

Perception-Action Loops of Multiple Agents: Informational Aspects and the Impact of Coordination

Philippe Capdepuy Daniel Polani

Chrystopher L. Nehaniv

Adaptive Systems Research Group

University of Hertfordshire

{p.capdepuy, d.polani, c.l.nehaniv}@herts.ac.uk

Received: date / Accepted: date

Abstract

Embodied agents can be conceived as entities perceiving and acting upon an external environment. Probabilistic models of this perception-action loop have paved the way to the investigation of information-theoretic aspects of embodied cognition. This formalism allows (i) to identify information flows and their limits under various scenarios and constraints, and (ii) to use informational quantities in order to induce the self-organization of the agent's behaviour without any externally specified drives. This paper extends the perception-action loop formalism to multiple agents. The multiple-access channel model is presented and used to identify the relationships between informational quantities of two agents interacting in the same environment. The central question investigated in this paper is the impact of *coordination*. Information-theoretic limits on what can be achieved with and without coordination are identified. For this purpose, different abstract channels are studied, along with a concrete example of agents interacting in space. It is shown that, under some conditions, self-organizing systems based on information-theoretic quantities have a tendency to spontaneously generate coordinated behaviour. Moreover, in the perspective of engineering such systems to achieve specific tasks, these information-theoretic limits put constraints on the amount of coordination that is required to perform the task, and consequently on the mechanisms that underly self-organization in the system.

1 Introduction

In the recent years there has been an increasing interest into information-theoretic quantities in the perception-action loop of embodied agents. Not only do they provide a unified and 'coordinate-free' framework for identifying limits on control [17, 18] and for studying how information flows [3] are structured under various scenarios [8, 10], but they also seem to be useful for guiding the self-organization of behaviour [19, 2, 12], the emergence of internal representations [9, 5] and interactive learning [14].

One such promising information quantity is *empowerment* [8]. It is defined as the maximum amount of information that an agent is able to send through the environment and retrieve through its sensors. Empowerment maximization for a single-agent leads to several interesting properties. It allows to evolve sensors and actuators, extract internal representations from sensorimotor data, or discover and perform 'natural' homeostatic behaviour.

However, apart from previous work on the ability of empowerment to generate complex spatial organizations [6], the interpretation and use of similar information-theoretic quantities in the case of multi-agent systems has been

mostly left aside. This work builds upon recent advances in information theory in order to provide a better understanding of the perception-action loop of multiple agents sharing the same environment and of agents having multiple actuators. Using a minimal model of such situations, the impact of coordination on information quantities is studied. Limits are identified and we discuss their consequences on information-driven self-organizing systems, on the tasks they can achieve, and on the corresponding constraints this imposes on the mechanisms underlying their self-organization.

This paper is structured as follows. The next section introduces the perception-action loop formalism that allows to express the information-theoretic quantities of interest. The cases of two agents interacting in the same environment and one agent using two distinct actuators are presented and stripped down to a minimal model that captures their interactions.

The resulting model, identified as the *multiple-access channel* (MAC) [7], is presented in section 3 along with the information-theoretic quantities of interest; the impact of coordination between agents on these quantities is theoretically studied. A subset of the space of possible MACs is numerically investigated in section 4, indicating that some channels have the peculiar property that coordination can increase information rates by an amount larger than invested in the coordination.

In section 5, we propose to classify MACs into different groups, depending on the impact that coordination has on their accessible information rates. Examples of MACs belonging to each class are presented. These classes are used in section 6 in order to distinguish between different situations that two agents may encounter while interacting in a very simple grid-world.

Section 7 concludes the paper by discussing the impact of these results on agents whose self-organization is guided by information-theoretic quantities, and mainly how empowerment-driven self-organization can result in agents that coordinate their actions.

2 Perception-Action Loops with Multiple Actuators

The framework of Causal Bayesian Graphs (CBG) is used in order to model the perception-action loop of embodied agents. The main idea is to model sensors (S), actuators (A), and the environment (R) through random variables¹. The interplay between these variables over time is described by conditional probability distributions which are represented as causal links of the graph².

A typical perception-action loop of a single agent is represented on Fig. 1.

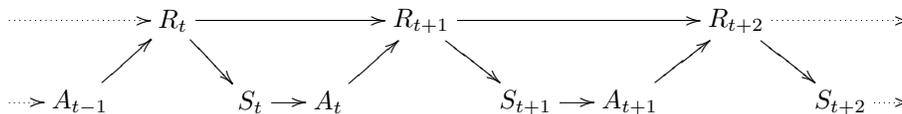


Figure 1: Representation of a typical perception-action loop as a causal Bayesian network unrolled in time with closed-loop usage. R stands for the state of the environment, S is the sensor of the agent and A its actuator. Other models are possible, for instance open-loop control can be represented by removing the arrows from S to A . The agent can also be equipped with a memory which would then be represented as an extra variable on the graph.

When two agents A and B are interacting in the same environment, their perception-action loops get intertwined. The resulting CBG is represented on Fig. 2 (left). Similarly, one can consider a single agent that has two distinct actuators (A and B), the corresponding CBG is depicted on Fig. 2 (right).

¹Other variables such as the agent's internal memory can also be modeled.

²Actually, the graph defines constraints on the dependencies between random variables, and therefore constraints on the possible conditional probability distributions.



Figure 2: **Left:** typical perception-action loop of two agents sharing the same environment. R : state of the environment, S, A : sensor and actuator of agent A, respectively T, B for agent B. **Right:** perception-action loop for a single agent using two distinct actuators A and B .

