2019 8th International Conference on Affective Computing and Intelligent Interaction Workshops and Demos (ACIIW),
Cambridge, United Kingdom, 2019, pp. 374-380.
DOI: 10.1109/ACIIW.2019.8925031

The Effects of Affective Social Bonds on the Interactions and Survival of Simulated Agents

Imran Khan, Matthew Lewis & Lola Cañamero
Embodied Emotion, Cognition, And (inter-)Action Lab
Department of Computer Science, University of Hertfordshire
Hatfield, United Kingdom
i.khan9 | m.lewis4 | l.canamero | @herts.ac.uk

***Abstract*—The formation and maintenance of affective social bonds plays a key role in the well-being of social agents. Oxytocin has been correlated with social partner preference, and it is hypothesised to influence prosocial behaviours. In this paper, we investigate the effects of modulating the preference of affective social bond partners through oxytocin during decisions related to food-sharing and grooming, in a society of simulated agents with different dominance ranks. Our results show survival benefits for agents with affective social bonds across a number of groups with different bond combinations. We observe a number of emergent social behaviours and suggest that our results bear some similarity with behaviors observed in biological agents.**

***Index Terms*—social bonds, oxytocin, embodied affect, social allostasis, social interaction, simulation**

I. INTRODUCTION

Affective social bonds are at the basis of social emotions. Throughout the evolution of all social species, adaptation and survival has largely been based on the ability to form and maintain meaningful relationships with others – to cooperate, protect and compete with others. The maintenance of affective social bonds for social agents plays an important role for survival. For instance, [1]–[3] have shown that, in a variety of species, individuals who maintain close social bonds are associated with longer, healthier lives. Social bonds act as a support system, and may be considered to be buffers against environmental stressors [4].

The neuromodulatory hormone, oxytocin (OT) is consistently linked with social behaviours (both prosocial and anti-social) in human and non-human animals. OT has been linked with trust [5], cooperation [6], generosity [7], reward-sharing [8] and empathy [6], [9], as well as anti-social behaviours such as defensiveness [10] and envy [11]. It is now widely understood that the effects of OT on social behaviours are context-dependent [12]–[14]. OT has been seen to play an important role in the formation of social bonds [15] and partner preference [7], [16]. Partner preference also exists in the presence of higher-ranked individuals [17], who may present more obvious survival-related benefits such as access to food. OT has also been seen to be associated with social bond cohesion and tolerance within animals [10], [18]. As [19] mentions, OT levels may be seen as a "biomarker for the valence and strength of the relationship between interaction partners." For an in-depth review on OT and social relationships, see [19].

It is hypothesised that OT may be responsible in activating a positive feedback loop of social behaviour. For instance, OT may promote the execution of one social behaviour, releasing more OT which will then increase the likelihood of performing that behaviour again [18], [20], creating affective states that may be related to cohesion, trust, or partner affiliation [18]. Alternatively, OT release may be associated with a reward system which may also influence future behaviours [21], [22]. In this way, OT mechanisms contribute to the rise of emotional states within social agents. The extent of this contribution, however, is still unknown.

For social agents to remain viable in the long-term, they must be able to adapt their physiology and behaviours to dynamic physical and social conditions. Adaptive mechanisms such as allostasis [23] (and, by extension, social allostasis [24]) permit this by adapting an individual's homeostatic mechanism [25] in response to, and anticipation of, changing world conditions. Given OT's facilitation of social behaviours, it is one hormone that is believed to play a role in social adaptation through social allostasis [24]. We use a model of OT within our work [26], [27] to explore mechanisms of social adaptation in dynamic world conditions.

In animal populations, dominance rank is associated with increased eating and grooming opportunities [28]. For social agents, survival and well-being is also faciliated by the maintenance of close social bonds [1]–[3]. The ability to control behaviours based on perception of relative rank, along with the presence of close social bonds, appears to be a crucial aspect of survival in a group setting [28] [17]. This trade-off is seen in [17], and is likely faciliated by mechanisms of OT [17], [19], [29].

Building on the findings by [17], we investigate the effects of modulating social bond partner preference through OT mechanisms in a society of agents with different dominance ranks. We use an artificial life approach to simulate the effects of OT mechanisms observed in several relevant chimpanzee studies [17], [18], [30]–[32], and by varying the social bond combinations of agents of different dominance ranks. We model one proposed effect of OT – that it promotes a positive affective loop for social behaviours [18], [20], [21] – and hypothesise that this positive affective system will benefit the

---

Imran Khan is supported by a Ph.D. studentship from the University of Hertfordshire, UK.

viability of socially bonded agents through increased access to food resources. We achieve this through a new iteration to our long-standing Action-Selection Architecture [33] (ASA) for embodied agents, which we have called the Social Assessment Component (SAC).

