The role of the environment in collective perception

Trendafilov, Dari; Almansoori, Ahmed; Carletti, Timoteo; Tuci, Elio

Published in:
ALIFE 2023: Ghost in the Machine

DOI:
10.1162/isal_a_00590

Publication date:
2023

The role of the environment in collective perception: A generic complexity measure

Dari Trendafilov, Ahmed Almansoori, Timoteo Carletti and Elio Tuci

Namur Institute for Complex Systems, University of Namur, Belgium
dari-borisov.trendafilov@unamur.be

Abstract

We propose a novel generic information-theoretic framework for characterizing the task difficulty in the Collective Perception paradigm. Our formalism builds on the notion of Empowerment - a task-independent, universal and generic utility function, which characterizes the level of perceivable control an embodied agent has over its environment. Series of simulations with an empowerment model of the collective perception scenario revealed a significant correlation between the levels of empowerment and the accuracy demonstrated by a set of standard collective decision-making strategies and a recent state-of-the-art neural network controller on nine benchmark patterns, used previously for assessing swarm performance. The results elucidate the key role of both the agent embodiment and the environmental pattern in characterising task difficulty, and justify the application of empowerment to analytically assess this role, which could help predict swarm performance and support the development of more efficient decision-making strategies.

Introduction

Swarm robotics studies multi-robot systems in which each robot has its own controller, perception is local and communication is based on spatial proximity (Hamann, 2018). The group-level response emerges from a self-organisation process (Camazine et al., 2001), based on the interaction between the robots and their physical environment. However, the autonomous nature of this process poses a challenge for designers, since it is notoriously difficult to infer which set of individual actions leads to the emergence of a desired collective response. Moreover, traditional design methods lack the ability to tackle problems and swarms of increasing complexity in uncertain and unpredictable environments. This further intensifies the need for fundamental and generic automated methodologies for modulating collective behaviour, with the potential to circumvent tedious trial-and-error model tuning. To enable large swarm-like robot collectives, generic measures of behavioural diversity could be highly beneficial. They could facilitate establishing the theoretical bounds on the complexities of individual robots, swarm and environment, and assessing their interactions and trade-offs. However, we argue that to date, the research community has not taken full advantage of what complexity measures can offer to swarm robotics.

We believe that information-theoretic approaches could address the above challenges in a generic fashion, by abstracting from implementation details and focusing on the interactions and dynamics related to information processing only (Roli et al., 2019). In this paper, we propose a novel generic measure for the characterization of task difficulty, based both on the environmental complexity and the robot's embodiment. To our knowledge, this is the first study to consider the agent's capabilities in the characterization of task difficulty in this domain. We apply this measure on a perceptual discrimination task, used in the swarm robotics literature to design mechanisms allowing swarms of robots to collectively decide which colour covers the majority of the arena floor. The aim is to provide a more rigorous theoretical treatment for the evaluation of the concept of task complexity, which could represent a useful metric to design progressively more effective solutions for swarm robotics. In our study, we explore the potential of the concept of empowerment to capture and predict the effect of different topological structures of the basic features on the task difficulty in the collective perceptual discrimination task. The information-theoretic concept of empowerment has been originally introduced by Klyubin et al. (2008), for providing a generic characterization of embodied agents and their environment. In order to validate our approach, we relate the empowerment levels to the performance achieved by four existing state-of-the-art decision-making strategies on the benchmark set of nine characteristic feature distribution patterns, illustrated in Figure 1 (see Bartashevich and Mostaghim, 2019; Almansoori et al., 2023).

Background

For designing large groups of robots, which coordinate and cooperatively perform a task, swarm robotics takes inspiration from natural self-organizing systems and attempts to recreate the emergence of collective behaviour from simple local interaction rules (see Kube and Zhang, 1993; Werfel et al., 2014). Through the design of individual robot

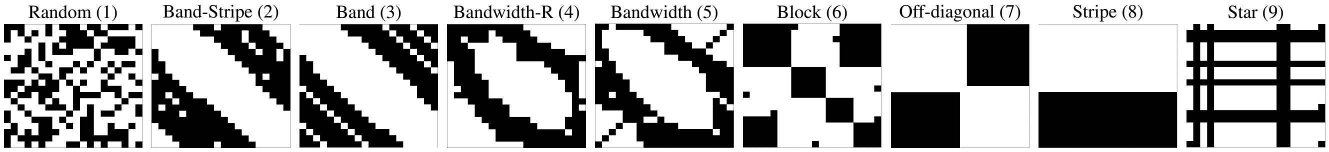

Figure 1: The layout of all nine test environmental conditions used in our study. Each environment is represented by a discrete 2-D grid of 20×20 cells. These bench-marking patterns have been proposed by Bartashevich and Mostaghim (2019) and employed in perceptual discrimination tasks (e.g., Almansoori et al., 2023) for assessing swarm performance and task difficulty.

behaviour, swarm robotics aims to achieve locally coordinated interaction that results in a self-organized collective behaviour (Ferrer et al., 2021; Boudet et al., 2021; Hasselmann et al., 2021). Robot collectives of fixed-size have been demonstrated in lab settings, typically employing a small number of robots (Nouyan et al., 2009). However, future swarms need to operate at many different scales.