The two situations described above have something in common, which is that the information flowing from the actuators A_t and B_t (whether these belong to one or two agents) ‘collides’ in the next environment state R_{t+1} . Put another way, the environment becomes a ‘choke point’ where all information flows coming from actuators have to go through before flowing into the next sensors.

When one considers information-theoretic driving principles, the amount of information that can potentially be injected in the environment is a crucial aspect. For instance, when using empowerment [8], one is interested into the amount of information that can flow from the actuators to the next sensor states. Because these flows have to go through the environment before reaching the sensors, the amount of information that can be injected in the environment is an upper limit on how much can be injected in the sensors (via the data processing inequality). Similarly, if one considers the principle of predictive information [2], it is the amount of information ‘flowing’ from past to future sensors that is of interest, therefore the amount of information flowing into the environment is again a limiting factor³.

Now that this choke point has been identified, we are going to make a few simplifying assumptions in order to study the information-theoretic quantities of interest:

- Because the sensors are only a degenerate copy of the state of the environment (through the sensory apparatus of the agents), the amount of information flowing in the environment is an upper limit on the amount of information flowing in subsequent sensors. Therefore, this paper will only study this quantity, and the agents are assumed to perceive the full state of the environment ($S = T = R$). In order to simplify the graphs and notation, S and T are dropped and only R will be used.
- Only one step of the perception-action loop will be studied, allowing us to drop the time indices (in particular, we will talk about actions A and B instead of A_t and B_t). Introducing more steps requires a heavier framework in which feedback has to be integrated, this is possible using the concept of directed information [11] and the feedback capacity [15], however this is out of the scope of this paper. Also it will be shown that considering a single step already brings in a lot of richness in the informational dynamics.
- The variables used for choosing the action (i.e. previous sensorimotor experience) will be aggregated in a single variable C (which stands for *common*) that encompasses all the information that is common to both agents. Note that when we will talk about jointly correlated actions A and B by two different agents, any such correlation is assumed to be mediated by the common causal predecessor (see also footnote 6) which we consider to be comprised by C . In particular, we will consider the support of C (i.e. the set of outcomes of C which have nonvanishing probability) to be suitably comprehensive as to accommodate all necessary correlations in the agents’ action variables A and B . This can be always guaranteed by constructing a suitable random variable C defined over the product space of the variables A and B .

³It has to be noted that, unlike empowerment, predictive information is not defined as an information flow but as standard mutual information. Therefore two contributions are at work: the information actually flowing from past sensors to future sensors through the route $S_t \rightarrow A_t \rightarrow R_{t+1} \rightarrow S_{t+1}$, and the mutual information that may result from common information in the state of the environment, i.e. $R_t \rightarrow S_t$ and $R_t \rightarrow R_{t+1} \rightarrow S_{t+1}$ (these two contributions may overlap). Nevertheless, similarly to empowerment, both contribution are impacted by the environment acting as a ‘choke point’ for multiple actuators.

- Furthermore, we also make the simplifying assumption that the environment is in a specific fixed state before the agents behave, i.e. the next state of the environment only depends on the actions of the agents. Put another way, C does not (need to) contain any information about the current state of the environment because there is only one starting state in which there is no variability. All scenarios we are considering in the following make this particular assumption.

This assumption simplifies the discussions in order to focus on uncovering the impact of coordination on information-theoretic quantities, rather than combining it with the issue of identifying hidden world information influencing the channel. In particular, we thereby avoid tracking the amount of information that each agent would have to know and use about the current state of the environment and how it changes the quantities of interest. Otherwise, even in the single agent case, this would already add some intricacies. We leave them out of this paper in order to focus solely on the coordination aspects.

With these assumptions, we converge on a minimal model of the perception-action loop with multiple actuators as depicted on Fig. 3.

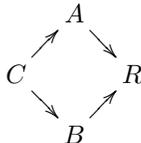


Figure 3: Minimal model of a perception-action loop with multiple actuators (belonging to a single agent or two different agents). Only one step is considered, all past information has been aggregated in a single variable C , agents are assumed to have perfect sensors (therefore R is used as the sensor), and the environment is assumed to have no memory.

This minimal core model forms the conceptual basis of all the following discussions. Important aspects of this simplification should be discussed: In the paper, we emphasize repeatedly the importance of embodiment. We note that, even in the minimal configuration of Fig. 3, with the following state R not depending on the previous state or directly on C , but only on the actions A and B , the aspect of embodiment is retained. Namely, in this simple model, embodiment is already represented in the way how actions A and B influence the world state R . We do not require a direct feedback of the environment on itself as a core property of embodiment. The spatial example from Sec. 6 will — among others — illustrate this principle.

3 Multiple-Access Channels and Coordination

If one removes the common information C , the minimal model presented on Fig. 3 is actually equivalent to the *multiple-access channel* (MAC) studied in the information theory literature (see [7] for an introduction). A two-users MAC is defined by a conditional probability distribution $p(r|a, b)$, and it is generally assumed that the inputs of each user are independent, therefore the space of input distributions is the set of distributions $p(a, b) = p(a)p(b)$. Let us denote this set as \mathcal{Q}_i , and the set of all joint distributions $p(a, b)$ as \mathcal{Q} .

Similarly to the single-user channel, the MAC has a capacity which is referred to as the *total capacity*. It is defined as:

$$Cap_{total} = \max_{\mathcal{Q}_i} I(A, B; R). \quad (1)$$

Algorithms for computing the total capacity and the corresponding capacity-achieving distributions of arbitrary discrete MACs has only been recently introduced by [13].

Another quantity of interest is the *joint capacity* of the MAC, i.e. the capacity when the inputs are not independent. This is equivalent to considering the graph $A, B \rightarrow R$, and therefore it is expressed as:

$$Cap_{joint} = \max_{\mathcal{Q}} I(A, B; R). \quad (2)$$

This quantity can be computed for an arbitrary discrete MAC using the standard Blahut-Arimoto algorithm [4, 1].

