A Reflection on Complexity *with Dana Kulić*

Sarah Bonnemaison

Organicism is very interesting. On the one hand, it is a philosophy of aesthetics that describes a way to learn from nature, to invent and create beautiful objects. Organicism also covers another trend of thoughts in the sciences that is still relevant today. It is a way of thinking about systems that builds one system on the next in increasing complexity. You work with Philip Beesley to build beautiful things, yet you have your own method of working, so I am interested to know how you think it relates to organicism.

Dana Kulić

Actually when I was reading the text about organicism and embryology, I thought it was interesting because it very much connects to some of the discussions that have been going on in robotics, and also in other research on autonomous agents and systems. When we are developing a control algorithm for controlling and directing many robots, should that system be centralized so there is one central master agent that is directing the activities of all of the agents, or should the system be decentralized so that each agent has their own control algorithm and some methods of communication, so that through those individual actions and the communications, a group behaviour emerges automatically? Researchers who have been proponents of this kind of distributed method have been very much inspired by biology. A few notable inspirations are ants and bees. These are very simple insects. Each individual has a small brain with limited capability, yet through communications between the agents the group can generate very complex behaviour, and very complex structure building can emerge.

And of course, learning from nature can take place at different levels of abstraction. We can think of learning from nature right at the neural level. Some researchers are trying to replicate exactly the brain's computational structure at that neuron level, creating computational models of the brain. There is a very famous researcher here at Waterloo, Chris Eliasmith, who is developing and testing such models. Other researchers are taking the inspiration from nature at a higher level, at a more functional level. Artificial agents and robots don't have the same neuronal structure as the biological system but we wish to implement the same or similar functionality. Instead of having neurons, robots have some programed behaviours that are implemented in silicon. There aren't neuron per se, but still, the kind of behaviours and the kind of communications are designed to emulate what we observe in the natural system.

In your work, are you trying to analyze something in nature and trying to copy it? How do you work?

So this question of centralized versus distributed is actually something we have been investigating in our collaboration with Philip. When you observe these systems, one of these installations, you can conceive of it as being a single entity. But because it is spatially distributed and consists of many different elements, you can also conceive of it as a group of agents, as a group of individuals that are grouped together in a kind of a forest. Both of these are computationally possible. When we are thinking about how do we design an algorithm that will control the behaviours of this system, we can design a centralized algorithm where you have a single algorithm that observes all of the sensors of the system. Based on that observation the centralized algorithm generates actions of the system, acting like a single entity. On the other hand, you can also conceive of the system as a group of agents, then you have a separate controller that only uses the sensor data of that single agent; it only controls the action of that single agent. And then you have some communications between the agents to create coordinated behaviours. It is something that we are very much investigating within the context of the work of LASG.

I spoke with Philip recently about the one that he is building right now, *Amatria*. He said it is more centralized.

Yes. But there are different ways of thinking about centralization. You can think about centralization at the implementation level, and centralization at the functional level. When you have a centralized system, one advantage is that the single agent has access to all the information. It can make the best decision because all of the information is available. It can make a globally optimized decision. On the other hand, having a centralized agent introduces a big potential failure mode. Because if the central agent fails, the whole thing is dead. When you think about it in terms of an exhibit, it is a big risk because you don't want to have a case where all of a sudden, you have the room full of visitors and the system is just not working.

Yes, it's like having only one plug.

Exactly. Now on the other hand, when you have a decentralized system, the advantage is that if any one of the agents, or units fails, it is almost imperceptible to the observers because everything else continues to operate. Two or three actors might not be operating, but you don't notice. Because the other agents are continuing to operate, it might appear that the failed agents are just in idle mode right now. But on the other hand, you have these additional costs because no one agent has all the information, so you have additional communication architecture that you require. You need to have these individuals communicating with each other – and that's a cost.

And also visually. There are a lot of little wires. Like a lot of guts – so to speak.

Exactly. And then maybe the decisions you make are not optimal because no agent has all the information. Every agent is just making a decision for itself locally but there is no guarantee that this local decision making will lead to a globally optimal decision. I should mention that part of the wiring is also to provide power.