To move away from hand-crafted solutions and designer-imposed bias, the multi-agent and robotics communities have recently turned their attention to information-theoretic measures. This is due to their ability to capture salient features of robot behaviour, based on generic information processing principles, while abstracting from system-specific details (see Roli et al., 2019). Information-theoretic approaches allow for a quantitative study of robot-environment systems (Beer, 1995; Smithers, 1995; Tarapore et al., 2006), and are fundamental in embodied systems research (Pfeifer and Scheier, 2001). Information theory has been used to formalise guided self-organization (Prokopenko, 2014; Ay et al., 2011; Polani et al., 2013; Prokopenko and Gershenson, 2014), relying on relatively simple local interactions from which complex global patterns emerge, stabilize and become more predictable as the information content decreases (see Polani, 2003, 2008; Prokopenko et al., 2006; Fernández et al., 2014). Generic information-theoretic complexity measures could capture non-linear relationships and have been used to study system dynamics (Lizier, 2013; Da Rold, 2018; Beer and Williams, 2015), to characterise information flows in the sensorimotor loop (Lungarella et al., 2005; Lungarella and Sporns, 2006), and to analyse robot behaviour (Roli et al., 2018). Shannon entropy-based measures, used to characterise self-organized emergent robot behaviour, range from mutual information (Salge and Polani, 2011; Sperati et al., 2014) and transfer entropy (Schreiber, 2000), to predictive and integrated information (Martius et al., 2013; Der et al., 2008; Balduzzi and Tononi, 2008). The potential of such measures, demonstrated by initial investigations, provides motivation for further exploration of information-theoretic approaches with respect to the automatic design of robot swarms and the analysis of their behavioural dynamics. One plausible option is the information-theoretic capacity of empowerment, introduced by Klyubin et al. (2008), which has previously been applied in various domains, e.g., to well-known problems in dynamical control systems (Jung et al., 2011), robotics (Salge et al., 2014), and human-computer interaction (Trendafilov and Murray-Smith, 2013). We believe that empowerment is a good candidate for providing a complexity measure in the perceptual discrimination task as illustrated in (Bartashevich and Mostaghim, 2019; Almansoori et al., 2023) for swarm robotics.

The collective perceptual discrimination task for swarm of robots has been originally introduced by Morlino et al. (2010), who used a binary version of this scenario to design and evaluate individual mechanisms underpinning the collective decision-making process. In this task, the swarm explored a close arena patched with tiles, randomly painted in black and white, with the aim to collectively decide which colour is dominant. In this task, the two colours are the options or features, and the proportion with which each colour covers the arena floor corresponds to the option/feature quality. The scientific challenge of this task is to develop individual opinion selection mechanisms that allow the swarm to converge on the desired consensus state (i.e., all robots selecting the opinion corresponding to the most represented colour on the arena floor). Various individual mechanisms for opinion selection have been developed, from the classical hand-crafted solutions, based on the voter model, the majority rule, and their variants (see Valentini, 2017), to more recent ones, based on the synthesis of artificial neural networks (Almansoori et al., 2023). The perceptual discrimination task has been used by Valentini et al. (2016) to investigate the performance of various decision-making strategies for swarm of robots while varying the options quality (i.e., the features ratio) for controlling task difficulty. Strobel et al. (2018) explored further variations of this task, characterised by the presence of byzantine robots, i.e., robots that communicate deceptive messages with the intent to entice the swarm to converge on a consensus to a non-optimal choice. Ebert et al. (2018) investigated scenarios with more than two options/features. While earlier research focused on the environmental feature ratio as a main characteristic of task difficulty, more recently, Bartashevich and Mostaghim (2019) suggested that the key determinant of the difficulty of the perceptual discrimination task for swarms of robots required to choose the best option (i.e., the most represented feature in the environment) is the features' distribution. Thus, Bartashevich and Mostaghim (2019) proposed a set of variations in the environmental topology and introduced measures for their characterization, however, disregarding robot's capabilities. To support and expand their work, we propose a universal and generic measure of task difficulty, which takes into account not only the environmental complexity (i.e., the features distribution), but also the agent's capabilities – arguably a key factor contributing to the overall swarm performance. We demonstrate the ability of the empowerment measure to quantify salient features of the environment, independent from the task or goal of the swarm, which makes our model directly applicable to further scenarios in this domain. Our study provides important insights about the generalizability of task difficulty with the proposed information-theoretic abstraction of the perceptual discrimination task.

Collective Perception

Our research is based on the collective perceptual discrimination task as described in (Bartashevich and Mostaghim, 2019; Almansoori et al., 2023), which is characterised by a square arena whose floor is covered by black and white tiles. In these studies, the most dominant colour, corresponding to the best quality option/feature, covers 55% of the arena floor, while the other colour covers the remaining 45%. Within each evaluation trial, the dominant colour can be either the black or the white. At the beginning of each trial, 20 robots are located in the arena without knowing whether the black or the white is the dominant colour. They are required to randomly explore the arena while perceiving the floor colour underneath their body, and by communicating their opinion on what is the dominant colour to spatially proximal robots. The objective of the swarm is to reach a consensus on which colour covers the largest proportion of the arena floor. Both in (Bartashevich and Mostaghim, 2019) and in (Almansoori et al., 2023), to induce significant variability in the task complexity (i.e., difficulty level), the mechanisms underpinning the collective decision-making process are evaluated on multiple environments that differ in the spatial distributions of the two features.

The most frequently used features' distribution in perceptual discrimination tasks is the random distribution of colour patches on the arena floor (see Figure 1/left most), which, however, has its limitations with respect to generalization of swarm behaviour; that is, decision-making strategies designed for randomly distributed patches are not equally successful in environments where features are distributed in a different way. In order to quantify the difference between various environmental patterns, Bartashevich and Mostaghim (2019) proposed a set of measures characterizing task difficulty, and validated them on the set of nine structurally different patterns shown in Figure 1. More recently, (Almansoori et al., 2023) used these patterns to evaluate the effectiveness of neural network-based decision-making mechanisms. Regardless of the nature of control mechanisms (hand-coded in Bartashevich and Mostaghim (2019), and neural network-based in Almansoori et al. (2023)), the results of the evaluation tests reveal that the swarm performance tends to deteriorate when the perceptual evidence is spatially arranged in distinctive clusters. In general, the less clustered the distribution of perceptual evidence, the higher the swarm accuracy in making the collective decision (see Bartashevich and Mostaghim, 2019; Almansoori et al., 2023). This observation can be explained by the potential emergence of an alignment in the opinions of spatially proximal robots exploring specific clusters, which hinders the swarm convergence to the correct option.