Because $\mathcal{Q}_i \in \mathcal{Q}$, it follows directly that

$$Cap_{total} \leq Cap_{joint}. \quad (3)$$

Therefore, for any channel where this inequality is strict, it means that coordinating the inputs can increase the overall transmission rate.

3.1 Common Information

Let us now reintroduce the common variable C as depicted on Fig. 3. Because of the structure of the graph, correlation of the inputs is only due to the common information. Put another way, we have $I(A; B|C) = 0$.

One may now ask whether coordination through a common source variable can increase the capacity of the MAC when the common information is discarded, i.e. we are interested in the quantity $I(A, B; R|C)$.

It turns out that this is not the case. Indeed, let us define \mathcal{P} as the set of all joint distributions $p(a, b, c)$ that can be written as $p(a|c)p(b|c)p(c)$ (according to the CBG of Fig. 3). The corresponding capacity can be written as:

$$\begin{aligned} Cap_{common} &= \max_{\mathcal{P}} I(A, B; R|C) \\ &= \max_{p(c)} \sum_c p(c) \max_{p(a|c), p(b|c)} I(A, B; R|c). \end{aligned}$$

For a given c , in the case where the environment has no memory, the inner maximization is equivalent to

$$\max_{p(a), p(b)} I(A, B; R)$$

which is equivalent to Cap_{total} . Therefore, because the result of the inner maximization does not depend on c , we have:

$$Cap_{common} = Cap_{total} \quad (4)$$

This means that coordination through a common information source does not increase the capacity of the channel when the common information is discarded. Put another way, agents that coordinate thanks to a common information source cannot send more ‘original’ information (in the sense that it is not already contained in C) than they would without any coordination.

This result seems to contradict the intuition that coordination should be beneficial. However, as we will show in the next subsection, there is still a benefit in coordination, only it does not appear when considering this capacity.

3.2 Capacity Regions

Besides the overall transmission rates presented in the previous section, it is also of interest to know the individual contributions of the different users of the MAC. These can be expressed using the *capacity region* (see Fig. 4). The capacity region is defined [7] as the set of accessible rate pairs (R_A, R_B) that satisfy:

$$\begin{aligned} R_A &\leq I(A; R|B), \\ R_B &\leq I(B; R|A), \\ R_A + R_B &\leq I(A, B; R). \end{aligned}$$

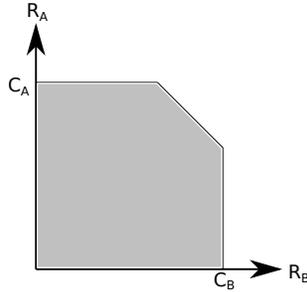


Figure 4: Example of a capacity region in a two-users MAC. C_A and C_B are the maximum capacities of each user. Grey area is the feasible region.

The capacity region shows that even though the overall transmission rate has a unique upper limit, the rates can be dispatched in various ways between the different users of the MAC. Generally, only one agent can send at the maximum rate, restricting the rates accessible to the other agent (this is not true for some channels as will be described in Sec. 5).

One has to note that the rates are absolute, in the sense that the other agent's action is known (and conditioned upon). We will refer to this region as the absolute region \mathcal{R}_{abs} .

However, in the situation of two embodied agents, agent A cannot know the action of agent B occurring at the same timestep, unless the latter contains information coming from an accessible common source. Therefore, if we consider that the agents cannot access each other's actions but have knowledge of a common variable C , we can define a new capacity region \mathcal{R}_{com} as the set of rate pairs (R_A, R_B) which satisfy the following inequalities:

$$\begin{aligned} R_A &\leq I(A; R|C), \\ R_B &\leq I(B; R|C), \\ R_A + R_B &\leq I(A, B; R|C). \end{aligned}$$

One has to note that, for both regions, the correlation between actions (if there is any) is properly 'discarded'. Indeed, for the absolute region, conditioning on the other agent's action removes any redundancy they may have. In the case of the common region, correlations between actions have to come from the common source C (as depicted on the CBG of Fig.3). Therefore conditioning on this common source removes any correlation between A and B (i.e $I(A; B|C) = 0$).

3.3 Relations Between Capacity Regions and Coordination

It was shown before that coordinating the inputs can increase the overall information rate in the channel $I(A, B; R)$, but not when the common information is discarded $I(A, B; R|C)$. It follows directly from this that the absolute capacity region \mathcal{R}_{abs} can be extended by coordinating actions. However, it is not clear yet whether this could also apply to the common capacity region \mathcal{R}_{com} .

Let us first define some sets of distributions. In the previous section we have introduced the set \mathcal{P} of joint distributions $p(a, b, c)$ that factorize as $p(c)p(a|c)p(b|c)$. For our study, it is useful to consider subsets of \mathcal{P} (see Fig. 5). Specifically we define:

- the independent set \mathcal{P}_i in which the coordination $I(A; B)$ is zero.

- the ‘pure’ set \mathcal{P}_p in which the information captured by A or B from C is also captured by B or A respectively. Because of the structure of the CBG on Fig. 3, this can also be expressed as $I(A; B) = I(A; C) = I(B; C)$. It is pure in the sense that no actuator contains part of the common information that the other actuator does not also contain.
- the independent pure set $\mathcal{P}_{i,p}$ in which no common information is captured by any actuator. This can be expressed as $I(A; B) = I(A; C) = I(B; C) = 0$.

