The Structure of an 8-state Finite Transducer Representation for Prisoner's Dilemma

Jeffrey Tsang
Department of Mathematics & Statistics
University of Guelph
Guelph, ON, Canada
jeffrey.tsang@ieee.org

ABSTRACT

The fingerprint operator generates a representation-independent functional signature of a game-playing strategy, which enables the automated analysis of evolved agents. With this, we attempt to study the structure of a relatively small representation — the 8-state finite transducers for Prisoner's Dilemma. Even then, there are almost 3×10^{20} strategies representable, and hence we sample 32,768 strategies uniformly at random for investigation. Accounting for phenotypic duplicates, there are 31,531 distinct strategies in the dataset; we compute all pairwise distances and use a variety of dimensionality reduction techniques to embed it into a manageable space. Results indicate no obvious cutoff scales, and a strong structural similarity with parallel studies on the entirety of even smaller state spaces.

Categories and Subject Descriptors

F.1.1 [**Computation by Abstract Devices**]: Models of Computation—*deterministic finite transducers*; I.2.1 [**Artificial Intelligence**]: Applications and Expert Systems—*games*

1. INTRODUCTION

[2, 1] presented the concept of fingerprinting, which turns arbitrary game-playing strategies into normal mathematical functions independent of representation by recording the strategy's average behaviour against a reference opponent. The model was updated in [4], improving upon several limitations; from [6], a metric has been defined on the space of fingerprints, which allows mathematical quantification of the *distance* between particular strategies.

Using the fingerprint metric, we investigate here a space of strategies that is just large enough to be used in actual studies, the 8-state automata. Since the space is too large to analyze exhaustively, we will use sampling and dimensionality techniques to consider the global structure (in the genotype space) imposed by the fingerprint distance (as phenotypic differences). For background on the fingerprint operator, see [4]; for the definition of the fingerprint distance and computational methodology, see [6].

2. EXPERIMENTAL DESIGN

We shall consider a deterministic 8-state finite transducer representation for playing iterated Prisoner's Dilemma, as a linear string of 34 integers: the transition (0–7) and the action (C/D), for 17 conditions — the initial move, and in each state, responding to the opponent's C or D. We take a sample of size 32,768 uniformly at random for analysis.

After combining all duplicates, there are 31,531 unique strategies in the sample; each was fingerprinted at $\alpha = 0.8$, a value found in previous studies to have good separation properties [5]. Pairwise distances are computed numerically and hierarchical clustering with the unweighted pair group with arithmetic mean method (UPGMA) is performed on this $31,531 \times 31,531$ distance matrix. We pick a level of 12,800 clusters and use that (weighted) distance matrix for analysis. Metric multidimensional scaling is used to embed these clusters into Euclidean space, with the stress majorization SMACOF algorithm [3], using the best fit chosen from over 1,000 runs starting at initial points i.i.d. uniformly random in $[0, 1]^n$. The algorithm is further repeated for embedding the data into anywhere from 1 to 10 dimensions.

3. RESULTS

The RMS error of embedding the distances, normalized by the RMS of the distances themselves (0.277029), is known as Kruskal's STRESS. Stress below 0.05 is considered good; our stress for embedding into $\mathbb{R}^2$ is 0.11043, into $\mathbb{R}^3$ is 0.05888, into $\mathbb{R}^5$ is 0.02789, into $\mathbb{R}^{10}$ is 0.02123. We conclude the data is mostly 5-dimensional, as further dimensions improve the embedding by a negligible amount.

We take the best-fit $\mathbb{R}^{10}$ MDS embedding of the distances and rotate it to principal coordinates; this is displayed as 12,800-point scatterplots in the top 3 components in Figure 1. A successful RGB colouring scheme, based on the expected single-round reactive strategy distribution used by the automata themselves, was developed in [6]; we will use this to colour each cluster.

The first major observation from the plots is that the colouring scheme is clearly reflected in the position of the points: there is a strong correlation with the colour green with the positive 1st component, red with negative 1st PC, blue with positive 3rd PC and black with negative 3rd PC. The two connected components in the 2nd component partition the strategies based on their initial move.

GECCO'14, July 12–16, 2014, Vancouver, BC, Canada.
ACM 978-1-4503-2881-4/14/07.
http://dx.doi.org/10.1145/2598394.2598498.