In this work, we draw inspiration from the above mentioned studies, by proposing an alternative way of measuring the task complexity with respect to the distribution of features characterising the collective perceptual discrimination task. Our view is grounded in the belief that the environmental complexity is intertwined with the robot's sensing and actuation abilities, and therefore must be taken into consideration when assessing task difficulty, i.e., a given task might be difficult for one type of robot and easier for another. Thus, for the characterization of task difficulty, our approach places a robot with a particular morphology into a specific environmental condition. Hence, we are interested in quantifying the complexity of the environment as perceived by the agent, which essentially depends on the agent's perception–action loop.

In order to carry out our task-complexity analysis, we make a few simplifications with respect to the original robot-based scenario as illustrated in (Bartashevich and Mostaghim, 2019; Almansoori et al., 2023). In particular, we consider a single agent placed in a discretized square grid of size of 20×20 cells. Each cell corresponds to a tile, that can be either black or white. The cells are coloured in order to recreate the nine features distribution patterns as shown in Figure 1. The agent can perceive the colour of the cell in which it is located and the colours of neighbouring cells. The number of the perceivable cells can vary from 5 (when the range is 1), to 61 (when the range is 5). The different neighbourhoods sizes with respect to the range are illustrated in Figure 2. The access to the colour of neighbouring cells intends to simulate the information generated by social influence. Within this metaphor, the different ranges model different values of the maximum robot–robot communication distance. There is, however, a direct mapping between our model and swarm robotics studies based on the e-puck2 platform, which has a communication range of 50 cm and the robots are placed in an environment of $2m \times 2m$, patched with black and white tiles, $10cm \times 10cm$ each (see, for example Almansoori et al., 2023). For the sake of simplicity, we follow the discrete grid boundaries in the definition of the neighbourhoods instead of using Euclidean distances, which would provide a better real-world representation in the form of smooth concentric circles, as we believe that such a loss

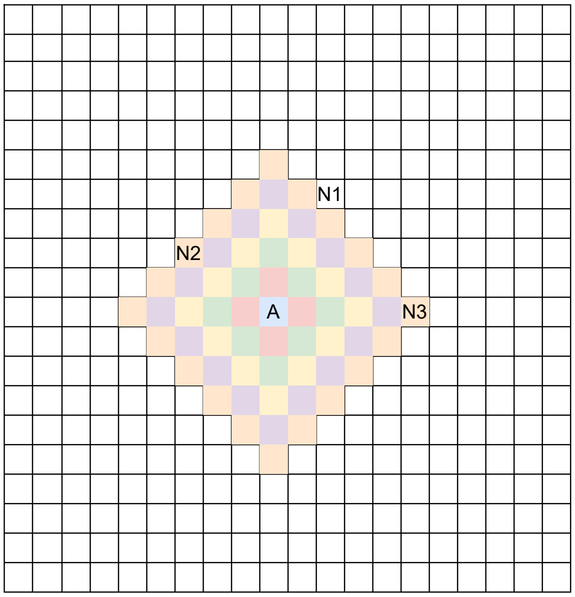

Figure 2: The experimental 2-D grid (20×20 cells) used in our study. The perceivable range of agent A is defined with the following colour map: range0 (blue), range1 (red), range2 (green), range3 (yellow), range4 (purple), and range5 (orange). E.g., the neighbours $N2$ and $N3$ are in range5, while $N1$ is out of range.

of resolution would not have a major impact on our results and our model is discrete by assumption.

To compute empowerment, for every feature distribution pattern and for each neighbourhood size (i.e., range), the agent is located in every cell of the grid. Thus, empowerment provides a measure of perceivable features with respect to the current position and range. For example, if the agent is placed on a black cell and all perceivable neighbouring cells are black, the empowerment will be zero (minimum empowerment value). Otherwise, if the agent is placed on black and at least one of the perceivable neighboring cells is white, then the empowerment will be one (maximum empowerment value). By computing this measure for all possible positions of the agent in the arena and for each of the nine feature distribution patterns, we obtain an estimate of the task complexity, which integrates both the environmental structure and the agent's sensory capabilities.

Empowerment Model

Our information-theoretic model of the Collective Perception paradigm is based on the empowerment formulation (Klyubin et al., 2008) of the perception–action loop of an embodied agent and its environment. According to that, the perception–action loop is represented as a communication channel in which when an agent performs an action, it injects information into the environment, and subsequently reacquires part of this information from the environment via its sensors. For stochastic dynamic systems in which transitions arise as the result of making a decision, empowerment measures the information an agent can inject into its environment and later perceive by its sensors. Formally, it is defined as the Shannon channel capacity from the sequence of actions $U_t, U_{t+1}, ..., U_{t+n-1}$ to the perception Y_{t+n} through the environment $X_{t+1}, X_{t+2}, ..., X_{t+n}$ after an arbitrary number of (n) time steps (Figure 3), using the following formula

$$C(U_t, ..., U_{t+n-1} \rightarrow Y_{t+n}) = \sup_{p(\vec{u})} I(U_t, ..., U_{t+n-1}; Y_{t+n})$$

where $\vec{u} = (u_t, ..., u_{t+n-1})$ and the mutual information between two discrete random variables U and Y is defined by

$$I(U;Y) = \sum_u p(u) \sum_y p(y|u) \log \frac{p(y|u)}{p(y)}. \tag{1}$$

Empowerment is a task and representation independent utility function, fully specified by the dynamics of the perception–action loop of the agent–environment coupling unrolled over time. It quantifies the maximal potential information flow from the agent's actuators to its sensors through the environment. Empowerment reflects the capacity of an agent to control or influence its environment as perceived by its sensors. It captures various sources and types of uncertainty (e.g., noise, delays, errors, etc.) in the perception–action loop in a single quantity, measured in source-independent uniform units (bits).

