Inferring Implicit Trait Preferences from Demonstrations of Task Allocation in Heterogeneous Teams

Extended Abstract

Vivek Mallampati
Georgia Institute of Technology
Atlanta, GA, USA
vmallampati6@gatech.edu

Harish Ravichandar
Georgia Institute of Technology
Atlanta, GA, USA
harish.ravichandar@gatech.edu

ABSTRACT

Task allocation in heterogeneous teams often requires reasoning about multi-dimensional agent traits (i.e., capabilities) and the demands placed on them. However, existing methods tend to ignore the fact that not all traits equally contribute to a task. We propose an algorithm to infer implicit task-specific trait preferences in expert demonstrations. We leverage the insight that the consistency with which an expert allocates a trait to a task across demonstrations reflects the trait's importance to that task. Further, inspired by findings in psychology, we leverage the fact that a trait's inherent diversity among the agents controls the extent to which consistency informs preference. Through detailed numerical simulations and the `FIFA 20` soccer dataset, we demonstrate that we can infer implicit trait preferences, and accounting for them leads to more computationally efficient and effective task allocation.

KEYWORDS

Task allocation; Coalition formation; Preference learning

ACM Reference Format:
Vivek Mallampati and Harish Ravichandar. 2023. Inferring Implicit Trait Preferences from Demonstrations of Task Allocation in Heterogeneous Teams: Extended Abstract. In *Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023), London, United Kingdom, May 29 – June 2, 2023*, IFAAMAS, 3 pages.

1 INTRODUCTION

Heterogeneous multi-agent systems are effective in solving complex problems in a wide variety of domains. To effectively coordinate such heterogeneous teams, researchers have developed a variety of task allocation frameworks [2], ranging from market-based approaches that auction tasks to agents based on their utilities [1], to the recent advances in trait-based approaches that satisfy task requirements specified in terms of agent capabilities [6, 8].

However, existing task allocation approaches often consider all capabilities to be equally beneficial to a task, ignoring any underlying preferences [6, 8]. Ignorance of such differences can lead to sub-optimal task allocations that attempt to satisfy irrelevant requirements, placing a burden on the team's limited resources.

In this work, we model and infer task-specific preferences over traits in contexts requiring the satisfaction of multi-dimensional requirements. We argue and demonstrate that inferring and explicitly

Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023), A. Ricci, W. Yeoh, N. Agmon, B. An (eds.), May 29 – June 2, 2023, London, United Kingdom.

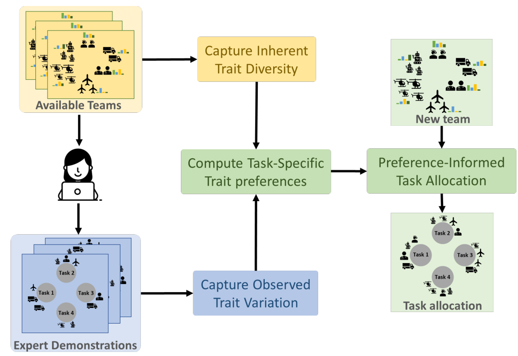

Figure 1: We infer implicit trait preferences in expert demonstrations to enable preferential task allocation.

accounting for such preferences will lead to more effective allocations, especially when capabilities are limited and task success requires surpassing minimum requirements.

We leverage two key ideas, both informed by psychology, to infer implicit trait preferences from expert demonstrations: i) the more important a trait, the most consistently it will be allocated by the expert, and ii) the abundance and diversity of a given capability in the dataset limit the extent of such consistency.

We evaluated our approach using extensive numerical simulations and a publicly-available soccer dataset (`FIFA 2020`) and compared its performance against that of a baseline that treats all traits equally. Our results conclusively demonstrate that inferring and accounting for trait preferences i) improves computational efficiency and allocation quality and ii) can significantly improve computational efficiency by eliminating irrelevant traits. Please see [5] for a full paper version of this extended abstract.

2 PROBLEM FORMULATION

Let us consider a heterogeneous team composed of S species (i.e., types of agents) with N_s agents in the s^{th} species. This heterogeneous team is required to perform M concurrent tasks denoted by $T = \{T_1, \dots, T_M\}$. See [5] for more details.