The rest of this paper is structured as follows: Section II describes the details of the simulated environment and agent model that we have used in our experiments. Section III discusses the experimental methodology and highlights our main results. We then discuss these results and observations in Section IV. Finally, we summarise this paper and propose future work in Section V.

II. SIMULATION

A. Environment

The simulated environment is set up as a bounded, two-dimensional world of size 99 * 99 patches. The outer walls of the world are represented by blue patches. Food resources are represented by yellow circles, and agents are represented by doughnut shapes of varying size. The size of an agent is a function of its dominance rank (Fig. (4)) and is determined at the beginning of each simulation. Food availability is sparse; existing in only two opposing corners in the world. Food size is dynamic, starting at a fixed size and reducing as the resource is consumed over time. It has a small growth rate of 1×10^{-3} per time step, regenerating over time when not being consumed. If both food resources are depleted before the simulation run ends, one additional food resource is regenerated in the top-right corner. Fig. (1) shows the set up of the world.

B. Agent Model

Extending on our previous work [26], [27], which used the long-standing model developed by [34] we endow each agent with an Action-Selection Architecture (ASA) that provides them with a real-time decision-making process, driven by maintaining the stability of two homeostatically-controlled variables. Each variable is responsible for driving one single motivation which are each responsible for driving their own behaviour. Table (I) shows the relationship between each variable, motivation and behaviour.

Agents have two homeostatically-controlled variables that they seek to maintain: $Energy$ and $SocialNeed$. Both of these variables have a lower (0), upper (1), and ideal value (1). Both variables experience a decay rate of 3×10^{-3} per time step, with both values moving towards 0 when their motivations are not being satisfied. $Energy$ represents a survival-based variable: when this drops to 0, an agent will die. On the other hand, while $SocialNeed$ drives a motivation, it is non-critical: agents will not die if this value drops to 0.

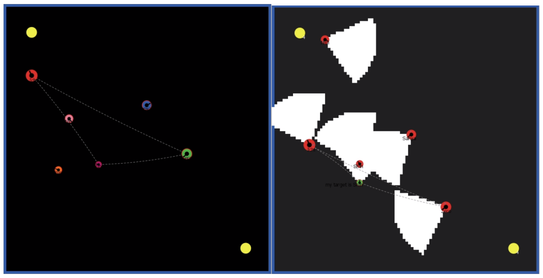

Fig. 1. Screenshots of the simulated world. Left: The simulation during initialisation. Right: The simulation during a simulation run. Agents are represented as doughnut-type shapes of varying size. White cones around agents represent their cone of vision. Yellow circles represent food sources.

TABLE I
RELATIONSHIP BETWEEN AGENTS' TWO INTERNAL VARIABLES, THEIR CORRESPONDING MOTIVATION, BEHAVIOURS AND THE STIMULUS REQUIRED TO EXECUTE THE BEHAVIOUR.

Internal Variable	**Motivation**	**Behaviour**	**Stimulus**	**Physiological Effect**
$Energy$	Hungry	Eat	Food	$+Energy$
$Social$ Need	Lonely	Groom	Agent	$+SocialNeed$

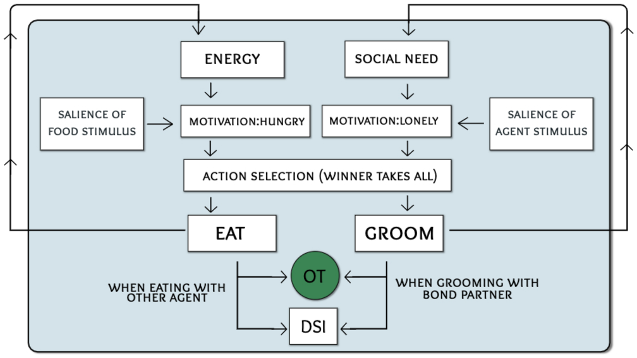

Fig. 2. The Action-Selection Architecture within the agents, used to select a behaviour that best satisfies an agent's physiological variables ($Energy$ and $SocialNeed$). The consummatory behaviours of Eat and $Groom$ satisfy one internal variable each (Table I), and also affect OT release and DSI strengthening in some cases (Table II).

The difference between the ideal value and the respective variable's current value produces an error which, along with the perceived availability of any external stimuli, is used to calculate the intensity (also seen as the "urgency") of each motivation. Agents then select the behaviour that can best satisfy that motivation. Fig. (2) gives a visual representation of this Action-Selection Architecture. Full notation of this model can be seen in [26].