With the system that we've been investigating, because no design is a clear winner – both designs have strengths and weaknesses – we are hedging our bets. We have the capability to implement both centralized and decentralized systems. We have a centralized system, and all of the wiring that supports all the data coming to a central server where it's been processed, and all of the command messages come out from that server. But then in each node, we also have local capability. So if the local system is disconnected from the central system, it can still operate autonomously. It doesn't just hang there doing nothing. The current system is kind of a hybrid.

Another advantage of having a central system, particularly with the new Indiana University installation *(Amatria)*, is the collaboration between the information science researchers who want to be able to visualize all the information that is happening in the system. This kind of data logging is facilitated if you have a central server where all the data is coming in, and then you can log it from there. That's also part of the reason in the Indiana system we have a centralized architecture, but the design of the architecture is actually flexible. We can have a decentralized system where the only centralized part is the data logging. The actual interactive behaviours and online interaction are decentralized. That is possible with the architecture that we're building.

Could you find the parallel like that in nature in this hybrid system? For example, the bees – if the Mother Bee dies, the whole thing is over.

The tendency is more towards decentralized systems, because decentralized systems tend to be more resilient. In nature, organisms are optimized for being resilient. Whereas since we're designing artificial systems, usually what we do is that we control environments very carefully, so the need for resilience is lessened.

Maybe a good example are organisms that live in protected environments. They lose their defensive capabilities because they are not required. Then when you have invaders, usually introduced by human visitors, that ends up in a bad outcome for those organisms.

Actually, I think humans are sort of like that. We are not as skilled as we used to be at defending ourselves because we are so protected. We could argue there are different levels. One is about the network, and your data collecting is about memory.

Exactly. The third advantage of a centralized system is ease of maintenance. For example, if we want to update our software implementation in the algorithm, having a centralized system would allow us to upload the software to the all of the nodes simultaneously. This is much easier in a centralized system.

I think your hybrid system sounds very interesting. How would you compare this with the article that you just read?

I'm very much on the side of the organicists in the debate. The kind of behaviours that can emerge from individuals collaborating is something that cannot be observed from the behaviour of a single individual. We see that also in some of the experiments that we have done with distributed systems. We implemented this version of a learning agent. Comparing our current robotic system to nature, one of the things that robotic systems and artificial systems lack is the ability to evolve, to improve, to repair, and to regain functionality over time. If you think about each individual bee, each individual ant, they get injured during their activities and they recover. They might learn where the new interesting food source is. Not just learn that for themselves, but also convey that to the rest of their group.

What we've been working on is this idea of imbuing each agent with a learning mechanism, so it can learn over time. We have each agent learning, and then we have some communication between the agents. So how does that impact what they learn, and how does the behaviour of the system as a whole evolve?

It turns out this communication is very important. If you don't have communication versus having some kind of communications, the evolution of the whole system looks very different. This is something very much on the side of the organicists. In our simple experiment, you get behaviours when you have communications among the individuals. You get behaviours that you would not observe, if you just have individuals that do not communicate with each other.

So, what happens when they don't communicate? Do they just keep the information for themselves? They keep on making the same mistake?

Just to give you an example, imagine one of Philip's installations. You have a person walk into the space. Each system in the installation is a learning agent, and it's been learning the relationships between its inputs and its outputs. What it learns when there is no visitor in the scene is the very simple relationship. If I turn the light on, my light sensors will be brighter. If I move my shape memory alloy, I will see a change in movement in the accelerometer. And no matter what I do, the infrared sensor which is sensing the presence of occupants will not change in value. Now a person walks in, and all of a sudden all of those models become violated. Because now the person is modifying the values of the proximity sensor. Now they move closer and further from the system. And if they touch the sculpture, they are changing the readings on the accelerometer, which wasn't happening before.

To learning agents, this is now a huge opportunity to learn. It starts generating actions to try to understand the new state of the environment when the visitor is present. If you have no communications between the agents, then only the agent that is nearby the person is involved. Because to all of the other agents, everything looks the same. When you have communications between agents, now the agents that start learning start communicating to their nearby agents. The neighbours are then also receiving new information, so they also start generating new behaviours. And now the visitor's interest might be also attracted not to this agent but to the nearby agent that is also activated. As a result, you have emergent behaviours between the agents and the visitors. Because the visitors might draw their attention to the nearby agent, might walk over there. So what you learn is different, how these systems behave is different, the entire evolution is very different when you have communication versus no communication.