Intuitively, empowerment quantifies the number of actions available to the agent on a logarithmic scale, the outcome of which it can perceive. At its extremes, it is zero, if regardless of the action the outcome will be the same, and is maximal, if every action has a distinct outcome (Figure 4). Empowerment depends on the agent's embodiment, i.e., its sensoric apparatus and motoric abilities, and on the degree of interaction between agents, i.e., agents need freedom to act and at the same time they need certain constraints imposed by other agents (Capdepuy et al., 2007).

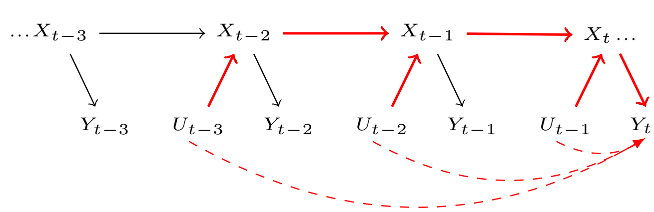

Figure 3: Perception–action loop as a causal Bayesian network – an agent performs an action U and injects information into the environment X, and subsequently reacquires part of this information via its sensors Y. Empowerment is the channel capacity from the action sequence to the resulting observation after n time-steps (e.g., from action sequence $U_{t-3}, U_{t-2}, U_{t-1}$ to perception Y_t).

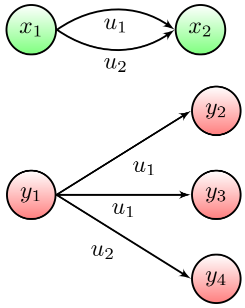

Figure 4: Transitions between perception states x and y when executing actions u. Intuitively, empowerment characterizes the number of actions available to the agent in a particular state, the outcome of which it can perceive. It is low, if regardless of the action the perception will be the same (in green), and it is high, if every action results in a distinct perception (in red). Note, that certain transitions can potentially be stochastic.

The decision-making mechanisms for collective perception are based on the agent's own perception and the opinions of its neighbours, which contain information about the environment at various remote positions and are transmitted from a distance within a specific communication range. This enables the agent to extend its sensing abilities and to acquire information about (perceive) the environment at distant locations. In a broader sense, from the agent's point of view, the environment is defined by the 2-D grid and the rest of the swarm. Since in this study, we use a single agent for the empowerment computation, in the following we will consider the environment to represent the 2-D grid only, and the swarm influence will be addressed in future work. Hence, in our model, the state space will consist of the position of a single agent in the grid. To translate the collective perception scenario to the empowerment formalism, we reframe the paradigm into a communication problem by using swarm communication as an action space and representing the action horizon with the communication range. For simplicity, we use only the main four orthogonal directions in our model, and build up a neighbourhood of a particular size, using an action space $\mathcal{U}$ consisting of the following five primitive actions

$$\mathcal{U} = \{north, south, east, west, idle\}.$$

The first four actions correspond to communicating with (polling the opinions of) the immediate neighbours in the four respective directions, while the last $(idle)$ action reflects the agent's own sensor reading, and n-step action sequences represent communication with agents in a neighbourhood of a particular range. The borders of the environment are hard and constrain the actions. Following this representation, Figure 2 depicts the perceivable range of agent A, in a blank 2-D grid, defined by a colour map – range0 (blue), range1 (red), range2 (green), range3 (yellow), range4 (purple), and range5 (orange).

Since we are interested in the overall task difficulty level, we compute empowerment in all positions across the grid, using the environmental features as sensor readings. For any state $x \in \mathcal{X}$ in the grid, empowerment is computed by

$$\mathfrak{E}(x) = \max_{p(\vec{u})} I(U_t, ..., U_{t+n-1}; Y_{t+n}|x),$$

where the action space $\mathcal{U}$ consists of the above five actions and the perception space $\mathcal{Y}$ is defined by a binary value

$$\mathcal{Y} = \{0, 1\},$$

representing the environmental features (black and white) in state $y \in \mathcal{X}$, where y is the resulting state after applying the action sequence $U_t, ..., U_{t+n-1}$ starting from x. Note, that x is a starting position on the 2-D grid and has two coordinates, while the perception $Y_{t+n} \in \mathcal{Y}$ is a single value representing the feature in the final position.

Results

Using the above model, we computed the empowerment levels for every starting position in the 2-D grid for all nine test environments (Figure 1), using a range of empowerment horizons from one to five, which corresponds to a discrete communication radius of one to five cells and is in line with previous swarm robotics studies in this scenario (see Almansoori et al., 2023).

The evolution of the empowerment levels, as the communication horizon increases, is presented in Figure 5 for one environmental pattern (Band(3)). It reveals the ability of empowerment to detect and characterize large homogeneous clusters in the environment with respect to the communication abilities (range) of the agent.

