

A Measure of Heterogeneity in Multi-Agent Systems

Philip Twu, Yasamin Mostofi, and Magnus Egerstedt

Abstract—Heterogeneous multi-agent systems have previously been studied and deployed to solve a number of different tasks. Despite this, we still lack a basic understanding of just what “heterogeneity” really is. For example, what makes one team of agents more heterogeneous than another? In this paper, we address this issue by proposing a measure of heterogeneity. This measure takes both the complexity and disparity of a system into account by combining different notions of entropy. The result is a formulation that is both easily computable and makes intuitive sense. An overview is given of existing metrics for diversity found in various fields such as biology, economics, as well as robotics, followed by a discussion of their relative merits and demerits. We show how our proposed measure of heterogeneity overcomes problematic issues identified across the previous metrics. Finally, we discuss how to apply the new measure of heterogeneity specifically to multi-agent systems by using the notion of a common task-space to compare agents with different capabilities.

I. INTRODUCTION

When creating a multi-agent system, one must make two important design choices: determining the types of agents that will make up the system, and the algorithms, e.g., control laws, perception algorithms, and inter-agent communication protocols that those agents will execute. Over the past decade, a significant amount of effort has been spent researching the latter. As a result, controllers now exist which drive teams of autonomous agents to carry out tasks such as reaching consensus (e.g., [1], [2], [3], [4]), achieving formations (e.g., [5], [6], [7]), providing sensor coverage (e.g., [8], [9], [10]), and so on. However, recent advancements in the standardization of agent-based software frameworks (e.g., [11], [12], [13]), has given rise to “heterogeneous” systems comprising of agents with different capabilities (e.g., legged robots, UAVs, AUVs, sensor nodes).

Many of these heterogeneous systems exhibit different dynamics and capabilities compared to their homogeneous counterparts. It is of no surprise then, that many recent research projects have focused on understanding the dynamics of such heterogeneous systems and controlling them (e.g., [14], [15], [16], [17], [18]). However, one important piece missing from the current research in heterogeneous multi-agent systems is a universally agreed-upon definition of heterogeneity that answers fundamental questions such as:

This work is partially supported by the Office of Naval Research through the MURI, Heterogeneous Unmanned Network Teams (HUNT), and in part by the NSF CAREER award # 0846483.

P. Twu and M. Egerstedt are with the School of Electrical and Computer Engineering, Georgia Institute of Technology, Atlanta, GA 30332, USA. Y. Mostofi is with the Department of Electrical and Computer Engineering, University of California Santa Barbara, Santa Barbara, CA, 93106, USA. ptwu@gatech.edu, ymostofi@ece.ucsb.edu, magnus@ece.gatech.edu

- 1) What is heterogeneity?
- 2) How can heterogeneity be quantified?
- 3) How heterogeneous is a multi-agent system?

In this paper, we take a step back from current research efforts on heterogeneous systems and instead attempt to answer these three questions. *It is the authors’ belief that only after a definition of heterogeneity has been established, can its effects on a multi-agent system be uncovered systematically.*

II. WHAT IS HETEROGENEITY?

A. Existing Metrics for Heterogeneity

There have been many past attempts to quantify heterogeneity. For example, the Atkinson, Gini, and Theil indices described in [19] are commonly used in economics to measure the inequality of wealth distribution within a population. Biologists have also proposed a number of measures to describe the biodiversity within an ecosystem. Examples from [20] include the Shannon index (i.e., entropy), Simpson index, and Berger-Parker index, all of which quantifies heterogeneity based on the percentage of the total population belonging to each species. However, they do not take into account any differences between the species.

Rao’s quadratic entropy was presented in [21] as a means to measure biodiversity, while incorporating the differences between species. Given a metric defined on the set of available species, the quadratic entropy calculates the expected distance-squared between any two randomly chosen members of the population. However, [22] showed that the distribution which maximizes Rao’s quadratic entropy only assigns the population to the pairs of species which are furthest away from one another and therefore does not promote species diversity. Within the robotics community, hierarchic social entropy was proposed in [23] to describe robot diversity for a multi-robot system. Like Rao’s quadratic entropy, hierarchic social entropy takes into account both the percentage of the total population within each species, and the differences between the species. However, the metric involves solving multiple clustering problems and is therefore difficult to compute analytically.

