Cyberoos2002: Confronting the Entropy of Joint Beliefs

Mikhail Prokopenko, Peter Wang

Intelligent Interactive Technology
CSIRO Mathematical and Information Sciences
Locked Bag 17, North Ryde, NSW 1670, Australia
{mikhail.prokopenko, peter.wang}@csiro.au

Current approaches to activity coordination in multi-agent systems (teams) range from strictly top down (plan-based coordination) to purely emergent (reactive coordination), with many hybrid variants, each having its specific advantages and disadvantages. It appears to be extremely difficult to rigorously compare various hybrid approaches to multi-agent coordination (and communication), given the lack of a generic semantics or some guidelines.

We studied some intuitive inter-agent communication policies and characterised them in terms of generic information-theoretic properties. In particular, the relative entropy of joint beliefs was suggested as an indicator of teams coordination potential. Our novel behaviour-based agent architecture (based on the Deep Behaviour Projection framework) enabled consistent reasoning about belief change, including beliefs about other agents. This allowed us to examine some of the identified communication policies empirically. The obtained results confirmed that there are certain interesting invariants — in particular, a change in team coordination (and overall performance) was shown to be within the boundaries indicated by the relative information entropy [10].

1 On entropy and multi-agent agreements

We characterised certain classes of multi-agent agreements, making as few assumptions as possible about the choice of inter-agent communication variables and periods of team synchronisation (extensively analysed by Stone and Veloso [12]). In particular, we studied *selfish* agreements covering "selfish" agents that communicate data about themselves only, *transitively-selfish* agreements ensuring that each "cooperative" agent always communicates the data about some other agent, and *mixed* agreements, where a team composition parameter determines the precise split between selfish and cooperative agents.

In order to capture the agreements in a formal information-theoretic setting we analysed the joint "output" of inter-agent communication after each period of team synchronisation. Then we estimated the *relative entropy* as a precise measure of the amount of freedom of choice (the degree of randomness) [11] contained in the resultant joint beliefs. We demonstrated that the relative entropy of joint beliefs in multi-agent teams may serve as a generic indicator of the team coordination potential. Clearly, the team following an agreement with near-zero entropy (almost no "misunderstanding" in joint beliefs) has a higher coordination potential than the team adherent to an agreement with near-maximal entropy (joint beliefs are almost random). In particular, the following representation results were obtained [10] for a particular protocol[1].

Theorem 1. *Selfish agreements attain minimal entropy.*

Transitively-selfish agreements without the selfish agents attain maximal entropy asymptotically when the number of agents $n \to \infty$.

The trajectory of the relative entropy in multi-agent teams ($n > 2$) following mixed agreements does not have a fixed-point in terms of the team composition parameter.

This theorem basically states that whenever team agents agree to communicate the data about themselves only, they eventually leave nothing to choice. In other words, they always

[1] A simple protocol $\mathcal{P}_1$ allows an agent to communicate data about only one agent precisely. In other words, each agent is able to encode either the data about itself or about the other agent. Without loss of generality, we may assume that the protocol $\mathcal{P}_1$ has enough symbols to encode n distinguishable objects and a single-object capacity for each communication message.

maximise their joint beliefs upon successful synchronisations. The obvious drawback is that while using single channels this saturation of joint beliefs requires that every agent *takes turns* in communication according to some schedule, and hence, large teams may take a while to minimise the entropy. The clear benefit, on the other hand, is that this minimisation is shown to be theoretically possible.

On the other hand, the "organisation" or "order" brought about by the transitively-selfish agreements is not sufficient to combat the entropy. Intuitively, the pair-wise "ignorance" of agents grows faster than the transitively-selfish agreement can cope with. Clearly, with the number of agents approaching infinity (and the entropy reaching its maximum asymptotically) the time to synchronise the team becomes infinite as well.

Obviously, the entropy of joint beliefs in multi-agent systems following mixed agreements exhibits some properties of both selfish and transitively-selfish configurations. We might expect that the selfish agents will bring in some order (as the compensation for potentially redundant information about themselves), while the cooperative (transitively-selfish) agents will lead to a higher degree of randomness (providing sometimes potentially non-trivial information about other agents).

Formally, the relative entropy produced by mixed agreements asymptotically approaches 1 with growth in the number of agents. In other words, the selfish agents "loose" the battle for order (asymptotically) when the number of agents is infinitely large. Interestingly, however, the lower limit is not zero, meaning that absolute order is never achievable regardless of the team split or the number of agents. In fact, our results showed that the joint beliefs obtainable in multi-agent teams with mixed agreements exhibit information-theoretic *complexity* in terms of the team composition. It has been recently pointed out in the literature (eg., by Suzudo [13]) that the entropy trajectory is a useful descriptor for a variety of self-organised patterns: eg., non-complex cellular automata (CA) have a fixed-point entropy trajectory and converge quickly to either very low or very high values. It should be noted that Suzudo considered the entropy of CA associated with the temporal pattern, while our analysis is focused on entropy of joint beliefs associated with the team composition parameter.