Since the agent can occupy any cell in the environment at any given moment, we averaged the empowerment levels across the grid and use the average values as characteristic for a particular experimental configuration. The average empowerment curves, for different horizons in all nine environmental conditions (Figure 6), are monotonically increasing with the horizon span, as expected, however, they reveal a significant variability between different environment types. The maximal empowerment level of 1 bit (since the agent can perceive at most two features and Equation 1 uses a binary logarithm), is reached for Random (1) in step 2 and for Star (9) in step 3, while Stripe (8) and Off-diagonal (7) have considerably lower empowerment at the presented horizon spans due to large homogeneous feature clusters. This insight emphasizes the key role of agent's abilities in characterizing task difficulty and predicting swarm performance.

The distribution of empowerment levels across the grid (Figure 7) reveals further details elucidating the significant difference between various patterns. The figure shows that

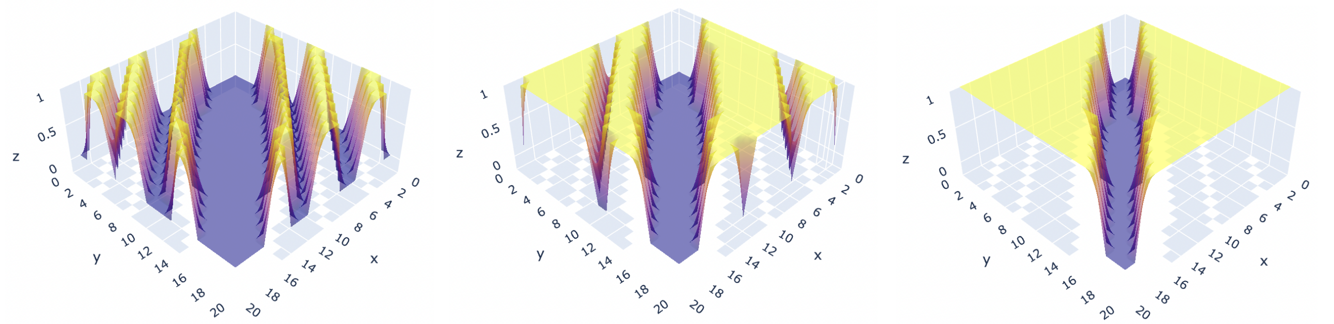

Figure 5: Empowerment maps for one- (left), two- (center) and three-step (right) horizons in one environmental condition. Empowerment (z) is presented as a function of the starting position in the grid (x-y coordinates). The pattern, corresponding to condition Band (3), is depicted in white and gray on the X-Y plane. On the left, the empowerment is highest around the transition edges between the two features and it is zero in homogeneous areas occupied by the same feature. As the horizon increases (center-right), it bridges relatively $narrow$ homogeneous areas, raising their empowerment level, while staying low in the $wider$ diagonal corridor.

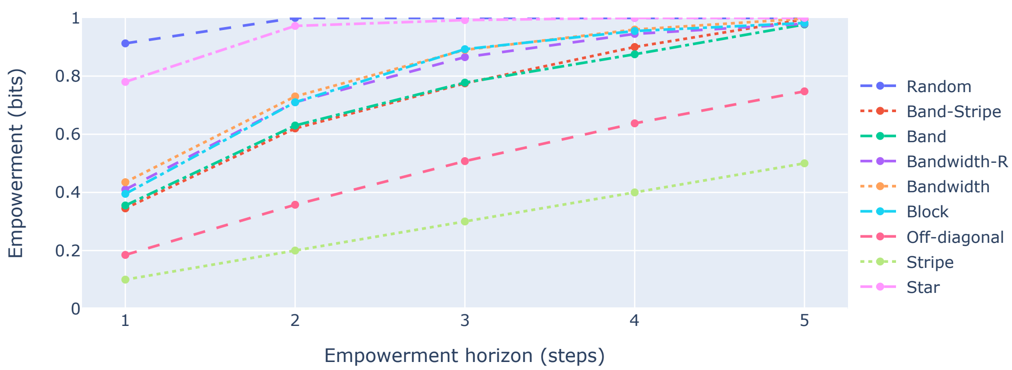

Figure 6: Empowerment curves, representing the average levels aggregated over the 2-D grid, for five discrete horizon spans in all nine environmental conditions. All trends are monotonically increasing, as expected, however with a significant variation between different environmental patterns. Note, that the maximal empowerment level for this scenario is 1 bit, which is reached for Random (1) in step 2 and for Star (9) in step 3. Large homogeneous feature clusters imply lower empowerment at these horizon spans for Stripe (8) and Off-diagonal (7).

the empowerment level is at maximum at the shortest horizon for condition Random (1), which is followed closely by Star (9). On the other extreme are conditions Off-diagonal (7) and Stripe (8), which do not converge overall to the maximum empowerment at this horizon range. These trends resonate well with the overall task difficulty levels measured empirically by swarm performance in previous studies.

In order to validate the proposed measure, we averaged the empowerment levels across the presented five horizons, since the agents communicate their opinion with neighbours that can be anywhere within the maximal communication range of five cells. We related the average values to the empirical swarm performance achieved by four existing state-of-the-art decision-making strategies – Voter model, Majority rule and Direct comparison, as reported by Bartashevich and Mostaghim (2019), as well as of an evolved neural network controller (NN), as reported by Almansoori et al. (2023) (Figure 8). We performed tests for correlation between the average empowerment and the accuracy of the above decision-making mechanisms. Due to the evolutionary bias in neural network controllers, leading to an obvious performance difference between black-dominant and white-dominant environments of the same type and ratio (Figure 8), we computed the correlation levels separately for