It should be clear by now that although heterogeneity has been observed in many scientific fields, there lacks a general consensus as to its definition and how it should be quantified. For the remainder of this section, we will present a new definition of heterogeneity that relates it to two key characteristics of a system: complexity and disparity. This fundamental relationship between the three properties will be developed into a new metric for heterogeneity in the next section and will be shown to overcome problematic issues identified in existing metrics.

B. Complexity and Disparity

We will illustrate what characteristics of a system are required for heterogeneity through a simple example.

Example 2.1: Consider the task of picking handfuls of colored marbles out of a bag. For simplicity, let the differences amongst the types (or species) of marbles be given by the difference in their shade of darkness. Suppose the first batch of marbles picked (Group A) was all white, as shown in Figure 1(a). Despite lacking a concise definition of heterogeneity at this point, it is clear that Group A is not heterogeneous at all. Now, suppose two more batches (Groups B and C) of marbles are picked from the bag as shown in Figures 1(b) and 1(c) respectively. Both are clearly more heterogeneous than Group A, but in different ways.

Group B is more heterogeneous than Group A since its marbles are distributed evenly amongst three different types, whereas all the marbles in Group A are of the same type. We will describe how distributed the marbles are to different types as the group’s “complexity”. Looking now at Group C, we see that its marbles are only distributed evenly amongst two types and so it is more complex than Group A but less so compared to Group B. However, notice that while Group B has more types present, the colors of the marbles are only slight variations of one another. Compare this to what is seen in Group C where different types of marbles have very distinct colors. We will describe this notion of how distinct the marbles are from one another as the group’s “disparity”.

Using these terms, we see that the homogeneous Group A has neither complexity nor disparity. Group B is more heterogeneous than Group A because it exhibits higher complexity. On the other hand, Group C is more heterogeneous than Group A because it shows higher disparity. Complexity and disparity thus serve as two key characteristics that must be present in a heterogeneous system. However, oftentimes an increase in one may decrease the other. Therefore, to maximize heterogeneity, a system must strike a balance between the two oftentimes competing properties as shown in Group D of Figure 1(d). Based on these observations, we propose the following relationship:

$$\boxed{\begin{aligned} \text{HETEROGENEITY} = \\ \text{COMPLEXITY} \times \text{DISPARITY}. \end{aligned}} \quad (1)$$

C. Implications of Heterogeneity on Expressiveness

The purpose of this paper is not to advocate heterogeneity, but to simply understand and describe it. From an engineering perspective, it would be useful if one could relate it to the “expressiveness” of a system, i.e., the range of tasks which it can accomplish. However, note that expressiveness can vary independently of heterogeneity. Consider two homogeneous teams of agents, one well-equipped and the other not. Clearly, one is more expressive than the other. Unveiling the connection between heterogeneity and expressiveness, if one exists, is a potentially fruitful research direction.

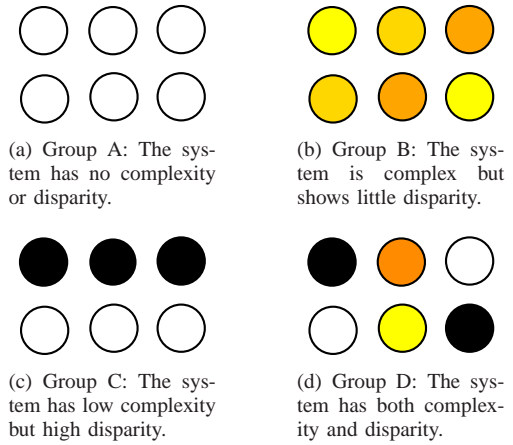


Fig. 1. Different groups of colored marbles used to illustrate the importance of having both complexity and disparity in a heterogeneous system.

III. HOW CAN HETEROGENEITY BE QUANTIFIED?

A. A New Measure of Heterogeneity

Assume that in a multi-agent system, each agent belongs to exactly one of M possible species (types). Let $p_i \in [0, 1]$ be the probability that a randomly chosen agent belongs to species i , for $i \in \mathcal{M} = \{1, \dots, M\}$, such that

$$\sum_{i=1}^M p_i = 1, \text{ and where } p = [p_1, \dots, p_M]^T$$

is the vector of probabilities that a randomly chosen agent will belong to each of the available species. Moreover, let

$$\mathcal{P}_M = \left\{ p \in \mathbb{R}^M \mid p_i \in [0, 1] \text{ for } i \in \mathcal{M}, \text{ and } \sum_{i=1}^M p_i = 1 \right\}$$

be the set of all probability distributions over M species.


