
Reinforcement Learning in Factored Action Spaces using Tensor Decompositions

Anuj Mahajan^{*†} Mikayel Samvelyan[‡] Lei Mao[§] Viktor Makoviychuk[§]
Animesh Garg[§] Jean Kossaifi[§] Shimon Whiteson[†] Yuke Zhu[§]
Animashree Anandkumar[§]

Abstract

We present an extended abstract for the previously published work TESSERACT [Mahajan et al., 2021], which proposes a novel solution for Reinforcement Learning (RL) in large, factored action spaces using tensor decompositions. The goal of this abstract is twofold: (1) To garner greater interest amongst the tensor research community for creating methods and analysis for approximate RL, (2) To elucidate the generalised setting of factored action spaces where tensor decompositions can be used. We use cooperative multi-agent reinforcement learning scenario as the exemplary setting where the action space is naturally factored across agents and learning becomes intractable without resorting to approximation on the underlying hypothesis space for candidate solutions.

1 Introduction

Reinforcement learning (RL) has experienced tremendous advancements in recent years towards coming up with agents which can solve complex tasks and demonstrate generally performant intelligent behaviour [Silver et al., 2017, Vinyals et al., 2019, DeepMind-OEL et al., 2021]. One dimension in which task complexity in RL has seen big growth is that of size of the action space over which an agent (or a group of agents) have to make decisions. Thus, a lot of research efforts are being made towards creating approximation methods for learning in large action spaces which would otherwise be unnameable for classical methods. This has set the stage for development of innovative methods, analysis and addressal of new challenges which arise from the above approximation. Factorisation of the action space is one such strategy for overcoming computational intractability and learning approximately optimal policies. Factorisation involves decomposing/partitioning the action space into smaller sets so that RL can be done tractably on these smaller sets. This naturally provides a multi-agent perspective to the problem as each factored component can be seen as an agent working jointly with other such agent towards maximising the expected cumulative rewards. Factorisation can be imposed on an RL problem for tractability or it may arise from the problem/environment constraints such as communication limitations, partial observability and multi-agentness. In this work we will be focussing on product based factorisation of the action space \mathbf{U} that can be decomposed as cross product over action sets $\{U_i\}_1^n$ such that $\mathbf{U} = \bigotimes_{i=1}^n U_i$. This form of decomposition naturally occurs in many application like robotics where action space is multi-dimensional (ex. a single robot has a joint-policy for locomotion and communication) and swarm robotics Buşoniu et al. [2010] where each robot signifies a factor in decomposition.

^{*}Correspondence to anuj.mahajan@cs.ox.ac.uk

[†]University of Oxford

[‡]UCL

[§]NVIDIA

We will next see how TESSERACT [Mahajan et al., 2021] utilises the structure present in the above factored action spaces using a novel tensorised form of the Bellman equation. This novel form allows for creating algorithms for RL which can provide exponential gain in sample efficiency in the number of factors under the PAC framework. Further it opens up opportunities for using various forms of tensor decomposition for approximation based on the problem structure.

2 Background

Factored action spaces We consider the simplest setting for factored action space (FAS) where the state of the system is completely observed and is available for deciding an action from each component in the action space factorisation (we refer to this as an ‘agent’ from hereon, implying the factored problem has n agents, one for each factor). We refer the reader to Mahajan et al. [2021] for the more general scenarios. For the above scenario, a FAS can be modelled using a multi-agent Markov Decision Process (MMDP) which is defined as the tuple: $\langle S, U, P, r, n, \gamma \rangle$. Here, S is the state space of the environment with joint action space $\mathbf{U} = U^n$. Note that we have assumed without loss of generality that each factor/agent admits the same set of actions U ie. $U_i = U \forall i \in \mathcal{A} \equiv \{1, \dots, n\}$. A policy in FAS is the mapping $\pi : S \rightarrow \mathcal{P}(\mathbf{U})$ which gives a distribution over the joint action space given a state. At each time step t , an action $u^i \in U$ is chosen for every agent $i \in \mathcal{A}$ using π which forms the joint action $\mathbf{u} \in \mathbf{U} \equiv U^n$. The joint *action-value function* given a policy π is defined as: $Q^\pi(s_t, \mathbf{u}_t) = \mathbb{E}_{s_{t+1}, \mathbf{u}_{t+1}, \dots} [\sum_{k=0}^{\infty} \gamma^k r_{t+k} | s_t, \mathbf{u}_t]$. In RL, the goal is to find the optimal policy π^* corresponding to the optimal action value function Q^* .

Reinforcement Learning Methods RL methods come in a wide variety, but can fundamentally be classified as model based methods and model free methods. Model based methods typically estimate the underlying dynamics of the MDP (P, r) whereas model free methods implicitly account for them. Another important dimension of variance is the algorithmic setup used for learning which differentiates the methods as policy based, value based or actor-critic (a hybrid of previous two). Sutton and Barto [2011] gives a comprehensive overview of these methods. Value-based and actor-critic methods which can attributed most of the recent progress in RL for large action spaces, both rely on an estimator for the action-value function Q^π given a target policy π . Q^π satisfies the (scalar)-Bellman expectation equation: $Q^\pi(s, \mathbf{u}) = r(s, \mathbf{u}) + \gamma \mathbb{E}_{s', \mathbf{u}'} [Q^\pi(s', \mathbf{u}')]]$, which can equivalently be written in vectorised form as:

$$Q^\pi = R + \gamma P^\pi Q^\pi, \quad (1)$$

where R is the mean reward vector of size S , P^π is the transition matrix. The operation on RHS $\mathcal{T}^\pi(\cdot) \triangleq R + \gamma P^\pi(\cdot)$ is the Bellman expectation operator for the policy π . However, Eq. (1) doesn’t expose the structure present in FAS problems, in Section 3 we discuss how TESSERACT generalises the Bellman expectation equation (and analogously the Bellman optimality equation) to a novel tensor form suitable for sample efficient learning in FAS.