This architecture is extended in this paper by introducing the first iteration of the Social Assessment Component (SAC). Whilst the ASA is responsible for selecting a winning motivation and the relevant behaviour to execute, the SAC is used to determine two aspects of behaviour execution: (a) whether to approach (and share), or avoid an occupied food resource (when the winning behaviour is Eat), and (b) which potential partner to select for grooming (when the winning behaviour is Groom). It achieves this by allowing each agent to take into account another agent's relative dominance rank, the existence and strength of any bond between them (called the DSI, described below), and the amount of OT within an agent's physiology. In this way, agents do not simply execute winning behaviours, but instead consider their local, social environment to determine how best to execute it. The specific effects of these SAC calculations are described in detail below. This can be seen in Fig. (3).

The dominance rank of an agent is a value representing its rank within the social hierarchy. Higher dominance rank grants an agent easier access to food resources. Higher-ranked agents are also more desirable to be selected by others as a grooming target.

The value of OT itself is increased as one of the social behaviours have been executed, which then increases the valence of other bonds; creating a positive feedback loop for bond preference. This is in line with the OT mechanisms proposed by [18], [30]. The strength of the bond (DSI) is also increased as these behaviours are executed.

We model a bond between two agents with a flag and a value (DSI) that represents the strength of the dyadic relationship between them. An existing bond is akin to a stable, mutually-positive relationship between the two agents, which biases decision-making in favour of bond partners through the two components of the SAC, discussed below.

1) SAC 1: Food Approach or Food Avoidance: Agents who look to approach and eat at a food resource take into account any other agents that may be at the food resource (called the "food owner", as described by [32]), and determines whether an agent will Approach that food resource, or Avoid it by turning around. By default, food owners (O) currently observe a passive approach towards approachers (A) of food, (considered a "neutral" approach to food sharing, as described by [32]). Approaching agents (A) decide whether to Approach or Avoid a food resource by using the relative rank between them (A) and the food owner (O), and using their OT to lend weighting to any dyadic bond (DSI) that may exist between them (Fig. (3)).

The approach value of agent A for food occupied by agent O is given by:

$$ApproachVal = rank_{diff} + (bond_{AO} \times (DSI_{AO} \times OT)) \tag{1}$$

where $bond_{AO}$ is the existence of a bond between agents A and O (1 if a bond exists, else 0), and DSI_{AO} is the strength of the dyadic bond between agents A and O. Agents have a fixed Approach threshold ($\theta = 0$): $ApproachVal \geq \theta$ will lead an agent to Approach the resource, $ApproachVal < \theta$ will lead an agent to Avoid the resource.

In the absence of a social bond between them, a lower-ranked agent cannot approach and access a food resource occupied by a higher-ranked agent. Additionally. the highest-ranked agent at a food resource becomes the food owner, regardless of whether or not they were the first to that food resource. In this way, dominant agents take over food ownership, and any other approaching agents will then calculate their $ApproachVal$ using the strength of the new food owner.

2) SAC 2: Dynamic Grooming Target Selection: Agents use a flexible decision-making process to assign a value to each available grooming partner in its vision. Using the SAC, agents (A) that are looking to Groom with a potential partner (i) will assign a normalised (0-1) rank $rank_{relative}$ to each agent in its field of vision in 0.25 increments: a value of 1 is assigned for the highest-ranked available agent, 0.75 for the second-highest-ranked agent, and so on. Agents also take into account the strength of any dyadic bond that may exist between them. An agent's OT value is then used to modulate the valence of a bond, and a final $PartnerValue$ is calculated for each agent. The visible agent with the highest $PartnerValue$ is selected as the grooming target.

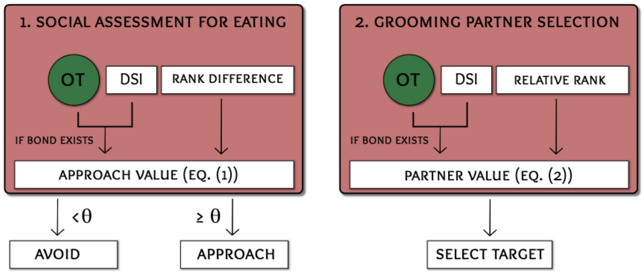

Fig. 3. Visualisation of the calculations within the SAC. Left: The assessment agents use when they decide to Avoid ($\theta < 0$) or Approach ($\theta \geq 0$) a food resource occupied by another agent. Right: The calculation used to select a grooming partner. The agent with the maximum $PartnerVal$ is selected.

$$PartnerVal_{Ai} = (rank_{relative} + (bond_{Ai} \times (DSI_{Ai} \times OT))) \tag{2}$$

where $bond_{Ai}$ is the existence of a bond between agents A and i (1 if a bond exists, else 0), and DSI_{Ai} is the strength of the dyadic bond between agents A and i. In the absence of any bonds, agent (A) will always look to groom with the highest-ranked partner (i) available.