Recall from the previous section that the complexity of a system describes how well spread out the agents are amongst the available species. This measure of disorder is exactly captured by the entropy of the system. Entropy gives the expected number of bits needed to describe which species an agent belongs to when using an optimal coding scheme. The precise definition of entropy is stated below.

Definition 3.1: Suppose that the M available species in a multi-agent system, as well as the probability distribution $p \in \mathcal{P}_M$ of agents belonging to each species, have both been established. The *entropy*¹, $E : \mathcal{P}_M \rightarrow \mathbb{R}_{\geq 0}$, of the multi-agent system is given by


$$E(p) = - \sum_{i=1}^M p_i \log(p_i). \quad (2)$$

Entropy is minimized when all the agents belong to the same species, and is maximized when the agents are evenly distributed across all species as illustrated in Figure 2(a). Note that entropy is purely a function of the percentage of agents belonging to each species.

¹Traditionally, in information theory, entropy is denoted by $H(p)$. However, in this paper, the letter H will instead be reserved for heterogeneity.



(a) The maximizing distribution p^* for entropy: $E(p)$.



(b) The maximizing distribution p^* for Rao's quadratic entropy $Q(p)$.

Fig. 2. Plots showing the maximizing distributions for entropy $E(p)$ and Rao's quadratic entropy $Q(p)$ in a multi-agent system with $M = 5$ uniform species, i.e., $d(i, j) = \alpha|i - j|$ for all $i, j \in \mathcal{M}$.

Next, we move on to quantifying a system's disparity. As was described in the previous section, disparity is a measure of how different agents are from one another. To compute such a quantity, it is necessary to first establish a metric amongst the set of all species. Disparity can then be computed as the expected distance between two randomly drawn agents by using Rao's quadratic entropy.

Definition 3.2: Suppose the entropy of a multi-agent system is well defined, and a metric $d : \mathcal{M} \times \mathcal{M} \rightarrow \mathbb{R}_{\geq 0}$ is established between the species. Rao's quadratic entropy, $Q : \mathcal{P}_M \rightarrow \mathbb{R}_{\geq 0}$, of the multi-agent system is given by

$$Q(p) = \sum_{i=1}^M \sum_{j=1}^M p_i p_j d(i, j)^2. \quad (3)$$

Note that just like with entropy, Rao's quadratic entropy is also minimized when all agents belong to the same species. The population distribution which maximizes Rao's quadratic entropy involves the agents being distributed only amongst the pairs of species which are the most distant from one another based on the metric d , as seen in Figure 2(b).

Having specified how to compute both the complexity and disparity of a system, heterogeneity can then be quantified using the relationship described in (1).

Definition 3.3: Suppose that the entropy $E(p)$ and Rao's quadratic entropy $Q(p)$ of a multi-agent system are both well-defined. The heterogeneity, $H : \mathcal{P}_M \rightarrow \mathbb{R}_{\geq 0}$, of a multi-agent system is given by

$$H(p) = E(p) Q(p). \quad (4)$$

B. Conservation of Species within Heterogeneity-Maximizing Populations

A major issue which prevented Rao's quadratic entropy from being used as a measure of biodiversity was that oftentimes, achieving the most diverse population required eliminating the population from some species. We will now show that our heterogeneity measure does not suffer from this problem. Before showing this key result, however, we must first establish that a heterogeneity-maximizing distribution even exists.

Lemma 3.1: There exists a probability distribution $p^* \in \mathcal{P}_M$ such that $H(p^*) = \sup_{p \in \mathcal{P}_M} H(p)$.

Proof: H is a continuous function over the compact set \mathcal{P}_M . Therefore, a maximizing $p^* \in \mathcal{P}_M$ must exist. ■

We are now ready to show a key result: that no species need to be eliminated in the most heterogeneous population according to H , i.e., H promotes species variety.

Theorem 3.1: Let $p^* \in \mathcal{P}_M$ be the probability distribution which maximizes H , then $p_i^* > 0$ for all $i \in \mathcal{M}$.

Proof: The proof of this theorem will be shown using induction for when the system contains k species. In the base case when only $k = 1$ species is available, the only way to distribute the population is to have $p_1^* = 1 > 0$. For the inductive hypothesis, assume that for $k = M$ species, the population distribution which maximizes H assigns a nonzero portion of the population to each species. We will show that the same holds with $k = M + 1$ species.