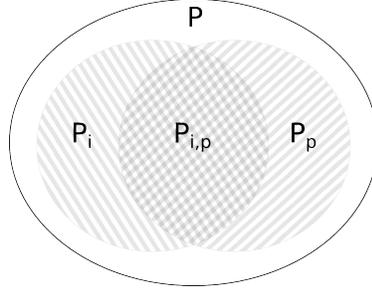


Figure 5: Relationship between the different sets of distributions $p(a, b, c) = p(a|c)p(b|c)p(c)$.

Let us now consider the mapping π that maps an element of \mathcal{P} , i.e. a joint distribution $p(a, b, c)$ that factorizes as $p(c)p(a|c)p(b|c)$ into an element of \mathcal{Q} , i.e. a joint distribution $p(a, b)$, using

$$p(a, b) = \sum_c p(a, b, c).$$

For non-correlated input distributions, we have the following equivalences:

$$\pi(\mathcal{P}_{i,p}) = \pi(\mathcal{P}_i) = \mathcal{Q}_i. \quad (5)$$

In the case of correlated inputs, we have the following relationships:

$$\mathcal{Q}_i \subseteq \pi(\mathcal{P}_p) = \pi(\mathcal{P}) \subseteq \mathcal{Q} \quad (6)$$

with equality for the left-hand part if $|C| = 1$.

The baseline, i.e. the standard MAC without common source and therefore independent inputs, is the capacity region for the set $\mathcal{P}_{i,p}$. In the case of the absolute capacity region \mathcal{R}_{abs} , we have the following results:

- Because \mathcal{P} and \mathcal{P}_p map to \mathcal{Q} (or at least to a set larger than \mathcal{Q}_i when $|C| > 1$), it follows that coordination alone (without any agent having access to extra information about the other agent’s actions) can increase the overall rate $I(A, B; R)$, and therefore extend the capacity region.
- Because \mathcal{P}_i maps to \mathcal{Q}_i , it follows that having extra knowledge of the other agent’s action but not coordinating with it does not extend the capacity region (which is obvious because we are already conditioning on the other agent’s action).

Now, if we consider the common capacity region \mathcal{R}_{com} , we can draw the following conclusions:

- Because \mathcal{P} and \mathcal{P}_i provide the agents with extra information about each other’s actions, the individual rates can be increased, but as shown before, the overall rate cannot. Nevertheless this extends the capacity region.

- On the other hand, \mathcal{P}_p does not provide them with extra information, therefore individual rates cannot be increased, so the region cannot be extended.

More generally what we can say is that coordinating actions can extend both the overall absolute rate $I(A, B; R)$ and the absolute capacity region \mathcal{R}_{abs} , however it does not impact the overall common rate $I(A, B; R|C)$ nor the corresponding region. On the other hand, having knowledge about the other agent's actions thanks to the common information allows to extend the common capacity region \mathcal{R}_{com} , making some new rate pairs accessible, but it does not allow to increase the overall rate $I(A, B; R|C)$ and does not change the absolute region.

4 Numerical Results

4.1 Limits on Capacity Gain

We have shown in the previous sections that the absolute capacity of the MAC $I(A, B; R)$ can be increased through coordination. One may ask how much can this increase be? In order to provide an answer to this question, we have conducted numerical experiments.

The methodology is the following. The space of regular (2,2;4) MACs (2 users, $|\mathcal{A}| = 2$, $|\mathcal{B}| = 2$, $|\mathcal{R}| = 4$) is discretized and sampled. The discretization is done by considering only the channels for which the conditional probability distribution can be expressed as $p(r|a, b) = \frac{n}{N}$ where $\{n \in \mathbb{N}; 0 \leq n \leq N\}$ and N is fixed⁴. For each MAC, three values are computed:

- The total capacity Cap_{total} using the algorithm from [13].
- The joint capacity Cap_{joint} using the standard Blahut-Arimoto algorithm.
- The amount of coordination $I^*(A; B)$ required to achieve the joint capacity.

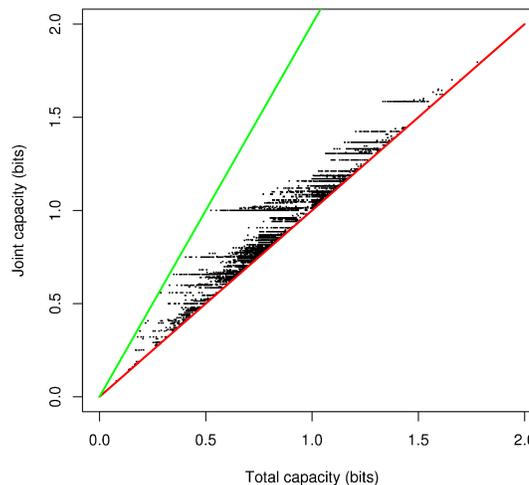


Figure 6: Numerical exploration of joint and total capacities in the space of (2,2;4)-MACs. The red and green lines respectively have equations $y = x$ and $y = 2x$. Discretization parameter for the sampling process is $N = 4$.

Results of this experiment are shown on Fig. 6. It seems from these results that coordination can at most double the capacity of the MAC (other searches with different alphabet sizes have been performed without finding any

⁴For $|\mathcal{R}| = 4$ and $N = 4$ we have 35 different distributions for a given a, b , and because $|\mathcal{A}||\mathcal{B}| = 4$, there are $35^4 = 1500625$ different distributions.

channel violating this rule).

A related conjecture is presented in [16] which shows that the total capacity of any MAC with white Gaussian noise can be at most doubled when the channel is used with feedback (even for more than two users). This might appear unrelated to the problem of coordination, but it actually is. Indeed, feedback in the case discussed in [16] allows the users to have a common source of information which they can use to correlate their messages. The author also shows that for a general MAC with k users the total capacity with feedback is bounded from above by kC where C is the total capacity without feedback. The paper concludes by conjecturing that the capacity doubling upper limit might hold for the general MAC with k users.