3) Role and Release of OT: In this model, OT modulates the preference of any existing social bond partners in one of two situations: deciding to approach food or selecting a grooming partner. The initial value of OT is set to 1. It has a lower bound of 0, but no upper bound as we looked to magnify the effects of these mechanisms within the model. The value of OT is subject to a decay rate of 5×10^{-3} per time step, and is endogenously released within agents at non-uniform rates depending on the context of behaviour execution. These contexts relate to the type of social behaviour (food sharing or grooming) and the existence of a social bond between the two agents.

4) Social Bond Strength: Dyadic Strength Index: The strength of a dyadic bond between a pair of agents is represented through a measure called the Dyadic Strength Index (DSI). Adopted from work done by [35], the DSI is a dynamic measure of bond strength between dyads which is updated after every social interaction. In this model, a social interaction is defined as either food sharing or grooming.

Socially-bonded agents have an internal DSI value for each of their two respective bonds, and use this value within the SAC calculations for both their food approach and partner

TABLE II

CONTEXT-DEPENDANT EFFECTS ON OT RELEASE AND DSI. THE MAGNITUDE OF EACH EFFECT DEPENDS ON THE BEHAVIOUR BEING EXECUTED, AND WHETHER THE TARGET IS A BOND PARTNER OR A NON-BOND. EACH ARROW REPRESENT A WEIGHT OF 0.25 FOR EACH BEHAVIOUR, USING THE CALCULATION FROM [35].

Behaviour	Bonded?	OT Effect	DSI Effect	Stimulus
Eat	-	-	-	Food
Eat	Bond	↑↑	↑↑	Food
Eat	Non-Bond	-	-	Food
Groom	Bond	↑	↑	Agent
Groom	Non-Bond			Agent

selection decisions (Fig. (3)). The DSI value is bi-directional, i.e. the strength of the bond from A3 to A4 is the same as A4 to A3. Like OT, the DSI of bonds experiences a decay rate over time if it is not renewed, and can be strengthened at a rate that is determined by the social behaviour being executed. The specific effects on the DSI can be seen in Table (II).

C. Agent Interactions

1) Agent Movement: Agents perceive the world through an 80 degree cone of vision of length, 20 patches in length. When an agent sees a stimulus that is relevant to their current motivation, it will face the stimulus and keep it in the centre of its cone of vision as it moves towards it. When looking to groom with another, agents will follow other agents by keeping them in their vision.

2) Passive Food Sharing: "Food sharing" is implemented as a passive approach by food owners (O) to an approacher's (A) advances. No direct facilitation of food transfers take place from (O) to (A), nor does (O) exhibit any socio-negative behaviours to approaching behaviours.

3) Grooming: Grooming is implemented as a single time step of tactile feedback between two agents. Agents can groom with another when they are in radius 2 of their selected partner. Grooming is unidirectional: only the physiology of the agent performing the grooming is affected through this behaviour.

III. EXPERIMENTS AND RESULTS

We modelled a small society of agents $(n = 6)$ that use a flexible decision-making process when looking to execute two prosocial behaviours: approaching food for food sharing and partner selection for grooming. In these situations, an agent's OT level modulates the preference of socially-bonded agents. Agents could either be bonded with two other agents or unbonded. Of the six agents in the environment, three would be bonded with each other, with the remaining three unbonded. These bonds would be pre-determined prior to the simulation run. Agents are motivated by the maintenance of their internal variables ($Energy$ and $SocialNeed$). When $Energy$ drops to 0, an agent will die.

In line with the previous iterations of this model [26], [27], experiments were performed in a simulated environment using the NetLogo platform, version 5.3.1 [36]. Simulation models provide several experimental and computational advantages for embodied agents which are discussed in detail in [27].

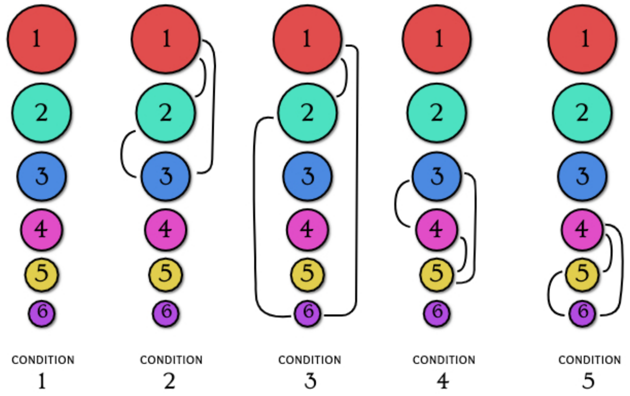

Fig. 4. Visual representation of bonded agents within each condition. Agent size is relative to dominance rank: 1 being the most dominant agent, 6 being the least dominant.