Let $p \in \mathcal{P}_M$, where $p_1, \dots, p_M > 0$, be any probability distribution for an arbitrarily chosen subset of M species out of the $M + 1$ that are available. We will show that H will always increase when some of the population from each of the M species is distributed to the $M + 1$ th species. Hence, the most heterogeneous distribution when $k = M + 1$ must have nonzero population in all $M + 1$ species. Start by defining a probability distribution $q(\epsilon) \in \mathcal{P}_{M+1}$, which is parameterized by $\epsilon \geq 0$, where

$$q_i(\epsilon) = p_i - \epsilon \text{ and } q_{M+1} = M\epsilon,$$

for $i = 1, \dots, M$, such that ϵ represents how much of the population from the first M species is being transferred to the $M + 1$ th species. The heterogeneity of the new distribution is $H(q(\epsilon)) = E(q(\epsilon)) Q(q(\epsilon))$, which is continuous for $\epsilon \in [0, \epsilon^*]$, where $\epsilon^* = \min\{\frac{1}{M}, p_1, \dots, p_M\}$ since $H(q(\epsilon))$ is only defined for $q(\epsilon) \in \mathcal{P}_{M+1}$. Taking the derivative with respect to ϵ reveals that $H(q(\epsilon))$ is continuously differentiable for $\epsilon \in (0, \epsilon^*)$.

Taking the right sided limit as ϵ goes to zero, we get that

$$\lim_{\epsilon \rightarrow 0^+} \frac{\partial H(q(\epsilon))}{\partial \epsilon} = \infty.$$

Therefore, there must be some $\hat{\epsilon} \in (0, \epsilon^*)$ for which $\frac{\partial H(q(\epsilon))}{\partial \epsilon}$ is positive for all $\epsilon \in (0, \hat{\epsilon})$ and hence, $H(p) < H(q(\hat{\epsilon}))$. Since any such $\hat{\epsilon}$ assigns a nonzero population to each of the $M + 1$ species, heterogeneity always can be increased by

distributing some of the population in the first M species to the $M+1$ th species. From this we see that when $k = M+1$, the heterogeneity-maximizing distribution involves having a nonzero population in each of the $M+1$ species. ■

C. Heterogeneity in Systems with Uniform Species

To gain better insight into what H considers the most heterogeneous population, we will temporarily focus on a special type of system. Suppose each species in a system is represented by its index, $1, \dots, M$, and the distance between two species is proportional to the absolute value of the difference between the two indices. We will refer to this situation collectively as the multi-agent system having uniform species.

Definition 3.4: A multi-agent system with *uniform species* consists of M available species, where the metric d between species is defined such that $d(i, j) = \alpha|i - j|$, for some constant $\alpha > 0$.

Although a system with uniform species is but a special case, it allows us to easily visualize population distributions in a simple scenario where not all species are equidistant. The intuition gained from analyzing this special case can then be applied to estimate what the maximally heterogeneous population distribution looks like in a general system.

We will start by showing that the most heterogeneous population distribution in a system with uniform species obeys a symmetry property.

Lemma 3.2: If p^* is the heterogeneity-maximizing distribution in a multi-agent system with uniform species, then it satisfies the symmetry property: $p_i^* = p_{M+1-i}^*$, for all $i \in \mathcal{M}$.

Proof: We will show that for every probability distribution p that does not satisfy the symmetry property, it is always possible to construct a distribution $q \in \mathcal{P}_M$, where

$$q_i = q_{M+1-i} = \frac{1}{2}(p_i + p_{M+1-i}), \text{ for all } i \in \mathcal{M},$$

such that $H(q) > H(p)$.

First, we show $E(q) > E(p)$ by noting that

$$E(q) = -\sum_{i=1}^M q_i \log(q_i) = -\sum_{i=1}^M p_i \log(q_i).$$

Suppose in the above expression we fix p and are allowed to choose $q \in \mathcal{P}_M$ to minimize $E(q)$. Using Lagrange multipliers to minimize $E(q)$ with the constraint that $g(q) = \sum_{i=1}^M q_i - 1 = 0$, it is simple to show that $\arg \min_{q \in \mathcal{P}_M} E(q) = p$. However, since $q \neq p$, we then conclude that $E(q) > E(p)$.