4.2 Coordination and Capacity Gain

We know that coordination can improve the capacity of the MAC. However, one may ask what is the relation between the amount of coordination and the capacity gain? It would seem intuitive at first sight to expect that the capacity gain $\Delta Cap = Cap_{joint} - Cap_{total}$ is bounded from above by the amount of coordination $I^*(A; B)$ used for the joint capacity achieving distribution. This is indeed the case for many of the channels that were sampled as can be seen on Fig. 7. However, there are a few of them for which this is not true.

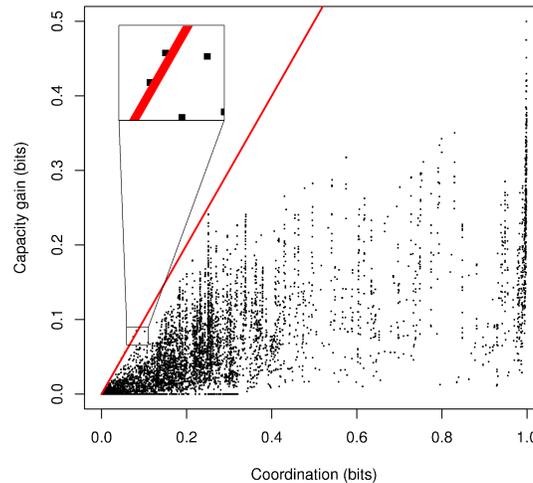


Figure 7: Numerical exploration of the space of (2,2;4)-MACs. Capacity gain ΔC versus coordination $I^*(A; B)$. The red line has equation $y = x$. Discretization parameter is $N = 4$. There are some points just above the $y = x$ line, for example at coordinates $(0.082, 0.085)$ as can be seen in the zoom (see text for details).

One such channel is the deterministic binary erasure channel. This channel and its capacity achieving distributions are described by the following table ($p^*(a, b)$ and $p_i^*(a, b)$ stand respectively for the joint and total capacity-achieving distributions):

$\mathbf{p}(\mathbf{r} \mathbf{a}, \mathbf{b})$	r	0	1	2	$\mathbf{p}^*(\mathbf{a}, \mathbf{b})$	$\mathbf{p}_i^*(\mathbf{a}, \mathbf{b})$
a	b					
0	0	1	0	0	$\frac{1}{3}$	$\frac{1}{4}$
0	1	0	1	0	$\frac{1}{6}$	$\frac{1}{4}$
1	0	0	1	0	$\frac{1}{6}$	$\frac{1}{4}$
1	1	0	0	1	$\frac{1}{3}$	$\frac{1}{4}$

The total capacity of this channel is $C(\mathcal{P}_{i,p}) = 1.5$ and its joint capacity is $C(\mathcal{P}) = \log_2(3) \approx 1.585$ bits. The joint capacity achieving distribution has coordination $I^*(A; B) \approx 0.082$. Therefore the capacity gain $\Delta C \approx 0.085$ is higher than the coordination. Obviously, this contradicts our initial intuition that the amount of coordination is a limiting factor to the capacity gain (even though it seems true in most other cases). It is not clear at this stage why this is the case. A more thorough analysis of the quantities involved and their subcomponents would be needed to understand this phenomenon. One could conjecture that the shared information is somehow used ‘twice’: the amount of coordination is added once as information transmitted through the channel, and a second time to improve the rate of transmission of ‘original’ information.

5 Classes of MACs

Understanding which MACs will provide a benefit from coordination is a key aspect for analyzing complex situations. In this section we identify different classes of MACs and provide examples. These classes are used in the next section to analyze a grid-world scenario with two agents. More generally their general structure can be related to more complex channels to give an intuition of how these would behave with respect to coordination.

5.1 Non-Interfering MACs

This first class of channels is the simplest one. It can be characterized by the fact that each source of the MAC sends information into independent degrees of freedom at the output. It is quite obvious that in such channels the amount of information independently sent by one user is unaffected by that of the other user. Put differently: any rate is reachable by any user independently of the other users. A consequence of this is that introducing correlations between sources can only decrease the total amount of information sent, for the simple reason that correlations reduce the degrees of freedom available to the users, and therefore reduce the rate at which they can send information through the channel.

A non-interfering MAC has the following properties:

- the joint and total capacities are the same.
- the absolute and common capacity regions are not extended by introducing coordination or knowledge of the other agent’s actions.

The most prototypical channel of this class is a composition of two binary symmetric channels. It is referred to as the binary symmetric (2,2;4)-MAC and described by the following conditional probability distribution:

$\mathbf{p}(\mathbf{r} \mathbf{a}, \mathbf{b})$	r	00	01	10	11
a	b				
0	0	1	0	0	0
0	1	0	1	0	0
1	0	0	0	1	0
1	1	0	0	0	1

5.2 Weakly-Interfering MACs

This class of channel is characterized by the following properties:

- the joint and total capacities are the same,

- the common capacity region \mathcal{R}_{com} is extended by introducing coordination and knowledge of agents' actions.

A good example of such a channel is the XOR (2,2;2)-MAC. Basically the output is computed as a XOR of the two inputs. It has the following conditional probability distribution:

$\mathbf{p}(\mathbf{r} \mathbf{a}, \mathbf{b})$	r		
	0	1	
a	b		
0	0	1	0
0	1	0	1
1	0	0	1
1	1	1	0

The common capacity region with and without coordination is represented on Fig 8.

