bound in monoidal category theory, string diagrams etc. Sequential composition of selection functions have been considered [2] but the utility function of the first game is modelled exogenously ... what happens in practice is that the utility function of the first game is actually computed endogenously. Open Games [6] address this problem perfectly. For the price of a little more algebraic structure, they form a compositional model of game theory which includes a number of operators for compositionally building large complex games from smaller and simpler games. Chief amongst these operators are sequential composition but also choice operators, and operators for iterated game theory [7]. This section generalises approximation to open games.

4.1 Open Games

The key concept which introduced open games is the following:

Definition 4.1 (Open Game). *Let X, Y, R and S be sets. An open game $G = (\Sigma_G, P_G, C_G, E_G) : (X, S) \rightarrow (Y, R)$ consists of:*

- a set Σ_G of strategy profiles,
- a play function $P_G : \Sigma_G \rightarrow (X \rightarrow Y)$,
- a coulity function $C_G : \Sigma_G \rightarrow (X \times R \rightarrow S)$, and
- an equilibrium function $E_G : X \times (Y \rightarrow R) \rightarrow \text{P}(\Sigma_G)$.

Intuitively, the set X contains the states of the game, Y the moves, R the utilities and S the coutilities. The set Σ_G contains the strategies we are trying to pick an optimal one from. The play function P_G selects a move given a strategy and a state, while the coutility function C_G computes the coutility extruded from the game, given a strategy, state and utility. Finally, if $\sigma \in E_G x k$, then σ is an optimal strategy in state x and with utility given by $k : Y \rightarrow R$. Just as lenses systematise the presentation of selection functions, they also systematise the presentation of open games.

Lemma 4.2. *An open game $G : (X, S) \rightarrow (Y, R)$ is given by a set Σ of strategies and, for each $\sigma \in \Sigma$: i) a lens $G_\sigma : (X, S) \rightarrow (Y, R)$; and ii) a predicate*

$$E_\sigma \subseteq X \times (Y \rightarrow R)$$

We now restrict all lenses occurring in open games to be short. This is important not only in functoriality as before, but also for the compositional treatment of approximation for sequential games. Notice the similarity of the lens predicate E_σ and a selection function. But notice also differences — a selection function is defined over a pair of sets, ie an object of Lens, while an open game is defined over a family of lens maps. There are other notions of open games, in particular Gavranovic Games [1] which deepen the relationship with lenses by defining open games to be parameterised maps in the category of lenses. We expect our techniques apply to Gavranovic games in the obvious way.

4.2 Approximation and Open Games

The methodology developed for selection function clearly generalises to open games as follows:

Definition 4.3. *Let $G : (X, S) \rightarrow (Y, R)$ be an open game with strategy set Σ and equilibrium predicates E_σ for $\sigma \in \Sigma$. Further let $\varepsilon \in \mathbb{R}$. Then define $\top_\varepsilon G : (X, S) \rightarrow (Y, R)$ to be the open game with i) the same strategy set Σ ; ii) lens structure $(\top_\varepsilon G)_\sigma = G_\sigma$; and iii) equilibrium given by*

$$x \in E_{\top_\varepsilon G, \sigma} k \text{ iff } (\exists k' \sim_\varepsilon k) x \in E_{G, \sigma} k'$$

The similarity of approximation for open games with that for selection functions is obvious and highlights a strength of the methodology developed here — it is likely to generalise to many other situations as discussed in the conclusions and future work. This similarity also means that lemma 3.6 carries over to the setting of open games. Next we consider functoriality of approximation for open games

Definition 4.4. *The category of open games Op has open games as objects, and a morphism from the open game $(\Sigma, G, E) : (X, S) \rightarrow (Y, R)$ to the open game $(\Sigma', G', E') : (X', S') \rightarrow (Y', R')$ consists of a function $f : \Sigma \rightarrow \Sigma'$, (short) lens maps $\alpha : (X, S) \rightarrow (X', S')$ and $\beta : (Y, R) \rightarrow (Y', R')$ such that*

$$(\forall \sigma \in \Sigma)(\forall x \in X)(\forall k : Y' \rightarrow R') \alpha x \in E'_{f\sigma} k \text{ implies } x \in E_\sigma(k.\beta)$$

Notice how the contravariance in maps of selection functions is replicated here in maps of open games. This will be another hallmark of a general theory of approximation for systems consisting of utility indexed predicates. And again notice the use of the shortness of the lens maps α and β . The proof of the following lemma is the natural generalisation of the analogous one for selection functions.