Tensor Decomposition Tensors are high dimensional analogues of matrices and tensor methods generalize matrix algebraic operations to higher orders. In the rest of this paper, we use $\hat{\cdot}$ to denote tensors. Formally, an order n tensor \hat{T} has n index sets $I_j, \forall j \in \{1..n\}$ and has elements $T(e), \forall e \in \times_{\mathcal{I}} I_j$ taking values in a given set \mathcal{S} , where \times is the set cross product and we denote the set of index sets by \mathcal{I} . Each dimension $\{1..n\}$ is also called a mode. An elegant way of representing tensors and associated operations is via tensor diagrams as shown in Fig. 1.

Tensor contraction generalizes the concept of matrix with matrix multiplication. For any two tensors \hat{T}^1 and \hat{T}^2 with $\mathcal{I}_\cap = \mathcal{I}^1 \cap \mathcal{I}^2$ we define the contraction operation as $\hat{T} = \hat{T}^1 \odot \hat{T}^2$ with $\hat{T}(e_1, e_2) = \sum_{e \in \times_{\mathcal{I}_\cap} I_j} \hat{T}^1(e_1, e) \cdot \hat{T}^2(e, e_2), e_i \in \times_{\mathcal{I}^i \setminus \mathcal{I}_\cap} I_j$. Using this building block, we can define tensor decompositions, which factorizes a (low-rank) tensor in a compact form. This can be done with various decompositions [Kolda and Bader, 2009, Janzamin et al., 2020], such as Tucker, Tensor-Train (also known as Matrix-Product-State), or CP (for Canonical-Polyadic). In this paper,




Figure 1: Left: Tensor diagram for an order 3 tensor \hat{T} . Right: Contraction between \hat{T}^1, \hat{T}^2 on common index sets I_2, I_3 .

we focus on the latter, which we briefly introduce here. Just as a matrix can be factored as a sum of rank-1 matrices (each being an outer product of vectors), a tensor can be factored as a sum of rank-1 tensors, the latter being an outer product of vectors. The number of terms in the sum is called the *CP-rank*. Formally, a tensor \hat{T} can be factored using a (rank- k) CP decomposition into a sum of k vector outer products (denoted by \otimes), as,

$$\hat{T} = \sum_{r=1}^k w_r \otimes^n u_r^i, i \in \{1..n\}, \|u_r^i\|_2 = 1.$$

3 Methodology

3.1 Tensorised Bellman equation

In this section, we provide the basic framework for Tesseract. We focus here on the discrete action space. The extension for continuous actions can be found in Mahajan et al. [2021]. Given a multi-agent problem $G = \langle S, U, P, r, n, \gamma \rangle$, let $\mathcal{Q} \triangleq \{Q : S \times U^n \rightarrow \mathbb{R}\}$ be the set of real-valued functions on the state-action space. We are interested in the *curried* Barendregt [1984] form $Q : S \rightarrow U^n \rightarrow \mathbb{R}, Q \in \mathcal{Q}$ so that $Q(s)$ is an order n tensor (We use functions and tensors interchangeably where it is clear from context). Algorithms in Tesseract operate directly on the curried form and preserve the structure implicit in the output tensor. (Currying in the context of tensors implies fixing the value of some index. Thus, Tesseract-based methods keep action indices free and fix only state-dependent indices.)

We are now ready to present the tensorised form of the Bellman equation shown in Eq. (1). Fig. 2 gives the equation where \hat{I} is the identity tensor of size $|S| \times |S| \times |S|$. The dependence of the action-value tensor \hat{Q}^π and the policy tensor \hat{U}^π on the policy is denoted by superscripts π . The novel **Tensorised Bellman equation** provides a theoretically justified foundation for the approximation of the joint Q -function, and the subsequent analysis (Theorems 1-3) for learning using this approximation.




Figure 2: **Tensorised Bellman Equation** for n agents. There is an edge for each factor in product decomposition/agent $i \in \mathcal{A}$ in the corresponding nodes $\hat{Q}^\pi, \hat{U}^\pi, \hat{R}, \hat{P}$ with the index set U^i .

3.2 TESSERACT Algorithms

It turns out that, constraining the underlying tensors for dynamics and rewards (\hat{P}, \hat{R}) is sufficient to bound the CP-rank of \hat{Q} (Theorem 1, Mahajan et al. [2021]). From this insight, a model-based RL version for TESSERACT can be constructed in Algorithm 1 (reproduced here from Mahajan et al. [2021]). The algorithm proceeds by estimating the underlying MDP dynamics using the sampled trajectories obtained by executing the behaviour policy $\pi = (\pi^i)_1^n$ satisfying Theorem 2. Specifically, we use a rank k approximate CP-Decomposition to calculate the model dynamics R, P as we show in Section 4. Next π is evaluated using the estimated dynamics, which is followed by policy improvement, Algorithm 1 gives the pseudocode for the model-based setting. The termination and policy improvement decisions in Algorithm 1 admit a wide range of choices used in practice in the RL community. Example choices for internal iterations which broadly fall under approximate policy iteration include: 1) Fixing the number of applications of Bellman operator 2) Using norm of difference between consecutive Q estimates etc., similarly for policy improvement several options can be used like ϵ -greedy (for Q derived policy), policy gradients (parametrized policy) Sutton and Barto [2011]

For large state spaces, where storage and planning using model parameters is computationally difficult, Mahajan et al. [2021] provide a model free version of the approach, details of which can be found in Appendix A with a sample of empirical results from the original paper in Appendix C. We next briefly revisit the main theoretical results from TESSERACT. Additional related works can be found in Appendix D.