Let $\mathcal{D} = \{X^{(i)}, Q^{(i)}\}_{i=1}^{N}$ denote a set of demonstrations, where $X^{(i)} \in \mathbb{Z}_{\geq 0}^{M\times S}$ and $Q^{(i)} \in \mathbb{R}_{\geq 0}^{S\times U}$ denote the task assignment and the team's capability matrices, respectively, associated with the i^{th} demonstration. We assume that the expert performed these allocations by minimizing the following weighted cost function

$$X^{(i)} = \arg\min_{X} ||Y^* - XQ_i||_{W^*}, \forall i = 1, \cdots, N \tag{1}$$

where Y^* and $W^* \in \mathbb{R}_{\geq 0}^{M \times U}$ are unknown ground-truth trait requirements and the associated trait preferences of the expert.

Problem Statement: Given demonstrations $\mathcal{D}$, infer implicit trait requirements $\hat{Y}^*$ and the associated trait preferences $\hat{W}$.

3 INFERRING TRAIT PREFERENCES

We extract implicit trait preferences in demonstrations and use them in task allocation. We first compute two key quantities:

Observed Variation (CV_{obs}) is computed as the variation of each aggregated trait across the demonstrations for each task: $CV_{obs} = \sigma(Y_{\mathcal{D}})/\overline{Y}_{\mathcal{D}}$, where $Y_{\mathcal{D}} \in \mathbb{R}^{N \times M \times U}$ is the collection of all aggregated trait matrices in the demonstrations, $\sigma(Y_{\mathcal{D}})$ is the standard deviation of $Y_{\mathcal{D}}$, and $\overline{Y}_{\mathcal{D}}$ is the mean of $Y_{\mathcal{D}}$, both computed across demonstrations. As such, observed variation captures how consistently a trait is allocated to a particular task. Our insight here is that the importance of a trait to a task is inversely proportional to its variability in the demonstrations.

Inherent Diversity (CV_{div}) is computed as a function of the diversity of each trait in the demonstrations $\mathcal{D}$: $CV_{div} = \sigma(Q_{\mathcal{D}})/\overline{Q}_{\mathcal{D}}$, where $Q_{\mathcal{D}} \in \mathbb{R}^{N \times S \times U}$ is the collection of all species-trait matrices in the demonstrations, $\sigma(Q_{\mathcal{D}})$ is the standard deviation of $Q_{\mathcal{D}}$, and $\overline{Q}_{\mathcal{D}}$ is the mean of $Q_{\mathcal{D}}$. Our definition is inspired by the fact that the more diverse a dataset, the more informative it is. For instance, if all agents in the dataset possess near-identical speed, we cannot rely solely on the inevitable low observed variation in speed and conclude that it is highly preferred. Studies in psychology support this intuition by demonstrating that consistent selection and human preference are related [7] and that children learn preferences better with more diverse examples [3].

We finally compute the preference over traits by combining Observed Variation and Inherent Diversity:

$$\hat{W} = \frac{CV_{div}}{\tau} cos(\alpha CV_{obs} + \beta) + c \tag{2}$$

where α, β, τ and c are constants that control the mixing. Note that the above definition ensures that the lower values of the observed variation are generally associated with higher preferences and that lower values of inherent diversity increase the influence of observed variation. As such, low (high) inherent diversity results in small (large) influences of the observed variance on inferred preferences.

Once inferred, we incorporate the trait preferences into task allocation. Specifically, we compute the allocation $\hat{X}$ that minimizes the weighted objective $\|\hat{Y}^* - \hat{X}Q\|_{\hat{W}}$, where $\hat{Y}^*$ is the average trait requirement matrix over all demonstrations, and Q denotes the team capability matrix. Refer to [5] for a more detailed description.

4 EXPERIMENTAL EVALUATION

We evaluated the efficacy and need for our approach using the following experiments. See [5] for a comprehensive description.

4.1 Numerical Simulations

We first evaluated our approach using numerical simulations involving three tasks, three traits, and four classes of unknown ground-truth trait preference, each containing 1000 demonstrations. Notably, our approach was able to infer the correct order of importance among traits for all tasks across all cases.