It's something that we experiment with to find out what the best architecture is. We've experimented with spatial neighbours, functional neighbours, or just random neighbours. What we found with some preliminary experiments was that for visitors, the spatial one seemed to be the most interesting. It gave the impression of something coherent. First you are here, then the thing that was right in front of you would respond, and then the neighbouring things would respond. This looks very similar to some of the things that Philip has implemented as preprogrammed behaviours where you see this rippling outwards. However, we didn't preprogram that. That was an emergent behaviour from having this group of individuals, each with their own algorithm and communicating with each other.

Each algorithm is the same. But it is a learning algorithm, so it's possible that different agents learn different things even though they would be running the same learning algorithm. A good example would be when you have an agent located right in the centre of the sculpture where there would be a lot of traffic. They might learn something different from agents somewhere in the corners that might not be accessible to people. So, no matter what this thing does, because people can't actually walk over there, it won't cause people to stand nearby because it is not accessible versus one that can attract people to stand nearby, which will learn different things. It is very much something that you can observe in biology; if you think about children, they learn based on the environment that they are in. If you have a baby who has English parents, that baby will learn English. A baby who has French parents will learn French. Just because that is what they are exposed to in their environment.

My main research question is, how can we develop systems that can be autonomous over the long term?

And in that respect, this idea of complexity is interesting. I don't think it is a straightforward matter of the longer you survive, the more complex you get. Certain environments foster complexity, and certain other environments degrade complexity. So if you have a very severe environment where it is harder to survive, then complexity is lessened. But if you have a very nurturing environment, then greater complexity can emerge. If you want to have a system that can survive for a long time, it needs to be able to modulate its complexity based on how much support there is in the environment.

In this installation, there are regions that are simpler while some are more complex, where there is more foot traffic from visitors. Are you interested in having anything react to the larger environment? Like the movement of the sun moving, the moon?

Yes, I am always pushing for more sensors. But of course, sensors are expensive. But this is exactly my idea, we want the system to learn not just from the human interaction but also from everything in the environment. Lots of the sensors, for example, the ones that can sense both the light generated by the sculpture itself (so in that sense it is a proprioceptive sensor) and the light changes in the environment. If the system is installed in a place that has natural light, you can sense the changes between daytime, nighttime, sunny day, cloudy day.

The same thing for microphones – of course you will hear the sound of people speaking. If you have an outdoor installation, you will hear the sound of the wind through the sculpture. I think that would be something very beautiful and something that could open the door to new emergent behaviour. Maybe if it's very windy, maybe you want to keep the fronds closed in, so they don't get damaged. And if it's not windy, you can open them up. Wind, for example, can be sensed by the accelerometer because the wind moves the system. All of these learning algorithms, they are learning the relationship between the sensors and the actuators. It is really kind of agnostic to the environment – anything that can cause a change in perception can be used as signals for learning.

Dana Kulić received the combined B. A. Sc. and M. Eng. degree in electro-mechanical engineering, and the PhD degree in mechanical engineering from the University of British Columbia, Canada, in 1998 and 2005, respectively. From 2002 to 2006, Dr. Kulić worked with Dr. Elizabeth Croft as a PhD student and a post-doctoral researcher at the CARIS Lab at the University of British Columbia, developing human-robot interaction strategies to quantify and maximize safety during the interaction. From 2006 to 2009, Dr. Kulić was a JSPS Post-doctoral Fellow and a Project Assistant Professor at the Nakamura-Yamane Laboratory at the University of Tokyo, Japan, working on algorithms for incremental learning of human motion patterns for humanoid robots. In 2009, Dr. Kulić established the Adaptive System Laboratory at the University of Waterloo, Canada, conducting research in human robot interaction, human motion analysis for rehabilitation and humanoid robotics. Since 2019, Dr. Kulić is a professor at Monash University, Australia. Her research interests include robot learning, humanoid robots, human-robot interaction and mechatronics.