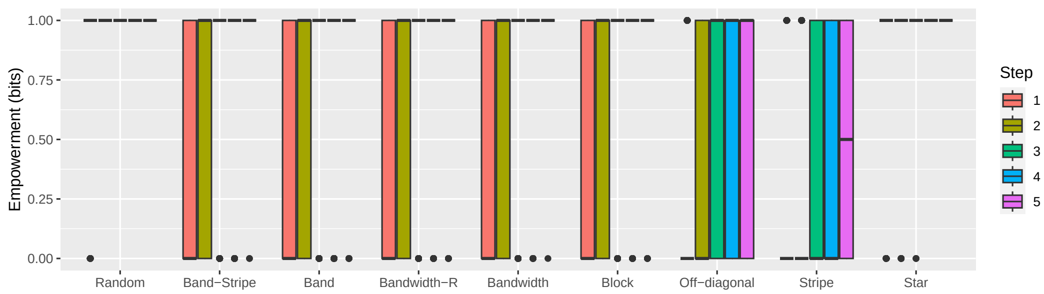

Figure 7: Box plots representing the empowerment distribution over the 2-D grid for horizon steps from one to five in all nine environmental conditions. The convergence to the maximal empowerment level of 1 bit is rather quick for conditions Random (1) and Star (9), and is much slower for conditions Off-diagonal (7) and Stripe (8). The bars reflect the balance between the number of 0 values and the number of 1 values, which the empowerment takes in this case.

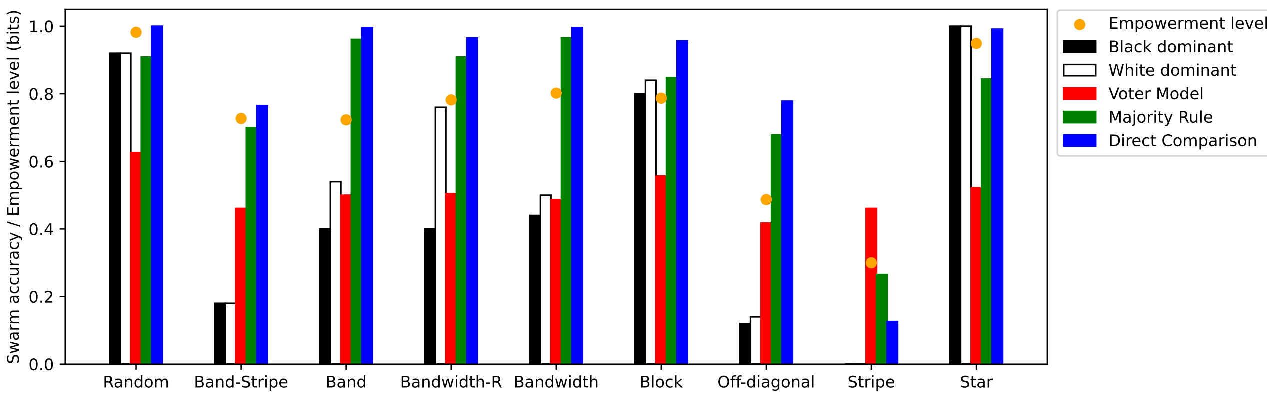

Figure 8: Average empowerment levels in all nine environmental conditions (in orange) and swarm accuracy of various collective decision-making strategies reported in previous work – an evolved neural network controller (Almansoori et al., 2023) (in black and white), and three standard hand-crafted solutions (Voter model, Majority rule and Direct comparison (Bartashevich and Mostaghim, 2019). Black and white bars stand for black-dominant and white-dominant environments respectively, which in the case of the neural controller have somewhat different accuracy due to evolutionary bias and are therefore analysed separately. The correlation between empowerment and accuracy is significant positive across strategies (Table 1).

NN black-dominant and NN white-dominant, as well as for their average performance. Pearson correlation analysis revealed significant positive correlation in all cases (Table 1), which suggests the relevance of the proposed measure for characterizing task difficulty and predicting swarm performance in the Collective Perception task. The only strategy with a lower confidence score is the Voter model.

Discussion

We have explored a novel approach for characterizing task difficulty, related to environmental topology, in the collective perception paradigm, based on a generic information-theoretic principle. We applied the empowerment formalism to derive a measure of environmental complexity relative only to the agent's perception-action abilities and independent of the task. Two key parameters influencing swarm performance in this scenario are the number of agents (i.e., agent density) and communication range. Our study investigated the effect of the communication range – represented by the empowerment horizon – on the proposed measure, which reveals monotonically increasing trends for all nine environments. Expanding the horizon increases empowerment, i.e., makes the task easier for the swarm at the expense of extended communication capabilities. The agent's embodiment is crucial in performing the task and essential in measuring task difficulty.

The results demonstrate a significant correlation between the empowerment levels and the accuracy of standard state-of-the-art decision-making strategies, which suggests the potential of the proposed measure to predict swarm performance based solely on properties of the environment and independent of the particular task. This makes the approach a good candidate for a universal complexity measure for supporting the design of robot swarms. The results reveal that

Empowerment vs.	ρ	$p - value$
NN black-dominant	0.86	0.003
NN white-dominant	0.88	0.002
NN average	0.88	0.002
Voter model	0.73	0.025
Majority rule	0.83	0.005
Direct comparison	0.86	0.003

Table 1: Pearson correlation between the average empowerment levels for all nine test environments and the empirical swarm accuracy using various standard collective decision-making strategies. Due to evolutionary bias in neural network controllers (Figure 8) we applied the test separately for NN black-dominant and NN white-dominant, as well as for the average NN performance. The rest of the strategies are bias-free.

the proposed method is able to quantify salient features of the environment from the agent's point of view, grounded in the embodied agent's morphology. Our approach does not explore the topological structure of the environment, e.g., number, size, shape and inter-connectivity of clusters, but instead, it explores the environment with the given agent morphology, which is critically relevant in determining task difficulty. Methods, quantifying task difficulty, based only on environmental features, as in (Bartashevich and Mostaghim, 2019), disregard the importance of the agent's capabilities for solving the task. Empowerment captures in a single measure salient features of the agent–environment perception-action loop, such as topology, morphology, noise in the sensing, actuation and communication channels, and it does so with a generic information-theoretic model. The empowerment curves (Figure 6), on the one hand, could be used for inferring the horizon span corresponding to a particular level of task difficulty, which could then serve as a predictor for swarm performance. The empowerment maps (Figure 5), on the other hand, could reveal critical salient points in the environment, which might inflict a significant drop in swarm performance and potentially raise designer's attention for a more careful consideration.