4 Analysis

Under mild assumptions on the underlying dynamics of the FAS problem (see Appendix B for these along with additional constants appearing in Theorem 2), we have that the following results hold, where

k_1, k_2 are upper bounds on CP rank of the reward and transition tensor respectively (**proofs for which can be found in the original text, Mahajan et al. [2021]**):

Theorem 1. For a finite FAS the action-value tensor satisfies $\text{rank}(\hat{Q}^\pi(s)) \leq k_1 + k_2|S|, \forall s \in S, \forall \pi$.

The next PAC result empowers Algorithm 1 by lower bounding the number of samples required to infer the reward and state transition dynamics for finite MDPs with high probability using sufficient approximate rank $k \geq k_1, k_2$:

Theorem 2 (Model based estimation of \hat{R}, \hat{P} error bounds). Given any $\epsilon > 0, 1 > \delta > 0$, for a policy π with the policy tensor satisfying $\pi(\mathbf{u}|s) \geq \Delta$, where

$$\Delta = \max_s \frac{C_1 \mu_s^6 k^5 (w_s^{\max})^4 \log(|U|)^4 \log(3k \|R(s)\|_F / \epsilon)}{|U|^{n/2} (w_s^{\min})^4}$$

and C_1 is a problem dependent positive constant. There exists N_0 which is $O(|U|^{\frac{n}{2}})$ and polynomial in $\frac{1}{\delta}, \frac{1}{\epsilon}, k$ and relevant spectral properties of the underlying MDP dynamics such that for samples $\geq N_0$, we can compute the estimates $\bar{R}(s), \bar{P}(s, s')$ such that w.p. $\geq 1 - \delta$, $\|\bar{R}(s) - \hat{R}(s)\|_F \leq \epsilon, \|\bar{P}(s, s') - \hat{P}(s, s')\|_F \leq \epsilon, \forall s, s' \in S$.

Theorem 2 gives the relation between the order of the number of samples required to estimate dynamics and the tolerance for approximation. Theorem 2 states that aside from allowing efficient PAC learning of the reward and transition dynamics of the factored action-space MDP, Algorithm 1 requires only $O(|U|^{\frac{n}{2}})$ to do so, which is a vanishing fraction of $|U|^n$, the total number of joint actions in any given state. This also hints at why a tensor based approximation of the Q -function helps with sample efficiency. Methods that do not use the tensor structure typically use $O(|U|^n)$ samples. The bound is also useful for off-policy scenarios, where only the behaviour policy needs to satisfy the bound. Given the result in Theorem 2, it is natural to ask what is the error associated with computing the action-values of a policy using the estimated transition and reward dynamics. This is addressed in the next result, which bounds the error of model-based evaluation using approximate dynamics in Theorem 3. The first component on the RHS of the upper bound comes from the tensor analysis of the transition dynamics, whereas the second component can be attributed to error propagation for the rewards.

Theorem 3 (Error bound on policy evaluation). Given a behaviour policy π_b satisfying the conditions in Theorem 2 and executed for steps $\geq N_0$, for any policy π the model based policy evaluation $Q_{\bar{P}, \bar{R}}^\pi$ satisfies:

$$|Q_{\bar{P}, \bar{R}}^\pi(s, a) - Q_{\hat{P}, \hat{R}}^\pi(s, a)| \leq (|1 - f| + f|S|\epsilon) \frac{\gamma}{2(1 - \gamma)^2} + \frac{\epsilon}{1 - \gamma}, \forall (s, a) \in S \times U^n$$

where f is $\frac{1}{1 + \epsilon|S|} \leq f \leq \frac{1}{1 - \epsilon|S|}$

Algorithm 1 Model-based Tesseract

- 1: Initialise rank $k, \pi = (\pi^i)_1^n$ and \hat{Q}
 - 2: Initialise model parameters \hat{P}, \hat{R}
 - 3: Learning rate $\leftarrow \alpha, \mathcal{D} \leftarrow \{\}$
 - 4: **for** each episodic iteration i **do**
 - 5: Do episode rollout $\tau_i = \{(s_t, \mathbf{u}_t, r_t, s_{t+1})_0^L\}$ using π
 - 6: $\mathcal{D} \leftarrow \mathcal{D} \cup \{\tau_i\}$
 - 7: Update \hat{P}, \hat{R} using CP-Decomposition on moments from \mathcal{D}
 - 8: **for** each internal iteration j **do**
 - 9: $\hat{Q} \leftarrow \mathcal{T}^\pi \hat{Q}$
 - 10: **end for**
 - 11: Improve π using \hat{Q}
 - 12: **end for**
 - 13: Return π, \hat{Q}
-

5 Conclusions & Future Work

In this extended abstract we discussed the main ideas introduced in TESSERACT [Mahajan et al., 2021], a novel framework utilising the insight that the action value function for RL problems with factored action space can be seen as a tensor. TESSERACT provides a means for developing new sample efficient algorithms and obtain essential guarantees about convergence and recovery of the underlying dynamics. We also revisited the main theoretical results of TESSERACT. There are several interesting open questions to address in future work, such as convergence and error analysis for rank insufficient approximation, and analysis of the learning framework under different types of tensor decompositions like Tucker and tensor-train [Kolda and Bader, 2009].