Lemma 4.5. *For $\varepsilon \in \mathbb{R}$, the assignment sending an open game G to an open game $\top_\varepsilon G$ extends to a functor $\top_\varepsilon : \text{Op} \rightarrow \text{Op}$*

The final part of the mathematical development of approximation for selection functions considered the monoidal product formalising Nash equilibria which showed that $\top_\varepsilon(G \otimes H) \subseteq \top_\varepsilon G \otimes \top_\varepsilon H$, and that \subseteq can be replaced by $=$ when moves are decidable. Again, this lemma generalises to open games — see [6] for exact the definition of the Nash tensor for open games.

Lemma 4.6. *The functor $\mathbb{T}_\varepsilon : \text{Op} \rightarrow \text{Op}$ is lax monoidal with respect to the Nash tensor on open games, and monoidal where the type of moves (inputs to utility) is decidable.*

Proof. The proof follows the same structure as Lemma 3.9. \square

One of the great advantages of open games is they compositionally model the sequential composition of games. Indeed, if we were to form a category whose objects are pairs of sets and whose morphisms are open games, the sequential composition of open games would be the composition in that category.

Definition 4.7. *Let $(\Sigma, G, E) : (X, S) \rightarrow (Y, R)$ be an open game and $(\Omega, H, B) : (Y, R) \rightarrow (Z, T)$ be another open game. Their sequential composition is the open game $H \circ G : (X, S) \rightarrow (Z, T)$ defined by*

- *Strategies of $H \circ G$ are $\Sigma \times \Omega$*
- *The lens $(H \circ G)(\sigma, \tau)$ is defined to be the lens composition $H\tau.G\sigma$*
- *The equilibrium predicate is defined by*

$$x \in E_{HG,(\sigma,\tau)}(k : Z \rightarrow T) \text{ iff } x \in E_{G,\sigma}(k.H\tau) \wedge G_\sigma x \in E_{H,\tau} k$$

Just like with the Nash product, we can ask ourselves the if the approximate equilibria of a sequential game has any structure, eg is it composed of equilibria for the sequential composition of two approximated games. Once more, the answer is yes!

Lemma 4.8. *Let $(\Sigma, G, E) : (X, S) \rightarrow (Y, R)$ and $(\Omega, H, B) : (Y, R) \rightarrow (Z, T)$ be open games. Then*

$$x \in E_{\mathbb{T}_\varepsilon(H \circ G),(\sigma,\tau)}(k : Z \rightarrow T) \text{ implies } x \in E_{\mathbb{T}_\varepsilon H \circ \mathbb{T}_\varepsilon G,(\sigma,\tau)}(k : Z \rightarrow T)$$

for any $x \in X$, $k : Z \rightarrow T$, $\sigma \in \Sigma$ and $\tau \in \Omega$

Proof. Assume $x \in X$, $k : Z \rightarrow T$, $\sigma \in \Sigma$ and $\tau \in \Omega$ and $x \in E_{\mathbb{T}_\varepsilon(H \circ G),(\sigma,\tau)}(k : Z \rightarrow T)$. By definition of \mathbb{T}_ε , we have a $k' : Z \rightarrow T$ such that $k \sim_\varepsilon k'$ and $x \in E_{H \circ G,(\sigma,\tau)}(k' : Z \rightarrow T)$. By definition, this means $x \in E_{G,\sigma}(k'.H\tau)$ and $G_\sigma x \in E_{H,\tau} k'$. Since $H\tau$ is short, we have $k.H\tau \sim_\varepsilon k'.H\tau$ and hence

$$x \in E_{\mathbb{T}_\varepsilon G,\sigma}(k.H\tau) \text{ and } G_\sigma x \in E_{\mathbb{T}_\varepsilon H,\tau} k$$

Thus $x \in E_{\mathbb{T}_\varepsilon H \circ \mathbb{T}_\varepsilon G,(\sigma,\tau)}(k : Z \rightarrow T)$ \square

Unlike with the Nash product, it is unlikely we can prove the reverse direction as, given $x \in E_{\mathbb{T}_\varepsilon G,\sigma}(k.H\tau)$, all we could conclude is there is a $k' \sim_\varepsilon k.H\tau$ such that $x \in E_{G,\sigma} k'$ but vital and needed structure is lost.