2 Agents situated in time and relativity of behaviours

The strength of the presented analysis, we believe, is in its generic nature. The results lay down some general guidelines in terms of team composition and suggest definite boundaries on the team coordination potential.

In the remainder of this section we focus on the agents ability to dynamically change their beliefs under different scenaria. In general, the agent's capability to maintain dynamic beliefs is based on another very important cognitive skill — the ability to remove itself from the current context. This ability is sometimes informally referred to as "imagination" [4] or as "possession of a reality simulator" [5]. Imagining or running a reality simulator allows the agent to reflect on past behaviour and project the outcome of future behaviour. For example, Joordens [5] makes a conjecture that higher mental states emerge as a result of a reality simulator: "an animal with no reality simulator basically lives in the present tense, and sees the world through only its eyes, at all times", while "the possession of a reality simulator may also allow an organism to experience many of the high-level cognitive processes that we identify with being human". Moreover, there is a possibility that cooperative behaviour of animals is related to their ability to run a reality simulator. The argument convincingly put forward by Joordens is that an organism with a reality simulator is more likely to engage in cooperative behaviour because of its ability to conceptualise rewards to others, and long-term rewards to itself.

We maintain that "world model" should appear in the architecture incrementally. In our previous work [7–9] we described the Deep Behaviour Projection (DBP) hierarchical framework. The DBP framework formally represents *increasing* levels of agent reasoning abilities, where every new level can be projected onto a deeper (more basic) behaviour. Put simply, a DBP behaviour can be present in the architecture in two forms: implicit (emergent) and explicit (embedded).

It is interesting at this stage to compare such behaviour duplication in DBP with the distinction between *automatic processes* and *controlled processes* in cognitive psychology. It is well-known that certain processes become highly automatic through repetition and are unconsciously triggered in the presence of certain stimuli, while controlled processes are mostly goal-oriented rather than reactive. With time and/or practice newly learned behaviours often shift from being controlled to automatic. Over the last few years, it has become apparent, however, that unconscious triggering can lead to strategic behavior — see, for example, a review by Bargh and Chartrand [1].

What the DBP approach suggests in addition, is that *the reactive/cognitive distinction is always relative* in a hierarchical architecture. The behaviour produced by the level l_k may appear reactive with respect to the level l_{k+1} but, at the same time, may look deliberate with respect to the level l_{k-1}. Let us exemplify this with the following three levels of the DBP agents:

- tropistic behaviour: Sensors → Effectors
- hysteretic behaviour: Sensors & Memory → Effectors
- tactical behaviour: Sensors & Memory & Task → Effectors.

The hysteretic behaviour is definitely more reactive when compared with the tactical behaviour, because the latter uses the task states in choosing the effectors. However, contrasted with a very basic tropistic behaviour, the hysteresis provided by (internal) memory states ensures a *degree* of cognition. More precisely, the hysteretic behaviour addresses some lagging of an effect behind its cause, providing a (temporary) resistance to change that occurred previously. For instance, in order to intercept a fast moving ball the agent needs to observe the shift in the ball positions and estimate its velocity before activating the effectors. Thus, the hysteretic behaviour is slightly more deliberate than the tropistic one (exemplified by a simple chase after the ball) — it better situates the agent in time (not only in space) and allows it to better respond to changes. Continuing with the example we re-iterate that the hysteretic intercept is a behaviour embedded explicitly, while the tropistic intercept is only possible as an emergent result of the recurring chase.

Thus, a reality simulator appears *incrementally* — starting from a basic ability to detect a change (eg., in direction) and moving towards a more and more comprehensive incorporation of the temporal asymmetry or "time's arrow" (eg., from direction-sensitive cells to a measurement of a shift in observed positions, to the notion of velocity emerging after a series of measurements, etc.). This means that an emergence of essentially new behavioural patterns always indicates a need for new elements in the agent architecture. At some stage, increasing levels of reasoning about change require an ability to consistently maintain the agent's beliefs — expand, contract or revise them according to some rational principles, such as the principle of minimal change (information economy) [3].