The key benefits of the applied information-theoretic treatment are that it is universal, general and could enable the analytical comparison of scenarios with different computational models. The proposed approach elucidates the trade-off between task difficulty (and swarm performance) and the cost of enabling particular agent capacities, and provides information-theoretic bounds, which are fundamental properties of agent–environment systems. Such theoretical bounds could provide guidelines and benchmarks for the evolution of optimal controllers by analytically evaluating task difficulty with a universal measure.

We believe that theories and tools from complex systems and information theory can successfully be applied for facilitating the automated design of robot collectives and for the analysis of their dynamics. Combining complexity measures with task-specific objective functions could enhance the swarm adaptivity and re-calibration in cases of environmental and task variations, and enable modulating the swarm complexity to the specific requirements. Generic complexity measures, such as empowerment, could provide a general-purpose tool for minimising designer-imposed bias. Leveraging classical Shannon's information theory by way of creating generative mathematical models and artificial simulations, empowerment offers a novel perspective for evolutionary swarm robotics, building on objective quantitative measures and analytical tools, which could support the automated design of robot swarms. The proposed formalism could contribute to the design of hybrid systems, combining model-free and model-based approaches, and offering a universal methodology that scales across domains.

Conclusion

This paper introduced a generic information-theoretic characterization of the environmental complexity in the Collective Perception paradigm for a homogeneous swarm consisting of robots with a particular morphology. It demonstrated an application of the information-theoretic capacity of empowerment to the field of swarm robotics and highlighted the benefits of utilising such a generic utility measure for predicting swarm performance. The proposed approach is task-independent, and the same model could be applied to various scenarios in this domain, e.g., shortest-path or site-selection problems. The results have shown a strong correlation between our measure and swarm performance, using standard strategies in this field. This suggests the potential of the approach in providing an analytical tool for making predictions and providing theoretical bounds, based solely on properties of the environment. Such a tool could support existing design frameworks in performance tuning, before resorting to costly empirical studies. Building on this formalism, future research will investigate the impact of various agent morphologies on the task difficulty level for specific environmental conditions. Another research direction is to extend and explore the formalism for heterogeneous swarms and for environments with multiple options. The results presented here are important for raising the awareness of the research community about the potential empowerment has in providing better theoretical foundations for the field of swarm robotics.

Acknowledgements

This project has received funding from the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No 101034383 and CERUNA doctoral fellowship by UNamur.

References

Almansoori, A., Alkilabi, M., and Tuci, E. (2023). Further investigations on the characteristics of neural network based opinion selection mechanisms for robotics swarms. *Proceedings of EvoStar – 26th European Conference on Applications of Evolutionary Computation*, pages 737–750.

Ay, N., Der, R., and Prokopenko, M. (2011). Guided self-organization: Perception-action loops of embodied systems. *Theory in biosciences*, 131:125–7.

Balduzzi, D. and Tononi, G. (2008). Integrated information in discrete dynamical systems: Motivation and theoretical framework. *PLOS Computational Biology*, 4(6):1–18.

Bartashevich, P. and Mostaghim, S. (2019). Benchmarking collective perception: New task difficulty metrics for collective decision-making. In *Progress in Artificial Intelligence: 19th EPIA Conference on Artificial Intelligence, EPIA 2019, Proceedings, Part I*, page 699–711.

Beer, R. D. (1995). A dynamical systems perspective on agent-environment interaction. *Artificial Intelligence*, 72(1):173–215.

Beer, R. D. and Williams, P. L. (2015). Information processing and dynamics in minimally cognitive agents. *Cognitive Science*, 39(1):1–38.

Boudet, J. F., Lintuvuori, J., Lacouture, C., Barois, T., Deblais, A., Xie, K., Cassagnere, S., Tregon, B., Brückner, D. B., Baret, J. C., and Kellay, H. (2021). From collections of independent, mindless robots to flexible, mobile, and directional superstructures. *Science Robotics*, 6(56).

Camazine, S., Deneubourg, J., Franks, N., Sneyd, J., Theraulaz, G., and Bonabeau, E. (2001). *Self-Organization in Biological Systems.* Princeton University Press, United States.

Capdepuy, P., Nehaniv, C., and Polani, D. (2007). Maximization of potential information flow as a universal utility for collective behaviour. In *Proceedings of the First IEEE Symposium on Artificial Life*, pages 207–213.

Da Rold, F. (2018). Information-theoretic decomposition of embodied and situated systems. *Neural Networks*, 103:94–107.

Der, R., Güttler, F., and Ay, N. (2008). Predictive information and emergent cooperativity in a chain of mobile robots. *Artificial life XI, 166-172*.

Ebert, J., Gauci, M., and Nagpal, R. (2018). Multi-feature collective decision making in robot swarms. In *17th International Conference on Autonomous Agents and Multiagent Systems, AAMAS 2018*, pages 1711–1719.