A total of five different affective social bond combinations were investigated in a world condition with scarce food resources (Fig. (1)), with three bonded agents per condition (Fig. (4)). The choice of bond combinations was inutitive, and we selected combinations where we would expect to see some differences (for instance, bonding the three most-dominant agents in Condition 2 and the three least-dominant in Condition 5). We report aggregated results across all simulation runs using Life Length as our primary viability indicator, as well as qualitative reports by watching the simulation runs. Life Length is defined as an agent's length of survival (keeping $Energy > 0$) as a percentage of the maximum simulation time.

Thirty simulation runs were performed for each bond combination, with a cut-off at 20,000 time steps. A time step is defined as one single iteration of the code, which translates graphically as a single frame update within the NetLogo GUI. This results in 150 simulation runs, which equated to approximately 10 hours of experimental runs in real-time. Data was captured at each time step and aggregated across for all simulation runs.

Tables (III) and (IV) show the results for individual life length and group life length respectively, averaged across all simulation runs. Across all conditions, we see Life Length of socially bonded agents to equal or outperform Control conditions (Condition 2: +16%; Condition 3: +11%; Condition 4: +9%; Condition 5: -1% (non-significant)). We identify some detrimental impact to the life length of unbonded agents vs. Control. This is through increased life length of more dominant agents (Condition 2), or the combined opposing effects of particular agents (A3 and A6) in conditions where the former is bonded (Condition 4). We also note that food consumed is a perfect indicator of overall life length (r=0.996).

The main results are as follows:

- In the Control condition, we see the two most dominant agents (A1 and A2) with the longest life lengths of all agents (A1: 55%; A2: 54%). However, A3 (the third-most dominant agent in the environment) had the shortest life

TABLE III
RESULTS OF LIFE LENGTH FOR INDIVIDUAL AGENTS ACROSS ALL CONDITIONS.

Agent	**Condition**				
	Control (1)	A1-A2-A3 (2)	A1-A2-A6 (3)	A3-A4-A5 (4)	A4-A5-A6 (5)
A1	55%	59%	59%	57%	54%
A2	54%	57%	59%	50%	55%
A3	39%	57%	47%	49%	46%
A4	49%	46%	44%	48%	47%
A5	44%	41%	43%	47%	42%
A6	43%	40%	51%	36%	46%

TABLE IV
RESULTS OF LIFE LENGTH FOR BONDED AND UNBONDED AGENTS ACROSS ALL CONDITIONS. CONTROL RESULTS DENOTE LIFE LENGTH OF THE BONDED AGENTS PER CONDITION WHEN NON-BONDED IN CONTROL.

Agent Group	Condition			
	A1-A2-A3 (2)	A1-A2-A6 (3)	A3-A4-A5 (4)	A4-A5-A6 (5)
Bonded	58%	57%	48%	45%
Unbonded	42%	45%	48%	52%
Control	50%	51%	44%	45%

length of all agents (39%).

- In Condition 2, where the three most dominant agents are bonded, life length of bonded agents improved by 16% compared to when they were non-bonded (Control). All bonded agents had increased life lengths (A1: +7%, A2: +4%, A3: +36%). There was a 36% difference in survival between bonded and unbonded agents (A4-A5-A6) in this condition.
- In Condition 3, we note a +19% increase in life length for A6 (the least dominant agent) when bonded with the two most dominant agents. We also see increases in life length for A1 (+8%) and A2 (+9%). Interestingly, we also see an increase in A3's life length (+20%) in this condition, despite being unbonded.
- Condition 4 saw overall group survival improve by 9% vs. Control. The least dominant agent, A6, saw an 18% drop in life length vs. the Control condition, while the lowest-ranked bonded agent (A5) saw a 12% increase in life length.
- Condition 5 saw the three least dominant agents bonded. Bonded agents saw no improvement in life length vs. their non-bonded performance in Control. This was a result of decreased life lengths of A4 (-4%) and A5 (-6%), but an increased life length of A6 (+8%).

IV. DISCUSSION

As [18] mentions, mechanisms of OT that promote prosocial behaviours may create a potential affective state between two agents through a positive feedback loop, and we have modelled basic mechanisms of this effect. We rely on additional data through observations to understand some of the potential affective behaviours that emerge from these mechanisms. We use these observations along with the life length data for a more holistic view of our findings, and to consider some additional hypotheses of mechanisms for social emotions. We propose that OT mechanisms that promote affective social bond preference may play a role in similar behaviours observed in biological agents, and that researchers of embodied agents should take such low-level mechanisms (such as hormones and social bonds) into consideration when studying (the emergence of) social emotions.