5 Metric Game Theory

Apart from solving the concrete problem of when is a move almost optimal in a selection function or open game, our Approximate Game Theory permits further theoretical developments. The one we sketch here we call *Metric Game Theory* which seeks to ask the question of how far apart are two games. This is important, eg in a compositional model we might want to know if replace sub-games/components of one game with other sub-game/components, can we quantitatively measure the effect. In more detail, we might ask question such as if $G \sim_\varepsilon G'$, what can we say about the distance between $G \otimes H$ and $G' \otimes H$. In this section we develop the basics of Metric Game Theory. We focus on selection functions as the theory seems to generalise to open games. We also only measure distances between selection functions over the same lens — we expect we can use fibrational structure of selection functions over lenses to enable the computation of distances between selection functions over different lenses. But we leave a proper treatment of *Metric Game Theory* to a full paper as the subject clearly warrants.

Definition 5.1. Let $s, s' \in \text{Sel}(X, V)$ be selection functions. We define

$$d(s, s') = (\inf \varepsilon : \mathfrak{K}) (s \subseteq T_\varepsilon s' \wedge s' \subseteq T_\varepsilon s)$$

In the above, we are essentially using the Hausdorff distance on subsets in the above definition. As a result, the above construction gives a pseudometric, ie there is no requirement that $d(x, y) = 0$ implies $x = y$. We list two obvious properties of the pseudometric on $\text{Sel}(X, V)$ showing how distances between games are preserved by compositional structure.

Lemma 5.2. Assuming a decidable equality on X, X' , the following are true

- Let s be a selection function. Then $d(s, T_\varepsilon s) \leq \varepsilon$
- Let $s, s' \in \text{Sel}(X, V)$ and $t, t' \in \text{Sel}(X', V')$. If $d(s, s') \leq \varepsilon$ and $d(t, t') \leq \varepsilon$, then $d(s \otimes t, s' \otimes t') \leq \varepsilon$

Proof. The first property follows from the graded structure of T_ε . For the second, note $s \subseteq T_\varepsilon s'$ and $t \subseteq T_\varepsilon t'$. Thus $s \otimes t \subseteq T_\varepsilon s' \otimes T_\varepsilon t' = T_\varepsilon (s' \otimes t')$ \square

Finally, note this states if we replace a subcomponent of a Nash product by one ε -different, the overall Nash product is no further than ε -apart. This is an ideal property to have in a quantitative compositional theory of games. The analogous property for the Nash product of open games also holds. However, note that for sequential composition of open games we only have $T_\varepsilon(H \circ G) \subseteq T_\varepsilon H \circ T_\varepsilon G$. Thus, replacing subgames with other subgames can increase distance between the overall games but (not shown here) in a bounded way — enough to give a theory of *uncertainty budgets* for approximation of open games.

6 Conclusions and Future Work

In summary, we have tackled the important problem of approximation for game theory. This is important as exact equilibria are often either impossible to determine due to inherent uncertainty in underlying data, or undesirable because the cost of exact computation is too high in comparison to approximate computation. In developing our approach to approximate game theory, it was vital to ensure our theory of approximation was compositional in that it respected the compositionality of game theory. Our fundamental idea was that a strategy was approximately optimal for a given utility function if it is optimal for an approximation of that utility function. This idea was developed for two models of game theory - firstly selection functions and secondly open games. In both, approximation was shown to preserve compositional structure. Finally, we showed how approximation and compositionality can be used to develop *Metric Game Theory* and then *Quantitative Compositionality* which allows one to replace approximately equal subcomponents and producing guarantees on how close the overall resulting systems are.