Next, we will show that $Q(q) \geq Q(p)$. Substituting q into the definition of Q and performing some simplifications, while keeping in mind that $d(i, j)^2 = d(M+1-i, M+1-j)^2$ because of uniform species,

$$Q(q) = \frac{1}{2}Q(p) + \frac{1}{2} \sum_{i=1}^M \sum_{j=1}^M p_i p_{M+1-j} d(i, j)^2.$$

Computing the difference between $Q(q)$ and $Q(p)$ gives

$$Q(q) - Q(p) = \alpha(2\epsilon_p - M - 1)^2 \geq 0,$$

where $\epsilon_p = \sum_{j=1}^M p_j j$ is the expected species index.

Because we have shown that $E(q) > E(p)$ and $Q(q) \geq Q(p)$, it follows that $H(q) > H(p)$. ■

Building on the symmetry result in the previous lemma, we will now derive a key relationship between the percentages of agents in each species within the most heterogeneous distribution.

Lemma 3.3: If p^* is the probability distribution which maximizes H in a multi-agent system with uniform species, then $p_i^* > p_k^* \iff i^2 - k^2 > (i - k)(M + 1)$.

Proof: The problem of finding the heterogeneity-maximizing distribution can be posed as a constrained maximization problem, where the goal is to maximize $H(p)$ with respect to the constraints: $g(p) = 0$ and $h_i(p) \geq 0$, for all $i \in \mathcal{M}$, such that

$$g(p) = \sum_{i=1}^M p_i - 1 \text{ and } h_i(p) = p_i,$$

for all $i \in \mathcal{M}$. The resulting Lagrange function is given by

$$\Lambda(p, \lambda, \mu) = H(p) + \lambda g(p) + \sum_{i=1}^M \mu_i h_i(p),$$

where $\lambda \neq 0$ and $\mu_i \geq 0$, for $i \in \mathcal{M}$, are Lagrange multipliers. However, since it was shown in Theorem 3.1 that $p_i^* > 0$ for all $i \in \mathcal{M}$, we know that all of the inequality constraints are inactive and so $\mu_i = 0$, for all $i \in \mathcal{M}$. Therefore, we can simplify the Lagrange function to:

$$\Lambda(p, \lambda) = H(p) + \lambda g(p).$$

The optimality conditions state that

$$\frac{\partial \Lambda(p, \lambda)}{\partial p_k} = \frac{\partial H(p)}{\partial p_k} + \lambda = 0, \text{ for all } k \in \mathcal{M}.$$

Calculating the partial derivative of $H(p)$ and substituting it into the previous expression, we get that for all $k \in \mathcal{M}$:

$$-(\log(p_k) + 1)Q(p) + 2E(p) \sum_{j=1}^M p_j d(j, k)^2 = -\lambda.$$

We will use the following shorthand notation for the summation in the expression above:

$$X_k^2 = \sum_{j=1}^M p_j d(j, k)^2.$$

Since λ is constant, the following holds for any $i, k \in \mathcal{M}$:

$$-\log(p_i)Q(p) + 2E(p)X_i^2 = -\log(p_k)Q(p) + 2E(p)X_k^2$$

Grouping the quantities together, we get the following optimality condition for each pair of species $i, k \in \mathcal{M}$:

$$\frac{2E(p)}{Q(p)}(X_k^2 - X_i^2) = \log\left(\frac{p_k}{p_i}\right).$$

Since the heterogeneity-maximizing distribution causes $H(p) > 0$ when $M \geq 2$, then $E(p) > 0$ and $Q(p) > 0$. Therefore, for the above expression to hold, it must be that

$$\text{sgn}(p_i - p_k) = \text{sgn}(X_i^2 - X_k^2), \quad (5)$$

for all $i, k \in \mathcal{M}$.

Calculating the term $X_i^2 - X_k^2$ for the special case of when the multi-agent system has uniform species gives

$$X_i^2 - X_k^2 = \alpha(i^2 - k^2 - 2(i - k)\epsilon_p).$$

However, because the heterogeneity maximizing probability distribution is symmetric by Lemma 3.2, we have that $\epsilon_p = \frac{M+1}{2}$. Substituting it back into the previous expression and combining this result with the expression in (5) gives

$$\text{sgn}(p_i - p_k) = \text{sgn}(i^2 - k^2 - (i - k)(M + 1)),$$

from which we conclude that

$$p_i > p_k \iff i^2 - k^2 > (i - k)(M + 1).$$

The previous lemma describes which species has more of the population distributed to it than others in the most heterogeneous population. Combined with the symmetry result of Lemma 3.2, they give a concise statement about the shape of the most heterogeneous population distribution in a system with uniform species.