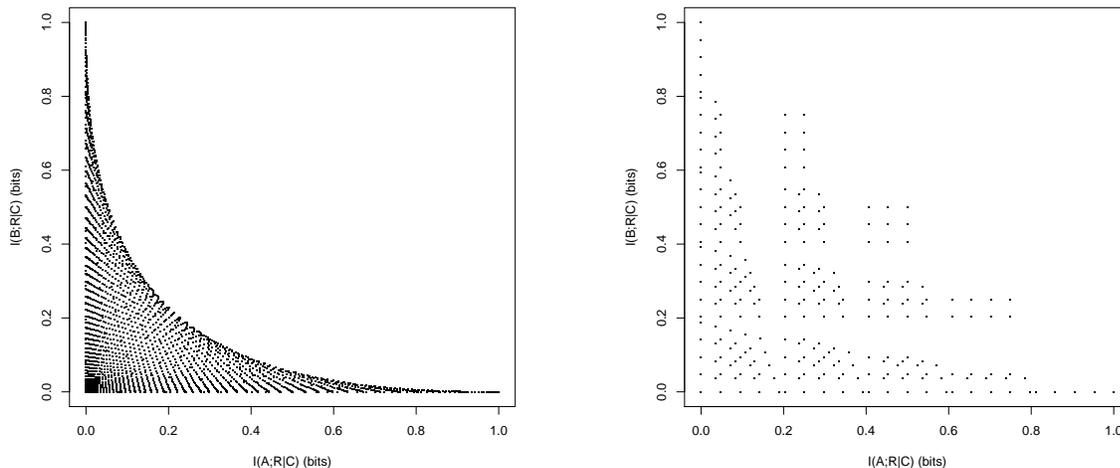


Figure 8: Numerical exploration of the space of input distributions conditioned on a common information source for the XOR MAC. Left: common capacity region for policies independent of C . Right: common capacity region for policies that depend on C .

5.3 Strongly-Interfering MACs

This class of channel is characterized by the following properties:

- the joint and total capacities are different: $C(\mathcal{P}_{i,p}) \leq C(\mathcal{P})$, and
- the common region for the independent pure set is a subset of the common region for the full set: $\mathcal{R}_{com}(\mathcal{P}_{i,p}) \subset \mathcal{R}_{com}(\mathcal{P})$.

An example of a strongly-interfering channel is the binary totalistic⁵ channel. Basically if both inputs are the same then the same output symbol is emitted, otherwise the output is a uniform distribution. The totalistic (2,2;2)-MAC is described by the following conditional probability distribution:

⁵We refer to it as a totalistic channel because the sum of both inputs is the relevant parameter, if the sum is equal to 0, then a 0 is transmitted, if the sum is 2 then a 1 is transmitted, if the sum is 1 then the output is uniformly distributed.

$\mathbf{p}(\mathbf{r} \mathbf{a}, \mathbf{b})$	r	0	1
a	b		
0	0	1	0
0	1	$\frac{1}{2}$	$\frac{1}{2}$
1	0	$\frac{1}{2}$	$\frac{1}{2}$
1	1	0	1

The capacity region with and without coordination is represented on Fig 9. Coordination in strongly-interfering channels makes some new rate pairs accessible between the extrema of the non-coordinated capacity region. This can be understood as a time-sharing mechanism mediated by the common information that allows agents to ‘take turns’ and transmit at their maximum rates without interfering with each other.

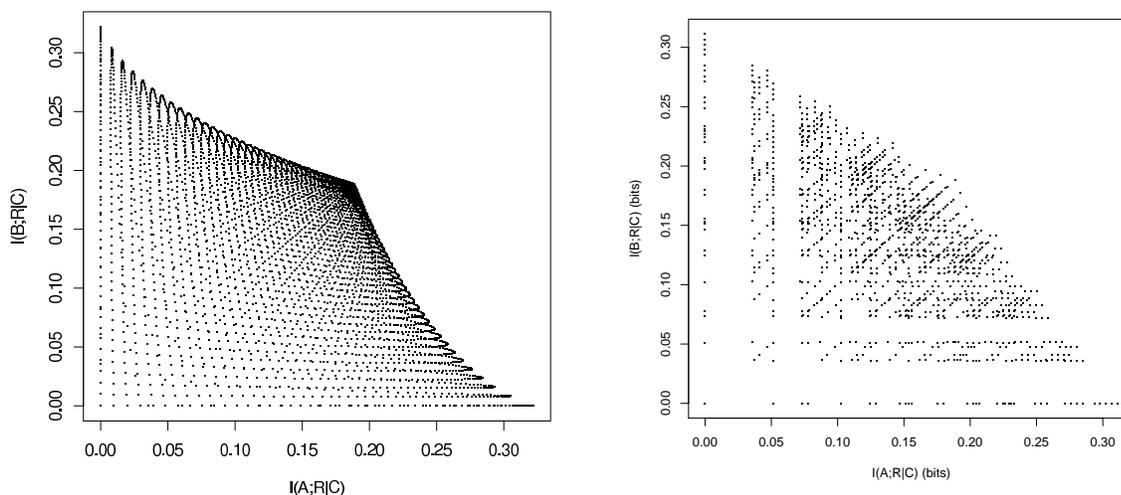


Figure 9: Numerical exploration of the space of input distributions conditioned on a common information source for the totalistic MAC. Left: capacity region for policies independent of C . Right: capacity region for policies that depend on C .

6 Spatial Example

We now turn to a more concrete example of embodied agents. Imagine two agents A and B (by which we at the same time denote their respective actions) in a one-dimensional grid-world space (see Fig. 10). They are allowed to move in this space but they cannot occupy the same tile, if they try they ‘collide’ with each other. The basic principles are the following:

- We consider a single time-step.
- Available actions are
 - *move one step to the left*
 - *move one step to the right*
 - *do not move* (in some of the experiments — see below).