In order to address this requirement, we explicitly introduced a domain model into the DBP architecture, resulting in the following hierarchy (a refinement of the architecture reported in [9]):

$$
\begin{array}{lll}
\langle S, E, & \mathit{tropistic\_behaviour} : S \to E, & \\
\quad I, & \mathit{hysteretic\_behaviour} : I \times S \to E, & \mathit{update} : I \times S \to I, \\
\quad T, & \mathit{tactical\_behaviour} : I \times S \times T \to E, & \\
 & \mathit{tactics} : I \times S \times T \to 2^T, & \mathit{decision} : I \times S \times T \to T, \\
\quad D, & \mathit{domain\_update} : I \times S \times D \to D, & \\
 & \mathit{domain\_revision} : I \times S \times D \to D, & \\
 & \mathit{domain\_projection} : I \times S \times D \to S \ \rangle &
\end{array}
$$

where S is a set of agent sensory states, E is a set of agent effectors, I is a set of internal agent states, T is a set of agent task states, and D is a set of domain model states.

The DBP agents extrapolate their domain model each simulation cycle with the *domain_update* function, and revise it with the *domain_revise* function whenever new information becomes available. The partition between update and revision corresponds to the well-known distinction between belief update and belief revision [6]. In particular, the belief *update* is appropriate when the world has changed and the agents need to accommodate this change into the previously correct beliefs. The belief *revision* should be used to incorporate new information about the same state of the world, in order to correct potential inconsistencies.

In the absence of new observations, the updated domain model $d^* = domain\_update(i, s, d)$ is the best approximation of the domain. In these cases, the domain model d^* is transformed by the *domain_projection* function into the agent's sensory state $s^* = domain\_projection(i, s, d^*)$. Very importantly, all the choices made by the agent based on s^* are not distinguishable from the choices it could have made if the same sensory state s^* was a result of the direct sensory input. Intuitively, the *domain_projection* function projects the results of the reality simulator and the agent *imagines* that these results have been observed directly. The projection function is needed only in the absence of new observations (more precisely, only at the stall cycles [2]), and should not be invoked at other times — the imaginative side of the agent is not needed when "live" information is available anyway.

In order to support our analysis of boundaries on the team coordination potential, we varied communication policies while leaving all other factors (agents skills and tactics) unchanged. The factors beyond our control (eg., a possible change in the opponent strategy) were minimised by repeated runs. The obtained empirical results (the experiment is ongoing) illustrated the dependency between communication policies, the information entropy of joint beliefs and the team coordination potential. Identification of this relation is a main contribution of the presented analysis, opening a new general perspective on reasoning about belief dynamics in multi-agent scenaria.

References

1. John A. Bargh and Tanya L. Chartrand. The unbearable automaticity of being. American Psychologist, 54(7), 1999.
2. Marc Butler, Mikhail Prokopenko and Thomas Howard. Flexible Synchronisation within RoboCup Environment: a Comparative Analysis. In RoboCup-2000: Robot Soccer World Cup IV, Springer, 2000.
3. Peter Gärdenfors. *Knowledge in Flux: Modeling the Dynamics of Epistemic States.* Bradford Books, MIT Press, Cambridge Massachusetts, 1988.
4. Steve Grand. *Creation: Life and How to Make It.* Weidenfeld and Nicolson, 2000.
5. Steve Joordens. Project Cetacea: A Study of High-Level Cognition In Toothed Whales. `http://www.psych.utoronto.ca/~joordens/courses/PsyD58/Cetacea.html`.
6. Pavlos Peppas, Abhaya Nayak, Maurice Pagnucco, Norman Foo, Rex Kwok and Mikhail Prokopenko. Revision vs. Update: Taking a Closer Look. In Proceedings of the 12th European Conference on Artificial Intelligence, 1996.
7. Mikhail Prokopenko and Marc Butler. Tactical Reasoning in Synthetic Multi-Agent Systems: a Case Study. In Proceedings of the IJCAI-99 Workshop on Nonmonotonic Reasoning, Action and Change, 57–64, Stockholm, 1999.
8. Mikhail Prokopenko, Marc Butler and Thomas Howard. On Emergence of Scalable Tactical and Strategic Behaviour. In RoboCup-2000: Robot Soccer World Cup IV, Springer, 2000.
9. Mikhail Prokopenko, Peter Wang and Thomas Howard. Cyberoos'2001: 'Deep Behaviour Projection' Agent Architecture. In RoboCup-2001: Robot Soccer World Cup V, Springer, 2001.
10. Mikhail Prokopenko and Peter Wang. Relating the Entropy of Joint Beliefs to Multi-Agent Coordination. Submitted to the RoboCup-2002 International Symposium.
11. Claude E. Shannon and Warren Weaver. *The Mathematical Theory of Communication.* University of Illinois Press, 1949.
12. Peter Stone and Manuela Veloso. Task Decomposition, Dynamic Role Assignment, and Low-Bandwidth Communication for Real-Time Strategic Teamwork. In Artificial Intelligence, volume 100, number 2, June 1999.
13. Tomoaki Suzudo. The entropy trajectory: A perspective to classify complex systems. In Proceedings of the International Symposium on Frontier of Time Series Modeling, Tokyo, 2000.