References

- Anuj Mahajan, Mikayel Samvelyan, Lei Mao, Viktor Makoviychuk, Animesh Garg, Jean Kossaifi, Shimon Whiteson, Yuke Zhu, and Animashree Anandkumar. Tesseract: Tensorised actors for multi-agent reinforcement learning. In *Proceedings of the 38th International Conference on Machine Learning*, volume 139, pages 7301–7312. PMLR, 2021. URL <https://proceedings.mlr.press/v139/mahajan21a.html>.
- David Silver, Julian Schrittwieser, Karen Simonyan, Ioannis Antonoglou, Aja Huang, Arthur Guez, Thomas Hubert, Lucas Baker, Matthew Lai, Adrian Bolton, et al. Mastering the game of go without human knowledge. *nature*, 550(7676):354–359, 2017.
- Oriol Vinyals, Igor Babuschkin, Wojciech M Czarnecki, Michaël Mathieu, Andrew Dudzik, Junyoung Chung, David H Choi, Richard Powell, Timo Ewalds, Petko Georgiev, et al. Grandmaster level in starcraft ii using multi-agent reinforcement learning. *Nature*, 575(7782):350–354, 2019.
- DeepMind-OEL, Adam Stooke, Anuj Mahajan, Catarina Barros, Charlie Deck, Jakob Bauer, Jakob Sygnowski, Maja Trebacz, Max Jaderberg, Michael Mathieu, Nat McAleese, Nathalie Bradley-Schmieg, Nathaniel Wong, Nicolas Porcel, Roberta Raileanu, Steph Hughes-Fitt, Valentin Dalibard, and Wojciech Marian Czarnecki. Open-ended learning leads to generally capable agents. *arXiv preprint arXiv:2107.12808*, 2021.
- Lucian Buşoniu, Robert Babuška, and Bart De Schutter. Multi-agent reinforcement learning: An overview. In *Innovations in multi-agent systems and applications-1*, pages 183–221. Springer, 2010.
- Richard S Sutton and Andrew G Barto. Reinforcement learning: An introduction. 2011.
- Tamara G Kolda and Brett W Bader. Tensor decompositions and applications. *SIAM review*, 51(3): 455–500, 2009.
- Majid Janzamin, Rong Ge, Jean Kossaifi, and Anima Anandkumar. Spectral learning on matrices and tensors. *arXiv preprint arXiv:2004.07984*, 2020.
- Henk P Barendregt. Introduction to lambda calculus. 1984.
- Evan Greensmith, Peter L Bartlett, and Jonathan Baxter. Variance reduction techniques for gradient estimates in reinforcement learning. *Journal of Machine Learning Research*, 5(9), 2004.
- Tingting Zhao, Gang Niu, Ning Xie, Jucheng Yang, and Masashi Sugiyama. Regularized policy gradients: direct variance reduction in policy gradient estimation. In *Asian Conference on Machine Learning*, pages 333–348. PMLR, 2016.
- Richard S Sutton. Learning to predict by the methods of temporal differences. *Machine learning*, 3(1):9–44, 1988.
- Kyunghwan Son, Daewoo Kim, Wan Ju Kang, David Earl Hostallero, and Yung Yi. Qtran: Learning to factorize with transformation for cooperative multi-agent reinforcement learning. *arXiv preprint arXiv:1905.05408*, 2019.