A. Dominance Rank is not Indicative of Survival

While we expected viability to be roughly in line with an agent's dominance rank in the control condition (due to an increased access to food, in theory), we found that this was not always the case. We highlight an interesting outcome of A3 (the third highest-ranked agent) in the control condition: that it actually survived for the least amount of time and that lower-ranked agents consistently outlived it. Watching the simulation runs, we find this to be due to a combination of a number of emergent behaviours. As there are only two food resources in the environment, they would often be occupied by the most dominant agents (A1 and A2) leaving A3 often wandering between resources until one could be approached. Furthermore, as agents are attracted to the highest-ranked available agents for grooming, A3 would follow A1 or A2 as it looked to groom. Often, A1 or A2 would find food resources. A3 would be unable to approach them due to the difference in rank, and would be forced to avoid that resource (see our Supplementary Material [37]). As A3 dies, the competitive social dynamics are greatly shifted for the less-dominant agents, and their access to food resources increases. This is further emphasised in results from Condition 3 (A1-A2-A6 bonded), where A1 and A2 would share food resources, allowing A3 to eat at the other available food resource. In this case, an absence of competition between more dominant agents may benefit others (such as A3).

B. Affective Social Bonds Lead to Increased Survival, even for Dominant Agents

Across three of the bonded conditions, we see that having affective social bonds increases the overall survival length of all agents with such bonds. This was also seen in conditions where the two most dominant agents (A1 and A2) were bonded with each other (Conditions 2 and 3). This is an interesting finding, particularly given that their dominance rank grants them priority access to food resources.

We consider whether a tolerant approach between the two most dominant individuals in a society is actually more beneficial to their other bonded partners, and even (some members of) the wider society. Considering the findings from Condition 3 (A1-A2-A6), the lack of competition between A1 and A2 greatly increased A3's average life length. Conversely, the Control condition (where A1 and A2 were not bonded) saw A3 experience its lowest life length. In this regard, these potential elements of cohesion between the most dominant agents may have trickle-down effects on the rest of their society. Further work is needed to understand whether this is an artefact of the world conditions or a result of the modelled OT mechanisms.

C. Dominant Agents Exhibit A "Parent-Child" Dynamic with Least-Dominant Individual

The highest-ranked agents, A1 and A2, would only actively groom with the lowest ranked agent (A6) when bonded with it (Supplemental Material [37]). A6 provides no (grooming or food sharing) benefit for the more dominant agents that the dominant agents could take at will. In some observations, we see that these higher-ranked agents (A1 and A2) would overlook the opportunity to groom with other similarly-ranked, non-bonded agents (such as A3) in favour of its bonded (yet much lower ranked) A6. This is despite the fact that the more dominant A3 would be more successful in reaching a food resource. For A2's viability in particular, this seems counter-intuitive. Therefore, a bond with a much lower ranked agent appears to provide no benefit for the higher-ranked, except for the potential maintenance of the bond itself, but provides benefits for the lower ranked as the affective social bond is strengthened between them. As bond strength is increased, the low-ranked A6 is more likely to access food occupied by the higher-ranked agents (A1 and A2).

This may be likened to an altruistic, parental dynamic; as a strong bond with higher-ranked agents provides the lowest–ranked agent additional opportunities to find and eat food. What is particularly interesting is that this "parent–child" bond also slightly promoted the life length of the higher-ranked agents, although the reasons for this are currently unclear. The underlying OT mechanisms may play a role in such a dynamic, and further observations in biological literature may shed light on this behaviour.

D. Altruism may be Detrimental for Lower-Ranked Agents

When the three lowest-ranked agents were bonded (Condition 5) the least dominant agent (A6) still showed viability benefits to being bonded with higher-ranked agents, albeit their (A4 and A5) relative rank vs. the other agents in the environment. However, A6's improved survival was at the expense of A4 and A5's survival. This suggests that, when an agent is not among the most dominant in a hierarchy, altruistic behaviours through passive food sharing with lower-ranked individuals may not be the ideal approach, particularly when the social and physical world poses a sufficient challenge.

Unlike A1 or A2, A4 or A5 are unable to completely monopolise a resource, which therefore creates a more challenging environment for them. Further work would look to biological literature to observe any potential relationships between dominance ranks and altruistic tendancies of behaviour in animals. We hypothesise that such social dynamics could give rise to anti-social or defensive behaviours in biological agents, even towards bond partners. Specifically, a challenging social environment may influence additional hormone release, such as cortisol, which may affect an agent's perception of social bonds through an interplay of hormones, as suggested by [21] [23] [24]. It may be the case that altruistic behaviours may not be as common in less-dominant individuals.