- Both agents take their respective action (A vs. B) simultaneously, not one after the other.
- If one agent tries to move toward a tile that is occupied by the other agent, and the other agent is not moving or is also moving toward the first agent, then the move has no effect, and both agents remain at their original position.
- If both agents try to move toward the same empty tile, then each of them has 50% chances of moving onto that tile while the other stays where it is.
- In all other situations, the moves are carried out.

In our particular experimental setup, we adopt a number of additional assumptions: The agents always perceive the full state of the environment R after the action is performed. Also, to make the analysis of the various influences more transparent, we are going to place the agents in different inequivalent fixed starting positions and study the MAC capacities for each of these starting constellations. This last assumption requires a discussion in the context of Fig. 3. The choice of a fixed starting constellation corresponds to an open-loop control problem, but as this is a one-step scenario, it still allows one to achieve full capacity since the precise state of the MAC channel is known when the action is taken⁶.

The corresponding MACs are studied for one step. Three different starting positions are considered (see Fig. 10): A and B are touching each other, they are one or two tiles away from each other. We refer to these channels as Touch, 1-Tile, and 2-Tiles.

We also study two different embodiments: the simple embodiment where agents have only two actions (left and right), and the full embodiment in which agents can move left, right, or stay on their tile. Considering both the starting positions and the embodiment, and assuming that the agents perceive the full state (i.e. both absolute positions), we obtain six different MACs.

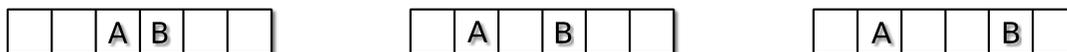


Figure 10: Two agents in a one dimensional space. Three starting positions: the agents are touching each other, or they are one or two tiles away from each other.

For each MAC, the joint and total capacities are computed, along with the capacity regions. This allows to identify the class each MAC belongs to. Results are presented in the following table:

⁶ Note that also in this scenario we assume the implicit existence of a common variable C as in Fig. 3. Its role is not to encode the initial world state which contains no variation, rather its purpose is solely to provide the causal source of joint correlation for the actions of the two agents that we formally demanded in Sec. 2. In a multi-step analysis, C would likely intertwine both the history of the run as well as the current positions of the agents to some extent to both provide a context state for the capacity computation, as in [5] and also a source for joint correlations of the individual agents' actions A and B . However, such an analysis would require a significantly more complex analysis beyond the scope of the present paper and we do not discuss it further, focusing on the one-step scenario only.

Embodiment Starting Position	Simple	Full
Touch	Non-Interfering $C(\mathcal{P}_{i,p}) = 2$ $C(\mathcal{P}) = 2$ \mathcal{R}_{com} unchanged	Strongly-Interfering $C(\mathcal{P}_{i,p}) = 2.5$ $C(\mathcal{P}) \approx 2.58$
1-Tile	Non-Interfering $C(\mathcal{P}_{i,p}) = 2$ $C(\mathcal{P}) = 2$ \mathcal{R}_{com} unchanged	Strongly-Interfering $C(\mathcal{P}_{i,p}) \approx 2.92$ $C(\mathcal{P}) = 3$
2-Tiles	Non-Interfering $C(\mathcal{P}_{i,p}) = 2$ $C(\mathcal{P}) = 2$ \mathcal{R}_{com} unchanged	Non-Interfering $C(\mathcal{P}_{i,p}) \approx 3.17$ $C(\mathcal{P}) \approx 3.17$ \mathcal{R}_{com} unchanged

As one can see in the table, an apparently minute change in the embodiment (adding a do-nothing action) has strong consequences on the nature of the corresponding MAC. When the agents are two or more tiles away from each other, their actions impact different degrees of freedom, i.e. each agent’s absolute location, whatever their embodiment. However, when the agents are close enough to each other to potentially ‘collide’ during the course of their action, coordination has an impact. It is interesting to note that the *Touch* and *1-Tile* starting positions with full embodiment have a structure partially similar to the binary erasure MAC mentioned earlier, with some extra degrees of freedom.⁷

7 Discussion

The results presented in this paper have to be put in the perspective of the simplifying assumptions used to get to the MAC model. It is not clear at this point if these results can be extrapolated to other situations. This prompts for further research with less simple assumptions, for example by considering that the agents do not sense the full state of the environment, that the environment has memory and that the coordination ‘signal’ has to pass through the environment, or that more than two agents have access to the channel.

Nevertheless, already with the minimal, trimmed-down models discussed here, the relationships between coordination and the information-theoretic quantities of embodied agents that have been presented in this paper have several consequences.

Let us consider the context of multiple agents that obey information-theoretic drives such as empowerment or predictive information maximization. In this case, the agents are expected to self-organize in a way such that their individual drives are maximized. However, because information rates or potential information rates (in the case of empowerment) are maximized when their actions are coordinated or at least when they share information about their future actions through the environment, it would be expected that they exhibit spontaneous coordination of their actions.

Another way to increase their information rates is to move to non-interfering situations if those are preferable, which is not necessarily the case or might not even be possible in some environments.

⁷One may be surprised that the simple embodiment *Touch* and *1-Tile* situations are classified as non-interfering. Indeed, because the agents can collide, it would seem natural to assume that they interfere. But this is only true with the common meaning of ‘interfering’, not for the information-theoretic quantities. For instance, in the *Touch* situation, denoted as $-AB-$, if agent B goes left, then, depending on agent A ’s action, two situations can be reached: $AB--$ and $-AB-$. If B goes right, then the possible end states are $A--B$ and $--AB$. Now, if we consider that B is acting non-deterministically, then A can reach two *groups* of end states: $A--B/AB--$ or $-AB-/--AB$, with in-group probabilities that depend only on B . In the end, the capacity for A is unchanged. A similar reasoning applies to the *1-Tile* situation.