- Mikayel Samvelyan, Tabish Rashid, Christian Schroeder de Witt, Gregory Farquhar, Nantas Nardelli, Tim GJ Rudner, Chia-Man Hung, Philip HS Torr, Jakob Foerster, and Shimon Whiteson. The Star-Craft Multi-Agent Challenge. In *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems*, 2019.
- Tabish Rashid, Mikayel Samvelyan, Christian Schroeder de Witt, Gregory Farquhar, Jakob Foerster, and Shimon Whiteson. QMIX: Monotonic Value Function Factorisation for Deep Multi-Agent Reinforcement Learning. In *Proceedings of the 35th International Conference on Machine Learning*, pages 4295–4304, 2018.
- Peter Sunehag, Guy Lever, Audrunas Gruslys, Wojciech Marian Czarnecki, Vinicius Zambaldi, Max Jaderberg, Marc Lanctot, Nicolas Sonnerat, Joel Z. Leibo, Karl Tuyls, and Thore Graepel. Value-Decomposition Networks For Cooperative Multi-Agent Learning Based On Team Reward. In *Proceedings of the 17th International Conference on Autonomous Agents and Multiagent Systems*, 2017.
- Yong Chen, Ming Zhou, Ying Wen, Yaodong Yang, Yufeng Su, Weinan Zhang, Dell Zhang, Jun Wang, and Han Liu. Factorized q-learning for large-scale multi-agent systems. *arXiv preprint arXiv:1809.03738*, 2018.
- Ming Tan. Multi-agent reinforcement learning: Independent vs. cooperative agents. In *Proceedings of the Tenth International Conference on Machine Learning*, pages 330–337, 1993.
- Anuj Mahajan, Tabish Rashid, Mikayel Samvelyan, and Shimon Whiteson. Maven: Multi-agent variational exploration. In *Advances in Neural Information Processing Systems*, pages 7611–7622, 2019.
- Carlos Guestrin, Michail Lagoudakis, and Ronald Parr. Coordinated reinforcement learning. In *ICML*, volume 2, pages 227–234. Citeseer, 2002a.
- Carlos Guestrin, Shobha Venkataraman, and Daphne Koller. Context-specific multiagent coordination and planning with factored mdps. In *AAAI/IAAI*, pages 253–259, 2002b.
- Eugenio Bargiacchi, Timothy Verstraeten, Diederik Roijers, Ann Nowé, and Hado Hasselt. Learning to coordinate with coordination graphs in repeated single-stage multi-agent decision problems. In *International conference on machine learning*, pages 482–490, 2018.
- Tarun Gupta, Anuj Mahajan, Bei Peng, Wendelin Böhmer, and Shimon Whiteson. Uneven: Universal value exploration for multi-agent reinforcement learning. *arXiv preprint arXiv:2010.02974*, 2020.
- Yaodong Yang, Jianye Hao, Ben Lu Liao, Kun Shao, Guangyong Chen, Wulong Liu, and Hongyao Tang. Qatten: A general framework for cooperative multiagent reinforcement learning. *ArXiv*, abs/2002.03939, 2020.
- Tonghan Wang, Tarun Gupta, Anuj Mahajan, Bei Peng, Shimon Whiteson, and Chongjie Zhang. Rode: Learning roles to decompose multi-agent tasks. *arXiv preprint arXiv:2010.01523*, 2020.
- Ryan Lowe, Yi Wu, Aviv Tamar, Jean Harb, OpenAI Pieter Abbeel, and Igor Mordatch. Multi-agent actor-critic for mixed cooperative-competitive environments. In *Advances in Neural Information Processing Systems*, pages 6379–6390, 2017.
- Jakob N Foerster, Gregory Farquhar, Triantafyllos Afouras, Nantas Nardelli, and Shimon Whiteson. Counterfactual multi-agent policy gradients. In *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.
- Anuj Mahajan and Theja Tulabandhula. Symmetry detection and exploitation for function approximation in deep rl. In *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems*, pages 1619–1621, 2017a.
- Anuj Mahajan and Theja Tulabandhula. Symmetry learning for function approximation in reinforcement learning. *arXiv preprint arXiv:1706.02999*, 2017b.
- Sham Machandranath Kakade. *On the sample complexity of reinforcement learning*. PhD thesis, UCL (University College London), 2003.

- Tor Lattimore, Marcus Hutter, and Peter Sunehag. The sample-complexity of general reinforcement learning. In *International Conference on Machine Learning*, pages 28–36. PMLR, 2013.
- Animashree Anandkumar, Rong Ge, Daniel Hsu, Sham M Kakade, and Matus Telgarsky. Tensor decompositions for learning latent variable models. *Journal of Machine Learning Research*, 15: 2773–2832, 2014.
- N. D. Sidiropoulos, L. De Lathauwer, X. Fu, K. Huang, E. E. Papalexakis, and C. Faloutsos. Tensor decomposition for signal processing and machine learning. *IEEE Transactions on Signal Processing*, 65(13):3551–3582, 2017.
- Yannis Panagakis, Jean Kossaifi, Grigorios G Chrysos, James Oldfield, Mihalis A Nicolaou, Anima Anandkumar, and Stefanos Zafeiriou. Tensor methods in computer vision and deep learning. *Proceedings of the IEEE*, 109(5):863–890, 2021.
- Nadav Cohen, Or Sharir, and Amnon Shashua. On the expressive power of deep learning: A tensor analysis. In *Conference on Learning Theory*, pages 698–728, 2016.
- Yu Cheng, Duo Wang, Pan Zhou, and Tao Zhang. A survey of model compression and acceleration for deep neural networks. *arXiv preprint arXiv:1710.09282*, 2017.
- Andrzej Cichocki, Anh-Huy Phan, Qibin Zhao, Namgil Lee, Ivan Oseledets, Masashi Sugiyama, Danilo P Mandic, et al. Tensor networks for dimensionality reduction and large-scale optimization: Part 2 applications and future perspectives. *Foundations and Trends® in Machine Learning*, 9(6): 431–673, 2017.
- Jean Kossaifi, Adrian Bulat, Georgios Tzimiropoulos, and Maja Pantic. T-net: Parametrizing fully convolutional nets with a single high-order tensor. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019.
- Jean Kossaifi, Antoine Toisoul, Adrian Bulat, Yannis Panagakis, Timothy Hospedales, and Maja Pantic. Factorized higher-order cnns with an application to spatio-temporal emotion estimation. In *IEEE CVPR*, 2020.
- Adrian Bulat, Jean Kossaifi, Georgios Tzimiropoulos, and Maja Pantic. Incremental multi-domain learning with network latent tensor factorization. 2020.
- Kamyar Azizzadenesheli, Alessandro Lazaric, and Animashree Anandkumar. Reinforcement learning in rich-observation mdps using spectral methods. *arXiv preprint arXiv:1611.03907*, 2016.
- Kamyar Azizzadenesheli. *Reinforcement Learning in Structured and Partially Observable Environments*. PhD thesis, UC Irvine, 2019.
- Akshay Krishnamurthy, Alekh Agarwal, and John Langford. Pac reinforcement learning with rich observations. In *Advances in Neural Information Processing Systems*, pages 1840–1848, 2016.
- Stefano Bromuri. A tensor factorization approach to generalization in multi-agent reinforcement learning. In *2012 IEEE/WIC/ACM International Conferences on Web Intelligence and Intelligent Agent Technology*, volume 2, pages 274–281. IEEE, 2012.