E. Weaker Agents May Need More Social Reinforcement

In some simulation runs where the weakest agent A6 did not have an opportunity to groom with bond partners early on in their life, their level of OT and DSI would be low. It would bias grooming partner selection to higher-ranked agents in its vision (Supplementary Material [37]): overlooking the availability of bonded (less-dominant) agents due to the low valence of the bonds, in favour of the more attractive, higher-dominant individuals. This lack of reinforcement of affective social bonds through grooming then reduced bond valence during food approach calculations, reducing A6's access to food and, ultimately, a lower life length.

This behaviour driven through low OT appeared to mirror potential states of distrust or a lack of social memory for bonded agents. This may highlight the importance of maintaining a strong bond with lower-ranked individuals, particularly if more dominant individuals are present, else they may attempt to form new relationships with more dominant members of society. Allowing for the emergence of new bonds in our model will help us investigate this further.

V. CONCLUSION AND FUTURE WORK

In this paper, we have investigated the effects of modulating the preference for bond partners through mechanisms of OT in a group of agents with different dominance ranks. We have carried out experiments with small (3-agent) groups of agents with different degrees of bonding, and assessed the effects that different combinations of bonded agents had on the survival of agents and their interactions.

Our results show viability benefits for agents with affective social bonds (particularly for lower-ranked agents) through increased food availability in a food-sparse environment. We note an interesting result: that the third most dominant agent in the group had the lowest rate of survival in control conditions, with all lower-ranked agents consistently outliving it. We also find that the most dominant agents survived for longer when a bond existed between them.

We observe agent behaviour related to altruistic states. In certain conditions, higher-ranked agents overlook grooming with similarly-ranked agents to select a bonded (yet much lower-ranked) partner. The decision to choose this bond partner over an agent with more access to food may be reminiscent of some altruistic, potentially (allo-)parental mechanisms between social agents. Conversely, these "altruistic" tendancies do not appear to be a viable approach when rank differences are much smaller among less-dominant individuals, suggesting that selflessness towards bond partners is not always the most viable approach. Finally, we observed lower-ranked agents exhibiting behaviours related to distrust towards bonded agents when OT was low: they turned towards higher-ranked agents for grooming, regardless of bond strength.

Future work will consider the implications of affectively-negative behaviours on both bond strength and social behaviours, and the formation of new bonds throughout an agent's lifecycle.

References

[1] O. Schülke, J. Bhagavatula, L. Vigilant, and J. Ostner, "Social bonds enhance reproductive success in male macaques," *Current Biology*, vol. 20, no. 24, pp. 2207–2210, 2010.

[2] E. Z. Cameron, T. H. Setsaas, and W. L. Linklater, "Social bonds between unrelated females increase reproductive success in feral horses," *Proceedings of the National Academy of Sciences*, vol. 106, no. 33, pp. 13 850–13 853, 2009.

[3] J. Holt-Lunstad, T. B. Smith, and J. B. Layton, "Social relationships and mortality risk: a meta-analytic review," *PLoS medicine*, vol. 7, no. 7, p. e1000316, 2010.

[4] M. Heinrichs, T. Baumgartner, C. Kirschbaum, and U. Ehlert, "Social support and oxytocin interact to suppress cortisol and subjective responses to psychosocial stress," *Biological psychiatry*, vol. 54, no. 12, pp. 1389–1398, 2003.

[5] M. Kosfeld, M. Heinrichs, P. J. Zak, U. Fischbacher, and E. Fehr, "Oxytocin increases trust in humans," *Nature*, vol. 435, no. 7042, p. 673, 2005.

[6] C. K. De Dreu and M. E. Kret, "Oxytocin conditions intergroup relations through upregulated in-group empathy, cooperation, conformity, and defense," *Biological psychiatry*, vol. 79, no. 3, pp. 165–173, 2016.

[7] J. R. Williams, T. R. Insel, C. R. Harbaugh, and C. S. Carter, "Oxytocin administered centrally facilitates formation of a partner preference in female prairie voles (microtus ochrogaster)," *Journal of neuroendocrinology*, vol. 6, no. 3, pp. 247–250, 1994.

[8] S. W. Chang, J. W. Barter, R. B. Ebitz, K. K. Watson, and M. L. Platt, "Inhaled oxytocin amplifies both vicarious reinforcement and self reinforcement in rhesus macaques (macaca mulatta)," *Proceedings of the National Academy of Sciences*, vol. 109, no. 3, pp. 959–964, 2012.

[9] J. A. Barraza and P. J. Zak, "Empathy toward strangers triggers oxytocin release and subsequent generosity," *Annals of the New York Academy of Sciences*, vol. 1167, no. 1, pp. 182–189, 2009.