Fernández, N., Maldonado, C., and Gershenson, C. (2014). Information measures of complexity, emergence, self-organization, homeostasis, and autopoiesis. In Prokopenko, M., editor, *Guided Self-Organization: Inception*, pages 19–51. Springer Berlin Heidelberg.

Ferrer, E. C., Hardjono, T., Pentland, A., and Dorigo, M. (2021). Secure and secret cooperation in robot swarms. *Science Robotics*, 6.

Hamann, H. (2018). *Swarm Robotics: A Formal Approach*. Springer.

Hasselmann, K., Ligot, A., Ruddick, J., and Birattari, M. (2021). Empirical assessment and comparison of neuro-evolutionary methods for the automatic off-line design of robot swarms. *Nat Commun*, 12(4345):1–11.

Jung, T., Polani, D., and Stone, P. (2011). Empowerment for continuous agent-environment systems. *Adaptive Behavior*, 19(1):16–39.

Klyubin, A., Polani, D., and Nehaniv, C. (2008). Keep your options open: An information-based driving principle for sensorimotor systems. *PLOS ONE*, 3(12):1–14.

Kube, C. R. and Zhang, H. (1993). Collective robotics: From social insects to robots. *Adaptive Behavior*, 2(2):189–218.

Lizier, J. (2013). *The Local Information Dynamics of Distributed Computation in Complex Systems*.

Lungarella, M., Pegors, T., Bulwinkle, D., and Sporns, O. (2005). Methods for quantifying the informational structure of sensory and motor data. *Neuroinformatics*, 3:243–62.

Lungarella, M. and Sporns, O. (2006). Mapping information flow in sensorimotor networks. *PLOS Computational Biology*, 2(10):1–12.

Martius, G., Der, R., and Ay, N. (2013). Information driven self-organization of complex robotic behaviors. *PLOS ONE*, 8(5):1–14.

Morlino, G., Trianni, V., and Tuci, E. (2010). Collective perception in a swarm of autonomous robots. *In: Proceedings of the International Joint Conference on Computational Intelligence*, 1:51–59.

Nouyan, S., Groß, R., Bonani, M., Mondada, F., and Dorigo, M. (2009). Teamwork in self-organized robot colonies. *IEEE Transactions on Evolutionary Computation*, 13:695–711.

Pfeifer, R. and Scheier, C. (2001). *Understanding Intelligence*. Cambridge, MA: The MIT Press.

Polani, D. (2003). Measuring self-organization via observers. In *Advances in Artificial Life*, pages 667–675. Springer Berlin Heidelberg.

Polani, D. (2008). Foundations and formalizations of self-organization. In Prokopenko, M., editor, *Advances in Applied Self-organizing Systems*, pages 19–37. Springer London.

Polani, D., Prokopenko, M., and Yaeger, L. S. (2013). Information and self-organization of behavior. *Advances in Complex Systems*, 16(2&3):1–12.

Prokopenko, M. (2014). *Guided self-organization: Inception*. New York: Springer.

Prokopenko, M., Boschetti, F., and Ryan, A. (2006). An information-theoretic primer on complexity, self-organisation and emergence. *Complexity*, 15.

Prokopenko, M. and Gershenson, C. (2014). Entropy methods in guided self-organisation. *Entropy*, 16(10):5232–5241.

Roli, A., Ligot, A., and Birattari, M. (2018). Complexity measures in automatic design of robot swarms: An exploratory study. *In Artificial Life and Evolutionary Computation (WIVACE) 2017, volume 830 of CCIS, eds M. Pelillo, I. Poli, A. Roli, R. Serra, D. Slanzi, and M. Villani (Cham: Springer),*, pages 243–256.

Roli, A., Ligot, A., and Birattari, M. (2019). Complexity measures: Open questions and novel opportunities in the automatic design and analysis of robot swarms. *Frontiers in Robotics and AI*, 6.

Salge, C., Glackin, C., Ristić-Durrant, D., Greaves, M., and Polani, D. (2014). Information-theoretic measures as a generic approach to human-robot interaction: Application in corbys project. In *Proceedings of the 2014 ACM/IEEE International Conference on Human-Robot Interaction*, page 282–283.

Salge, C. and Polani, D. (2011). Local information maximisation creates emergent flocking behaviour. *ECAL*, pages 688–696.

Schreiber, T. (2000). Measuring information transfer. *Phys. Rev. Lett.*, (85):461–464.

Smithers, T. (1995). On quantitative performance measures of robot behaviour. *Robotics and Autonomous Systems*, 15(1):107–133.

Sperati, V., Trianni, V., and Nolfi, S. (2014). Mutual information as a task-independent utility function for evolutionary robotics. *Guided Self-Organization: Inception*, pages 389–414.

Strobel, V., Ferrer, E., and Dorigo, M. (2018). Managing byzantine robots via blockchain technology in a swarm robotics collective decision making scenario. *In: Proceedings of the 17th International Conference on Autonomous Agents and Multi-Agent Systems*, page 541–549.

Tarapore, D., Lungarella, M., and Gómez, G. (2006). Quantifying patterns of agent–environment interaction. *Robotics and Autonomous Systems*, 54(2):150–158.

Trendafilov, D. and Murray-Smith, R. (2013). Information-theoretic characterization of uncertainty in manual control. In *Proc. IEEE SMC*.

Valentini, G. (2017). Achieving consensus in robot swarms: Design and analysis of strategies for the best-of-n problem. *SCI*, 706.

Valentini, G., Brambilla, D., Hamann, H., and Dorigo, M. (2016). Collective perception of environmental features in a robot swarm. In *Swarm Intelligence*, pages 65–76.

Werfel, J., Petersen, K., and Nagpal, R. (2014). Designing collective behavior in a termite-inspired robot construction team. *Science*, 343(6172):754–758.