Figure 3: Tesseract architecture reproduced from original paper Mahajan et al. [2021]

A Model free Tesseract

For large state spaces where storage and planning using model parameters is computationally difficult (they are $\mathcal{O}(kn|U||S|^2)$ in number), Mahajan et al. [2021] propose a model-free approach using a deep network where the rank constraint on the Q -function is directly embedded into the network architecture. Fig. 3 gives the general network architecture for this approach and Algorithm 2 the associated pseudo-code. Each agent in Fig. 3 has a policy network parameterized by θ which is used to take actions in a decentralised manner. The observations of the individual agents along with the actions are fed through representation function g_ϕ whose output is a set of k unit vectors of dimensionality $|U|$ corresponding to each rank. The output $g_{\phi,r}(s^i)$ corresponding to each agent i for factor r can be seen as an action-wise contribution to the joint utility from the agent corresponding to that factor. The joint utility here is a product over individual agent utilities. For partially observable settings, an additional RNN layer can be used to summarise agent trajectories. The joint action-value estimate of the tensor $\hat{Q}(s)$ by the centralized critic is:

$$\hat{Q}(s) \approx T = \sum_{r=1}^k w_r \otimes^n g_{\phi,r}(s^i), i \in \{1..n\}, \quad (2)$$

where the weights w_r are learnable parameters exclusive to the centralized learner. In the case of value based methods where the policy is implicitly derived from utilities, the policy parameters θ are merged with ϕ . The network architecture is agnostic to the type of the action space (discrete/continuous) and the action-value corresponding to a particular joint-action ($u^1..u^n$) is the inner product $\langle T, A \rangle$ where $A = \otimes^n u^i$ (This reduces to indexing using joint action in Eq. (2) for discrete spaces). More representational capacity can be added to the network by creating an abstract representation for actions using f_η , which can be any arbitrary monotonic function (parametrised by η) of vector output of size $m \geq |U|$ and preserves relative order of utilities across actions; this ensures that the optimal policy is learnt as long as it belongs to the hypothesis space. In this case $A = \otimes^n f_\eta(u^i)$ and the agents also carry a copy of f_η during the execution phase. Furthermore, the inner product $\langle T, A \rangle$ can be computed efficiently using the property

$$\langle T, A \rangle = \sum_{r=1}^k w_r \prod_1^n \langle f_\eta(u^i) g_{\phi,r}(s^i) \rangle, i \in \{1..n\}$$

which is $O(nkm)$ whereas a naive approach involving computation of the tensors first would be $O(km^n)$. Training the Tesseract-based Q -network involves minimising the squared TD loss Sutton and Barto [2011]:

$$\mathcal{L}_{TD}(\phi, \eta) = \mathbb{E}_\pi[(Q(\mathbf{u}_t, s_t; \phi, \eta) - [r(\mathbf{u}_t, s_t) + \gamma Q(\mathbf{u}_{t+1}, s_{t+1}; \phi^-, \eta^-)])^2],$$

where ϕ^-, η^- are target parameters. Policy updates involve gradient ascent w.r.t. to the policy parameters θ on the objective $\mathcal{J}_\theta = \int_S \rho^\pi(s) \int_U \pi_\theta(\mathbf{u}|s) Q^\pi(s, \mathbf{u}) d\mathbf{u} ds$. More sophisticated targets can be used to reduce the policy gradient variance [Greensmith et al., 2004, Zhao et al., 2016] and propagate rewards efficiently [Sutton, 1988]. Note that Algorithm 2 does not require the individual-global maximisation principle [Son et al., 2019] typically assumed by value-based MARL methods in the CTDE setting, as it is an actor-critic method. In general, any form of function approximation and compatible model-free approach can be interleaved with Tesseract by appropriate use of the projection function Π_k .

Algorithm 2 Model-free Tesseract

- 1: Initialise rank k , parameter vectors θ, ϕ, η
 - 2: Learning rate $\leftarrow \alpha, \mathcal{D} \leftarrow \{\}$
 - 3: **for** each episodic iteration i **do**
 - 4: Do episode rollout $\tau_i = \{(s_t, \mathbf{u}_t, r_t, s_{t+1})_0^L\}$ using π_θ
 - 5: $\mathcal{D} \leftarrow \mathcal{D} \cup \{\tau_i\}$
 - 6: Sample batch $\mathcal{B} \subseteq \mathcal{D}$.
 - 7: Compute empirical estimates for $\mathcal{L}_{TD}, \mathcal{J}_\theta$
 - 8: $\phi \leftarrow \phi - \alpha \nabla_\phi \mathcal{L}_{TD}$ (Rank k projection step)
 - 9: $\eta \leftarrow \eta - \alpha \nabla_\eta \mathcal{L}_{TD}$ (Action representation update)
 - 10: $\theta \leftarrow \theta + \alpha \nabla_\theta \mathcal{J}_\theta$ (Policy update)
 - 11: **end for**
 - 12: Return π, \hat{Q}
-

B Assumptions used for analysis

The following assumptions are used in TESSERACT [Mahajan et al., 2021] which we reproduce here for reference:

Assumption 1. For the given MMDP $G = \langle S, U, P, r, n, \gamma \rangle$, the reward tensor $\hat{R}(s), \forall s \in S$ has bounded rank $k_1 \in \mathbb{N}$.