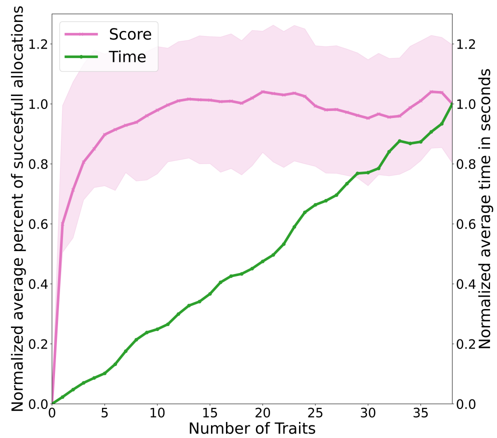

Figure 2: Inferring trait preferences significantly reduces dimensionality without sacrificing allocation quality.

We then compared our approach against two baselines (one ignores preference, and another ignores inherent diversity) in terms of allocation quality as measured by trait mismatch. We observed that our method results in statistically significantly lower trait mismatch than the baselines across all four classes of trait preferences.

4.2 Soccer Dataset

We also evaluated our approach on a soccer dataset, based on the FIFA 2020 game [4] involving 3600 players and 37 traits. We consider the four playing positions (forward, midfield, defense, and goalkeeper) as tasks and the players as agents. The allocation problem thus involves assigning positions to players. We ran experiments with 5-fold cross-validation with 40 teams.

Our results indicate that inferring and accounting for traits leads to a statistically significant increase in allocation success compared to the baseline that ignores preferences and considers all 34 traits equally. Note that, unlike the numerical simulations, we do not have access to the ground-truth preferences. Thus, we consider an allocation successful if a given player to allocated to their ground-truth playing position.

We also observed that our approach only requires 12 traits to match the performance of the baseline method that employs all 34 traits and is around 208% faster (see Fig. 2). Further, compromising on less than 5% of baseline accuracy can result in our method being 4X faster than the time taken by the baseline.

5 CONCLUSION

We demonstrated that we could successfully infer implicit preferences over traits from expert demonstrations based on how consistently traits appear in allocations and their diversity among agents. Our experiments reveal that inferring preferences and utilizing them in task allocation can improve allocation quality and computational efficiency.

REFERENCES

[1] Korsah Ayorkor, Stentz Anthony, and Dias Bernardine. 2013. A comprehensive taxonomy for multi-robot task allocation - G. Ayorkor Korsah, Anthony Stentz, M. Bernardine Dias,. https://journals.sagepub.com/doi/10.1177/0278364913496484

[2] Brian P. Gerkey and Maja J. Matarić. 2004. A Formal Analysis and Taxonomy of Task Allocation in Multi-Robot Systems. *The International Journal of Robotics Research* 23, 9 (Sept. 2004), 939–954. https://doi.org/10.1177/0278364904045564

[3] Tamar Kushnir, Fei Xu, and Henry M. Wellman. 2010. Young children use statistical sampling to infer the preferences of others. *Psychological science* 21, 8 (Aug. 2010), 1134–1140. https://doi.org/10.1177/0956797610376652

[4] Stefano Leone. [n.d.]. *FIFA 20 complete player dataset, Kaggle.* https://bit.ly/3Sd516U

[5] Vivek Mallampati and Harish Ravichandar. 2023. Inferring Implicit Trait Preferences for Task Allocation in Heterogeneous Teams. *arXiv preprint arXiv:2302.10817* (2023).

[6] Andrew Messing, Glen Neville, Sonia Chernova, Seth Hutchinson, and Harish Ravichandar. 2022. GRSTAPS: Graphically Recursive Simultaneous Task Allocation, Planning, and Scheduling. *The International Journal of Robotics Research* 41, 2 (Feb. 2022), 232–256. https://doi.org/10.1177/02783649211052066

[7] Sampo V. Paunonen. 1988. Trait Relevance and the Differential Predictability of Behavior. *Journal of Personality* 56, 3 (1988), 599–619.

[8] Harish Ravichandar, Kenneth Shaw, and Sonia Chernova. 2021. STRATA: A Unified Framework for Task Assignments in Large Teams of Heterogeneous Agents. *Journal of Autonomous Agents and Multi-Agent Systems (J-AAMAS)* (2021).