Theorem 3.2: Let p^* be the probability distribution which maximizes the heterogeneity measure H in a multi-agent system with uniform species, then $p_i^* > p_{i+1}^*$ for all $i < \frac{M}{2}$, and $p_k^* = p_{M+1-k}^*$ for all $k \in \mathcal{M}$.

Proof: Letting $k = i + 1$ in Lemma 3.3 gives

$$p_i^* > p_{i+1}^* \iff i^2 - (i + 1)^2 > (i - (i + 1))(M + 1),$$

where upon simplifying the inequality on the right side,

$$p_i^* > p_{i+1}^* \iff \frac{M}{2} > i.$$

That shows the first part of the theorem. The second part of the theorem is simply a restatement of Lemma 3.2. ■

Theorems 3.1 and 3.2 together paint a picture of what the most heterogeneous population distribution looks like with uniform species. With this information, we can now compare the heterogeneity measure H with entropy and Rao's quadratic entropy. Figure 3 shows the heterogeneity maximizing distribution in a multi-agent system with uniform species, consisting of $M = 5$ species total. Notice that consistent with Theorem 3.1, the maximally heterogeneous distribution assigns agents to each species, unlike with Rao's quadratic entropy as seen in Figure 2(b). Furthermore, the heterogeneity-maximizing distribution does not assign the same number of agents to each species in all situations, unlike with entropy as seen in Figure 2(a). Instead, because the differences between species are taken into account, more agents are assigned to the species whose indices are closer to the two extremes: 1 and 5, as described by Theorem 3.2. Thus, the most heterogeneous distribution retains some of the population in all species, and strikes a balance amongst the system's complexity and disparity.

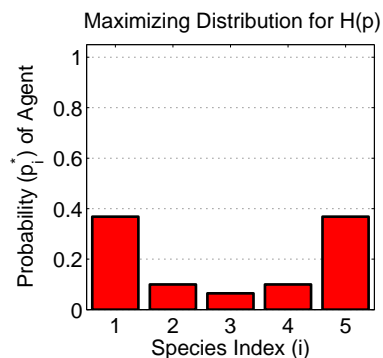


Fig. 3. Plot showing the heterogeneity maximizing distribution for H in a multi-agent system with uniform species, where $M = 5$. Notice that as stated in Theorem 3.1, all species have a nonzero portion of the total population assigned to it. Moreover, the shape of the distribution agrees with the description provided by Theorem 3.2.

IV. HOW HETEROGENEOUS IS A MULTI-AGENT SYSTEM?

The computation of H , or more specifically Q , in (4) requires that a metric d be established amongst the species. However, this choice may seem rather arbitrary when dealing with agents in a general setting. For example, what is the distance between a robot dog that can only wag its tail and a robot parrot that can only sing? Should the two be compared by their weight, size, vocal abilities, physical abilities, color, or something else? Clearly, the correct answer to this question depends on the task at hand. We therefore use the notion of a task-space to describe a set of tasks of interest. Such a set allows us to describe and compare the capabilities of agents from different species on common ground.

Definition 4.1: A task-space², (T, γ) , is a pair containing a non-empty countable set T of tasks, and a weight function $\gamma : T \rightarrow \mathbb{R}_{>0}$ for each task such that

$$\sum_{t \in T} \gamma(t) = 1.$$

The purpose of the weight function is to place greater emphasis on certain tasks over others. This could be for various reasons, such as if one task occurs more frequently than another. Accompanying a task-space is a task-map which describes what tasks each species of agents can accomplish.

Definition 4.2: A task-map, $\omega : \mathcal{M} \rightarrow 2^T$, associated with a multi-agent system and its task-space is a mapping from species index to a set of tasks.

These definitions establish a common ground for which the capabilities of different agent species can be described and compared on. To perform such a comparison, we present the notion of a functional distance as a measure of the difference in two species' capabilities with respect to a task space.