Moreover, even in very simple concrete scenarios such as two agents colliding in a one-dimensional space, the corresponding MACs have interesting coordination properties (such as being strongly interfering). The classes of channels presented in this paper can be related to more complex situations than those given as examples, allowing us to quickly identify whether these situations will lead to coordinated behaviour or not. Another important result is that slight changes in the embodiment can have an important impact on the underlying MAC, potentially changing its class from non-interfering to strongly-interfering.

If instead of multiple independent agents we consider the situation of multiple agents acting as a group, or one agent having multiple actuators, the same principles should push the agents towards either separating the degrees of freedom in which they or their actuators operate, or generate a coordination mechanism so that the control abilities of the group or of the multi-actuated agent are maximized.

Moreover, the numerical experiments presented above have shown that in some cases the amount of capacity gained from using coordination can be greater than the actual amount of coordination. Other numerical results (not presented in this paper) have shown that this effect is even stronger when below-capacity information rates are considered.

In a context where we want to engineer a multi-agent or multi-actuator system towards achieving a specific task, the underlying mechanisms have to allow for the required amount of coordination to occur in the system. For instance, if the task is to control the output of the totalistic channel presented in Sec. 5.3 in order to obtain a specific non-uniform binary sequence at the output, it is necessary for the agents to coordinate their actions. Without coordination, the expected sequence cannot be produced with probability 1.

It would be interesting in the future to look at more complex situations, especially extended over time, in order to study how agents can setup an effective coordination mechanism, i.e. generate and use a coordination signal. Not only is this interesting in order to identify the kind of coordination signals that are constructed by such systems, but also, because this signal has to be conveyed through the environment, agents would have to trade-off the available bandwidth between the actual information transmitted and the coordination signal. Studying such situations could bring important insights into the limits on the amount of coordination that can be achieved in practice.

Acknowledgements

The authors gratefully acknowledge the comments of reviewers and editors which considerably improved the paper.

References

- [1] S. Arimoto. An algorithm for computing the capacity of arbitrary discrete memoryless channels. *IEEE Transactions on Information Theory*, 18(1):14–20, 1972.
- [2] N. Ay, N. Bertschinger, R. Der, F. Güttler, and E. Olbrich. Predictive information and explorative behavior of autonomous robots. *European Physical Journal B*, 63(3):329–339, June 2008.
- [3] N. Ay and D. Polani. Information flows in causal networks. *Advances in Complex Systems*, 11(1):17–41, 2008.
- [4] R. Blahut. Computation of channel capacity and rate distortion functions. *IEEE Transactions on Information Theory*, 18(4):460–473, Jul 1972.

- [5] P. Capdepuuy, D. Polani, and C. L. Nehaniv. Constructing the basic umwelt of artificial agents: An information-theoretic approach. In *Proceedings of the Ninth European Conference on Artificial Life*, volume 4648 of *LNCS/LNAI*, pages 375–383. Springer, 2007.
- [6] P. Capdepuuy, D. Polani, and C. L. Nehaniv. 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. IEEE, 2007.
- [7] T. M. Cover and J. A. Thomas. *Elements of Information Theory 2nd Edition*. Wiley Series in Telecommunications and Signal Processing. Wiley-Interscience, July 2006.
- [8] A. Klyubin. *Organization of Information Flow Through the Perception-Action Loop*. PhD thesis, School of Computer Science, University of Hertfordshire, UK, 2007.
- [9] A. S. Klyubin, D. Polani, and C. L. Nehaniv. Representations of space and time in the maximization of information flow in the perception-action loop. *Neural Computation*, 19(9):2387–2432, 2007.
- [10] M. Lungarella and O. Sporns. Mapping information flow in sensorimotor networks. *PLoS Computational Biology*, 2(10):e144+, October 2006.
- [11] J. L. Massey. Causality, feedback and directed information. In *Proceedings of the International Symposium on Information Theory and its Applications*, 1990.
- [12] M. Prokopenko, V. Gerasimov, and I. Tanev. Evolving spatiotemporal coordination in a modular robotic system. In *From Animals to Animats 9: 9th International Conference on the Simulation of Adaptive Behavior (SAB 2006)*, volume 4095 of *Lecture notes in computer science*, pages 558–569. Springer, 2006.
- [13] M. Rezaeian and A. Grant. Computation of total capacity for discrete memoryless multiple-access channels. *IEEE Transactions on Information Theory*, 50(11):2779–2784, 2004.
- [14] S. Still. Information-theoretic approach to interactive learning. *EPL (Europhysics Letters)*, 85(2):28005 (6pp), 2009.
- [15] S. Tatikonda and S. Mitter. The capacity of channels with feedback. *IEEE Transactions on Information Theory*, 55(1):323–349, 2009.
- [16] J. Thomas. Feedback can at most double gaussian multiple access channel capacity. *IEEE Transactions on Information Theory*, 33(5):711–716, 1987.
- [17] H. Touchette and S. Lloyd. Information-theoretic limits of control. *Physical Review Letters*, 84:1156, 2000.
- [18] H. Touchette and S. Lloyd. Information-theoretic approach to the study of control systems. *Physica A*, 331:140, 2004.
- [19] K. Zahedi, N. Ay, and R. Der. Higher coordination with less control - a result of information maximisation in the sensori-motor loop. *CoRR*, abs/0910.2039, 2009.