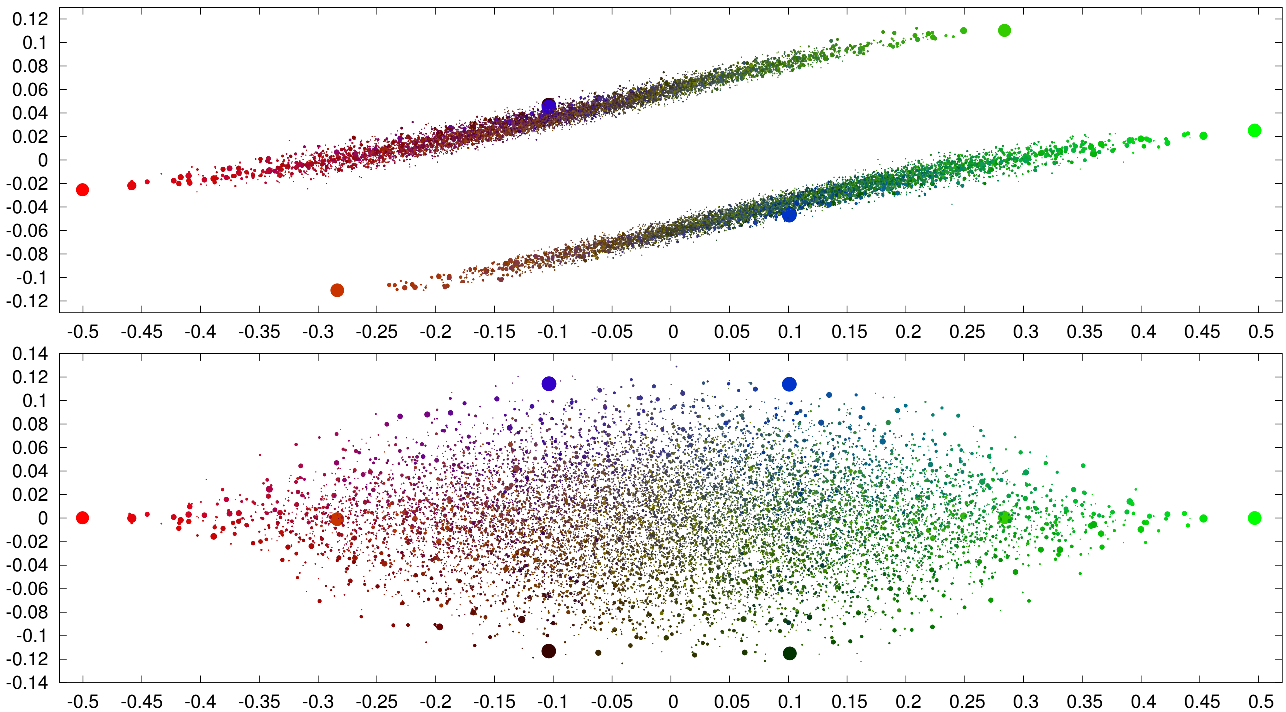

Figure 1: Scatter plot of all 32,768 strategies, reduced to 12,800 clusters with UPGMA hierarchical clustering, projected into 10 dimensions with metric MDS. Point size (area) is directly proportional to cluster size; for colouring see Section IV-B in [6]; axes are rotated to principal components, positive orientation is arbitrary. Top: plot of component 1 vs 2. Bottom: plot of component 1 vs 3.

Note also that the largest 8 clusters, which directly correspond to each of the 1-state strategies, form corners in the space. From the combination of the two figures, it is clear they are on the edges of the space, even if not in the dimensions reflected individually — this also means that the dimensions happen to separate different pairs of these strategies, further confirming their importance.

4. PREDICTING THE COORDINATES

Even though the representation is drastically larger, by comparison with [4] we can see a high similarity in structure. We thus test the hypothesis therein that the principal components in the embedded clusters correspond to *cooperativity* (probability of cooperating minus that of defecting), *responsiveness* (the correlation between your move and your opponent's last move) and *initialism* (difference in cooperativity in the first move vs. later moves).

We can write down predictors for the coordinate positions of the clusters using the colour components as defined in [6]. For cooperativity, we use the function Green − Red (bounded between -1 and 1), for responsiveness the predictor Blue − Black. As to initialism, we separate out the initial move, creating a fifth "colour" corresponding to whether the automaton cooperates on the initial move. A scaling factor of $(1+\alpha)(\text{Green}-\text{Red})\mp\alpha$ (depending on the initial move) was used to correct for the unequal weight of these moves.

Testing with the 10-D MDS embedding, the Pearson correlation (bounded in [-1,1]) between Green − Red and the first principal component of the points is 0.996790, between Blue − Black and the third component 0.881563, between the initialism predictor and the second component 0.991127. As 1 indicates a *perfect* linear relationship, these values are incredibly high; thus the predictors are *quantitative* explanations of the principal components.

5. REFERENCES

[1] D. Ashlock and E.-Y. Kim. Fingerprinting: Visualization and automatic analysis of Prisoner's Dilemma strategies. *IEEE Trans. Evol. Comput.*, 12(5):647–659, 2008.

[2] D. Ashlock, E.-Y. Kim, and W. K. von Roeschlaub. Fingerprints: enabling visualization and automatic analysis of strategies for two player games. In *2004 IEEE Congr. Evol. Comput.*, pages 381–387, 2004.

[3] J. de Leeuw. Applications of convex analysis to multidimensional scaling. In J. R. Barra, F. Brodeau, G. Romier, and B. van Cutsen, editors, *Recent Developments in Statistics*, pages 133–146. North-Holland Press, Amsterdam, Netherlands, 1977.

[4] J. Tsang. The parametrized probabilistic finite-state transducer probe game player fingerprint model. *IEEE Trans. Comput. Intel. AI Games*, 2(3):208–224, 2010.

[5] J. Tsang. The structure of a depth-3 lookup table representation for Prisoner's Dilemma. In *2010 IEEE Conf. Comput. Intel. Games*, pages 54–61, 2010.

[6] J. Tsang. The structure of a 3-state finite transducer representation for Prisoner's Dilemma. In *2013 IEEE Conf. Comput. Intel. Games*, pages 307–313, 2013.