Definition 4.3: The functional distance between two species $i, j \in \mathcal{M}$, with respect to a task-space (T, γ) and

²We realize that this definition of task-space is somewhat different from the ones found in existing robotics literature.

task-map ω , is given by $\delta(T, \gamma, \omega) : \mathcal{M} \times \mathcal{M} \rightarrow [0, 1]$, where

$$\delta(T, \gamma, \omega)(i, j) = \frac{\sum_{t \in (\omega(i) \cup \omega(j)) \setminus (\omega(i) \cap \omega(j))} \gamma(t)}{\sum_{u \in \omega(i) \cup \omega(j)} \gamma(u)}. \quad (6)$$

Suppose $\gamma(t)$ describes the percentage of the mission which requires an agent to perform task t . Then the functional distance is simply the ratio of the number of tasks which only one of the two agents can accomplish, to the number of tasks that either can do. When two agents are capable of the exact same tasks, the functional distance is 0. Moreover, when two agents perform sets of tasks which are disjoint, the functional distance is 1.

To make this slightly more concrete, consider a $|T|$ -dimensional space where each coordinate represents one of the tasks and $|T|$ denotes the size of the set T . We next look at the representation of the agent capabilities in this task space. This visualization can help us better understand the heterogeneity of the agents. It can also be utilized to extend the definition of the functional distance to the case where the task-map is not binary, i.e., a node can do a task with a certain level of quality.

Let S_i for $i \in \mathcal{M}$ be a $|T| \times 1$ vector that represents the capabilities of agent i in the $|T|$ -dimensional task space. For instance, for the task-map of Definition 4.2, S_i will be zero everywhere except for the rows corresponding to the tasks that it can accomplish where it would be one. We can easily extend this representation to non-binary task-map cases. Let $\langle u, v \rangle_\Gamma$ and $\|u\|_\Gamma^2$ represent the weighted inner product and the weighted Euclidean norm-squared respectively, for arbitrary $|T| \times 1$ vectors u and v and a diagonal weight matrix Γ with the diagonal elements of $\gamma(t)$ for $t \in T$. Consider a mission that consists of n tasks. Then, the weighted distance-squared between s_i and s_j in the $|T|$ -dimensional task space, $\|s_i - s_j\|_\Gamma^2$, represents the ratio of the number of tasks (out of the total n tasks) that can be accomplished by only one of two agents i and j (and not both) during the mission. The functional distance of Definition 4.3 can then be expressed, based on the $|T|$ -dimensional space representation of tasks, as follows:

$$\delta(T, \gamma, \omega)(i, j) = \frac{\|s_i - s_j\|_\Gamma^2}{\|s_i\|_\Gamma^2 + \|s_j\|_\Gamma^2 - \langle s_i, s_j \rangle_\Gamma}. \quad (7)$$

Having established a distance measure with respect to a task space, we can compute the heterogeneity of a multi-agent system with respect to that space. Since this measure is based on the agents' ability to carry out tasks specific to the task-space, we refer to it as the system's functional heterogeneity.³

Definition 4.4: The functional heterogeneity $\mathcal{H}(p, T, \gamma, \omega)$ of a multi-agent system with respect to a task-space (T, γ) and task-map ω , is the heterogeneity $H(p)$ of the system using the functional distance $\delta(T, \gamma, \omega)$ as the species metric.

³As part of the future work, we envision that the representation and visualization of the agents in the $|T|$ -dimensional task space can be further utilized for characterizing heterogeneity in more general cases (e.g., case of a non-binary task-map).