Intuitively, a small k_1 in Assumption 1 implies that the reward is dependent only on a small number of intrinsic factors characterising the actions.

Assumption 2. For the given MMDP $G = \langle S, U, P, r, n, \gamma \rangle$, the transition tensor $\hat{P}(s, s'), \forall s, s' \in S$ has bounded rank $k_2 \in \mathbb{N}$.

Intuitively a small k_2 in Assumption 2 implies that only a small number of intrinsic factors characterising the actions lead to meaningful change in the joint state. Assumption 1-2 always hold for a finite MMDP as CP-rank is upper bounded by $\prod_{j=1}^n |U_j|$, where U_j are the action sets.

Assumption 3. The underlying MMDP is ergodic for any policy π so that there is a stationary distribution ρ^π .

Next, we define coherence parameters, which are quantities of interest for our theoretical results: for reward decomposition $\hat{R}(s) = \sum_r w_{r,s} \otimes^n v_{r,i,s}$, let $\mu_s = \sqrt{n} \max_{i,r,j} |v_{r,i,s}(j)|$, $w_s^{\max} = \max_{i,r} w_{r,s}$, $w_s^{\min} = \min_{i,r} w_{r,s}$. Similarly define the corresponding quantities for $\mu_{s,s'}, w_{s,s'}^{\max}, w_{s,s'}^{\min}$ for transition tensors $\hat{P}(s, s')$. A low coherence implies that the tensor’s mass is evenly spread and helps bound the possibility of never seeing an entry with very high mass (large absolute value of an entry).

C Experiments

In this section we reproduce the main empirical results for TESSERACT on the StarCraft domain for the decentralised multi agent scenario. Complete experimental analysis and setup details can be found in the original paper. The experiments use the model-free version of TESSERACT (Algorithm 2).

StarCraft II A challenging set of cooperative scenarios from the StarCraft Multi-Agent Challenge (SMAC) [Samvelyan et al., 2019] is considered in this experiment. Scenarios in SMAC have been classified as **Easy, Hard and Super-hard** according to the performance of exiting algorithms on

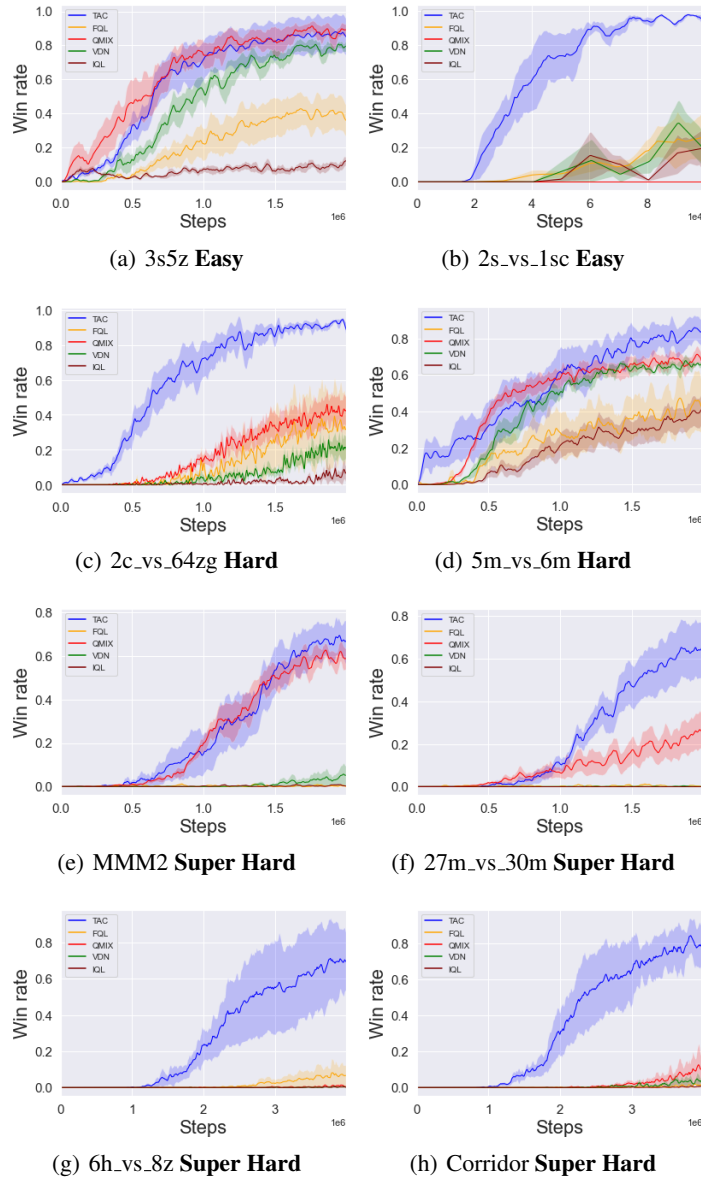


Figure 4: Performance of different algorithms on different SMAC scenarios: **TAC**, **QMIX**, **VDN**, **FQL**, **IQL**.