By opening up a new field, this paper has much that can be done to expand it. Firstly, there are a number of other operators in open games which the methods developed here could be applied to. These include choice, contextualisation and infinite iteration [7]. Secondly, these methods could also be extended to more advanced forms of open games, eg i) probabilistic open games using distance metrics such as Kantorovich on probability distributions; or ii) games based upon optics [13]. Thirdly, the generality of our ideas means applications to other forms of games are natural. For example, in Markov Decision Processes [3], one can consider coalgebras

$$\langle u, t \rangle : X \rightarrow R \times (A \rightarrow DX)$$

where X is a state space, R is a type of utility, A is a set of actions, D is the probability monad, $u : X \rightarrow R$ is a utility function and $t : X \times A \rightarrow DX$ is a probabilistic transition function. Within MDPs one defines optimal policies which we can denote Opt . Our methodology suggests defining an approximate version

$$\pi \in \text{Opt } T_\varepsilon \langle u, t \rangle \text{ iff } \pi \in \text{Opt } \langle u', t \rangle$$

where $u' \sim_\varepsilon u$. One could even be more general and consider tweaks to the structure of t ! Most excitingly, there is ample evidence that machine learning is fundamentally the same thing as game theory apart from in the former one looks for small step improvements in performance while in the latter one seeks exact equilibria. Thus it is possible that the ideas developed here could align with the learning based upon back propagation [4] of losses based upon utility and, in particular, the categorical deep learning [5] programme advocated recently (and as a special case, geometric deep learning). This would be a contribution to the huge problem of XAI. It would be good to see what an implementation in code of our ideas looks like, eg in the implementation of open games [8]. The breadth of application areas also suggests more theoretical work unifying these disparate examples, eg via a theory of approximation for utility indexed predicates, or maybe predicates fibred over metric spaces, or just fibrations over metric spaces!

References

- [1] Matteo Capucci, Bruno Gavranović, Jules Hedges & Eigil Fjeldgren Rischel (2021): *Towards foundations of categorical cybernetics*. *arXiv preprint arXiv:2105.06332*.
- [2] Martín Escardó & Paulo Oliva (2010): *Sequential games and optimal strategies*. *Proceedings of the Royal Society A* 467, pp. 1519–1545.
- [3] Frank Feys, Helle Hvid Hansen & Lawrence S. Moss (2018): *Long-Term Values in Markov Decision Processes, (Co)Algebraically*. In: *14th International Workshop on Coalgebraic Methods in Computer Science (CMCS), Coalgebraic Methods in Computer Science LNCS-11202*, pp. 78–99.
- [4] Brendan Fong, David I. Spivak & Rémy Tuyéras (2019): *Backprop as Functor: A compositional perspective on supervised learning*. In: *LICS'19, IEEE*, pp. 1–13, doi:10.1109/LICS.2019.8785665.
- [5] Bruno Gavranović, Paul Lessard, Andrew Dudzik, Tamara von Glehn, João G. M. Araújo & Petar Veličković (2024): *Categorical Deep Learning: An Algebraic Theory of Architectures*. *arXiv:2402.15332*.
- [6] Neil Ghani, Jules Hedges, Viktor Winschel & Philipp Zahn (2018): *Compositional Game Theory*. In: *LICS'18, ACM*, pp. 472–481, doi:10.1145/3209108.3209165.
- [7] Neil Ghani, Clemens Kupke, Alasdair Lambert & Fredrik Nordvall Forsberg (2018): *A compositional treatment of iterated open games*. *Theor. Comput. Sci.* 741, pp. 48–57, doi:10.1016/j.tcs.2018.05.026.
- [8] Jules Hedges, Philipp Zahn, André Videla & Sjoerd Visscher: *Open Games in Haskell*. Available at <https://github.com/jules-hedges/open-games-hs>.
- [9] Bart Jacobs (2016): *Introduction to Coalgebra: Towards Mathematics of States and Observation*. Cambridge Tracts in Theoretical Computer Science, Cambridge University Press, doi:10.1017/CBO9781316823187.
- [10] Kevin Leyton-Brown & Yoav Shoham (2008): *Essentials of Game Theory: A Concise Multidisciplinary Introduction*. Morgan & Claypool.
- [11] Saunders MacLane (1971): *Categories for the working mathematician*. Springer-Verlag. Graduate Texts in Mathematics, Vol. 5.
- [12] Martin Osbourne & Ariel Rubinstein (1994): *A course in game theory*. MIT Press.
- [13] Mitchell Riley (2018): *Categories of optics*. *arXiv preprint arXiv:1809.00738*.
- [14] Joel Watson (2002): *Strategy: an introduction to game theory*. 139, WW Norton New York.