REFERENCES

- [1] R. Olfati-Saber, J. A. Fax, and R. M. Murray, "Consensus and cooperation in networked multi-agent systems," *IEEE Proceedings*, vol. 95, no. 1, pp. 215–233, Jan. 2007.
- [2] W. Ren and R. W. Beard, "Consensus seeking in multiagent systems under dynamically changing interaction topologies," *IEEE Trans. Autom. Control*, vol. 50, no. 5, pp. 655–661, May 2005.
- [3] A. Jadbabaie, J. Lin, and A. S. Morse, "Coordination of groups of mobile autonomous agents using nearest neighbor rules," *IEEE Trans. Autom. Control*, vol. 48, no. 6, pp. 988–1001, Jun. 2003.
- [4] H. G. Tanner, A. Jadbabaie, and G. J. Pappas, "Flocking in fixed and switching networks," *IEEE Trans. Autom. Control*, vol. 52, no. 5, pp. 863–868, May 2007.
- [5] H. G. Tanner, G. J. Pappas, and V. Kumar, "Leader-to-formation stability," *IEEE Trans. Robot. Autom.*, vol. 20, no. 3, pp. 443–455, Jun. 2004.
- [6] A. Rahmani, M. Mesbahi, and F.Y. Hadaegh, "Optimal balanced-energy formation flying maneuvers," *Journal of Guidance Control and Dynamics*, vol. 29, no. 6, pp. 1395, 2006.
- [7] J.A. Fax and R.M. Murray, "Information flow and cooperative control of vehicle formations," *Automatic Control, IEEE Transactions on*, vol. 49, no. 9, pp. 1465–1476, 2004.
- [8] J. Cortes, S. Martinez, T. Karatas, and F. Bullo, "Coverage control for mobile sensing networks," *Robotics and Automation, IEEE Transactions on*, vol. 20, no. 2, pp. 243–255, 2004.
- [9] J.M. McNew, E. Klavins, and M. Egerstedt, "Solving coverage problems with embedded graph grammars," *Hybrid Systems: Computation and Control*, pp. 413–427, 2007.
- [10] V. de Silva and R. Ghrist, "Coverage in sensor networks via persistent homology," *Algebraic and Geometric Topology*, vol. 7, pp. 339–358, 2007.
- [11] Z. Kulis, V. Manikonda, B. Azimi-Sadjadi, and P. Ranjan, "The distributed control framework: a software infrastructure for agent-based distributed control and robotics," in *American Control Conference, 2008. IEEE, 2008*, pp. 1329–1336.
- [12] M. Quigley, B. Gerkey, K. Conley, J. Faust, T. Foote, J. Leibs, E. Berger, R. Wheeler, and A. Ng, "Ros: an open-source robot operating system," in *ICRA Workshop on Open Source Software*, 2009.
- [13] P. Martin, J.P. de la Croix, and M. Egerstedt, "Pancakes: A software framework for distributed robot and sensing network applications," in *10th International Symposium on Distributed Autonomous Robotics Systems*. IEEE, Nov. 2010.
- [14] B. Grocholsky, J. Keller, V. Kumar, and G. Pappas, "Cooperative air and ground surveillance," *Robotics & Automation Magazine, IEEE*, vol. 13, no. 3, pp. 16–25, 2006.
- [15] S. Balakirsky, S. Carpin, A. Kleiner, M. Lewis, A. Visser, J. Wang, and V.A. Ziparo, "Towards heterogeneous robot teams for disaster mitigation: Results and performance metrics from robocup rescue," *Journal of Field Robotics*, vol. 24, no. 11–12, pp. 943–967, 2007.
- [16] L. Pimenta, V. Kumar, R.C. Mesquita, and G. Pereira, "Sensing and coverage for a network of heterogeneous robots," in *Decision and Control, 2008. CDC 2008. 47th IEEE Conference on*. IEEE, 2008, pp. 3947–3952.
- [17] B.S. Smith, J. Wang, M. Egerstedt, and A. Howard, "Automatic formation deployment of decentralized heterogeneous multi-robot networks with limited sensing capabilities," in *Robotics and Automation, 2009. ICRA'09. IEEE International Conference on*. IEEE, 2009, pp. 730–735.
- [18] J. Kiener and O. Von Stryk, "Towards cooperation of heterogeneous, autonomous robots: A case study of humanoid and wheeled robots," *Robotics and Autonomous Systems*, vol. 58, no. 7, pp. 921–929, 2010.
- [19] F.A. Cowell, "Measurement of inequality," *Handbook of income distribution*, vol. 1, pp. 87–166, 2000.
- [20] R. Southwood and P.A. Henderson, *Ecological methods*, Wiley-Blackwell, 2000.
- [21] C.R. Rao, "Diversity and dissimilarity coefficients: A unified approach," *Theoretical Population Biology*, vol. 21, no. 1, pp. 24–43, 1982.
- [22] S. Pavoine, S. Ollier, and D. Pontier, "Measuring diversity from dissimilarities with Rao's quadratic entropy: Are any dissimilarities suitable?," *Theoretical Population Biology*, vol. 67, no. 4, pp. 231–239, 2005.
- [23] T. Balch, "Hierarchic social entropy: An information theoretic measure of robot group diversity," *Autonomous Robots*, vol. 8, no. 3, pp. 209–238, 2000.