them. TESSERACT (**TAC** in plots) is compared with, **QMIX** [Rashid et al., 2018], **VDN** [Sunehag et al., 2017], **FQL** [Chen et al., 2018], and **IQL** [Tan, 1993], the latter being algorithms specifically designed for cooperative decentralised multi agent reinforcement learning. VDN and QMIX use monotonic approximations for learning the Q-function. FQL uses a pairwise factorized model to capture effects of agent interactions in joint Q-function, this is done by learning an inner product space for summarising agent trajectories. IQL ignores the multi-agentness of the problem and learns an independent per agent policy for the resulting non-stationary problem. Fig. 4 gives the win rate of the different algorithms averaged across five random runs. Fig. 4(c) features 2c_vs_64zg, a hard scenario that contains two allied agents but 64 enemy units (the largest in the SMAC domain) making the action space of the agents much larger than in the other scenarios. TESSERACT gains a huge lead over all the other algorithms in just one million steps. For the asymmetric scenario of 5m_vs_6m Fig. 4(d), TESSERACT, QMIX, and VDN learn effective policies, similar behavior occurs in the heterogeneous scenarios of 3s5z Fig. 4(a) and MMM2 Fig. 4(e) with the exception of VDN for the latter. In 2s_vs_1sc

in Fig. 4(b), which requires a ‘kiting’ strategy to defeat the spine crawler, TESSERACT learns an optimal policy in just 100k steps. In the **super-hard** scenario of 27m_vs_30m Fig. 4(f) having largest ally team of 27 marines, TESSERACT again shows improved sample efficiency; this map also shows TESSERACT’s ability to scale with the number of agents. Finally in the **super-hard** scenarios of 6 hydralisks vs 8 zealots Fig. 4(g) and Corridor Fig. 4(h) which require careful exploration, TESSERACT is the only algorithm which is able to find a good policy. It is observed that IQL doesn’t perform well on any of the maps as it doesn’t model agent interactions/non-stationarity explicitly. FQL loses performance possibly because modelling just pairwise interactions with a single dot product might not be expressive enough for joint-Q. Finally, VDN and QMIX are unable to perform well on many of the challenging scenarios possibly due to the monotonic approximation affecting the exploration adversely [Mahajan et al., 2019].

D Related Work

Previous methods for modelling multi-agent interactions include those that use coordination graph methods for learning a factored joint action-value estimation Guestrin et al. [2002a,b], Bargiacchi et al. [2018], however typically require knowledge of the underlying coordination graph. To handle the exponentially growing complexity of the joint action-value functions with the number of agents, a series of value-based methods have explored different forms of value function factorisation. VDN Sunehag et al. [2017] and QMIX Rashid et al. [2018] use monotonic approximation with latter using a mixing network conditioned on global state. QTRAN Son et al. [2019] avoids the weight constraints imposed by QMIX by formulating multi-agent learning as an optimisation problem with linear constraints and relaxing it with L2 penalties. MAVEN Mahajan et al. [2019] learns a diverse ensemble of monotonic approximations by conditioning agent Q -functions on a latent space which helps overcome the detrimental effects of QMIXs monotonicity constraint on exploration. Similarly, Uneven Gupta et al. [2020] uses universal successor features for efficient exploration in the joint action space. Qatten Yang et al. [2020] makes use of a multi-head attention mechanism to decompose Q_{tot} into a linear combination of per-agent terms. RODE Wang et al. [2020] learns an action effect based role decomposition for sample efficient learning. Policy gradient methods, on the other hand, often utilise the actor-critic framework to cope with decentralisation. MADDPG Lowe et al. [2017] trains a centralised critic for each agent. COMA Foerster et al. [2018] employs a centralised critic and a counterfactual advantage function. These actor-critic methods, however, suffer from poor sample efficiency compared to value-based methods and often converge to sub-optimal local minima. While sample efficiency has been an important goal for single agent reinforcement learning methods Mahajan and Tulabandhula [2017a,b], Kakade [2003], Lattimore et al. [2013], in this work we shed light on attaining sample efficiency for cooperative multi-agent systems using low rank tensor approximation.

Tensor methods have been used in machine learning, in the context of learning latent variable models Anandkumar et al. [2014], signal processing Sidiropoulos et al. [2017], deep learning and computer vision Panagakis et al. [2021]. They provide powerful analytical tools that have been used for various applications, including the theoretical analysis of deep neural networks Cohen et al. [2016]. Model compression using tensors Cheng et al. [2017] has recently gained momentum owing to the large sizes of deep neural nets. Using tensor decomposition within deep networks, it is possible to both compress and speed them up Cichocki et al. [2017], Kossaifi et al. [2019]. They allow generalization to higher orders Kossaifi et al. [2020] and have also been used for multi-task learning and domain adaptation Bulat et al. [2020]. In contrast to prior work on value function factorisation, TESSERACT provides a natural spectrum for approximation of action-values based on the rank of approximation and provides theoretical guarantees derived from tensor analysis. Multi-view methods utilising tensor decomposition have previously been used in the context of partially observable single-agent RL Azizzadenesheli et al. [2016], Azizzadenesheli [2019]. There the goal is to efficiently infer the underlying MDP parameters for planning under rich observation settings Krishnamurthy et al. [2016]. Similarly [Bromuri, 2012] use four dimensional factorization to generalise across Q-tables whereas here we use them for modelling interactions across multiple agents.