[10] C. K. De Dreu, L. L. Greer, M. J. Handgraaf, S. Shalvi, G. A. Van Kleef, M. Baas, F. S. Ten Velden, E. Van Dijk, and S. W. Feith, "The neuropeptide oxytocin regulates parochial altruism in intergroup conflict among humans," *Science*, vol. 328, no. 5984, pp. 1408–1411, 2010.

[11] S. G. Shamay-Tsoory, M. Fischer, J. Dvash, H. Harari, N. Perach-Bloom, and Y. Levkovitz, "Intranasal administration of oxytocin increases envy and schadenfreude (gloating)," *Biological psychiatry*, vol. 66, no. 9, pp. 864–870, 2009.

[12] J. A. Bartz, J. Zaki, N. Bolger, and K. N. Ochsner, "Social effects of oxytocin in humans: context and person matter," *Trends in cognitive sciences*, vol. 15, no. 7, pp. 301–309, 2011.

[13] R. A. Bethlehem, S. Baron-Cohen, J. van Honk, B. Auyeung, and P. A. Bos, "The oxytocin paradox," *Frontiers in behavioral neuroscience*, vol. 8, p. 48, 2014.

[14] C. H. Declerck, C. Boone, and T. Kiyonari, "The effect of oxytocin on cooperation in a prisoner's dilemma depends on the social context and a person's social value orientation," *Social cognitive and affective neuroscience*, vol. 9, no. 6, pp. 802–809, 2013.

[15] H. E. Ross and L. J. Young, "Oxytocin and the neural mechanisms regulating social cognition and affiliative behavior," *Frontiers in neuroendocrinology*, vol. 30, no. 4, pp. 534–547, 2009.

[16] T. R. Insel and T. J. Hulihan, "A gender-specific mechanism for pair bonding: oxytocin and partner preference formation in monogamous voles." *Behavioral neuroscience*, vol. 109, no. 4, p. 782, 1995.

[17] R. M. Wittig, C. Crockford, T. Deschner, K. E. Langergraber, T. E. Ziegler, and K. Zuberbühler, "Food sharing is linked to urinary oxytocin levels and bonding in related and unrelated wild chimpanzees," *Proceedings of the Royal Society B: Biological Sciences*, vol. 281, no. 1778, p. 20133096, 2014.

[18] L. Samuni, A. Preis, A. Mielke, T. Deschner, R. M. Wittig, and C. Crockford, "Social bonds facilitate cooperative resource sharing in wild chimpanzees," *Proceedings of the Royal Society B*, vol. 285, no. 1888, p. 20181643, 2018.

[19] C. Crockford, T. Deschner, and R. M. Wittig, "The role of oxytocin in social buffering: What do primate studies add?" in *Behavioral Pharmacology of Neuropeptides: Oxytocin*. Springer, 2017, pp. 155–173.

[20] C. Crockford, T. Deschner, T. E. Ziegler, and R. M. Wittig, "Endogenous peripheral oxytocin measures can give insight into the dynamics of social relationships: a review," *Frontiers in behavioral neuroscience*, vol. 8, p. 68, 2014.

[21] K. D. Broad, J. P. Curley, and E. B. Keverne, "Mother–infant bonding and the evolution of mammalian social relationships," *Philosophical Transactions of the Royal Society B: Biological Sciences*, vol. 361, no. 1476, pp. 2199–2214, 2006.

[22] D. Scheele, A. Wille, K. M. Kendrick, B. Stoffel-Wagner, B. Becker, O. Güntürkün, W. Maier, and R. Hurlemann, "Oxytocin enhances brain reward system responses in men viewing the face of their female partner," *Proceedings of the National Academy of Sciences*, vol. 110, no. 50, pp. 20 308–20 313, 2013.

[23] P. Sterling, "Allostasis: a model of predictive regulation," *Physiology & behavior*, vol. 1, p. 5–15, 2012.

[24] J. Schulkin, "Social allostasis: anticipatory regulation of the internal milieu," *Frontiers in evolutionary neuroscience*, vol. 2, p. 111, 2011.

[25] W. B. Cannon, "Organization for physiological homeostasis," *Physiological reviews*, vol. 9, no. 3, pp. 399–431, 1929.

[26] I. Khan, M Lewis and L. Cañamero, "Adaptation and the social salience hypothesis of oxytocin: Early experiments in a simulated agent environment." in *Proceedings of the 2nd Symposium on Social Interactions in Complex Intelligent Systems (SICIS), Liverpool, UK*, 2018, pp. 2–9.

[27] I. Khan and L. Cañamero, "Modelling adaptation through social allostasis: Modulating the effects of social touch with oxytocin in embodied agents," *Multimodal Technologies and Interaction*, vol. 2, no. 4, p. 67, 2018.

[28] N. Watanabe and M. Yamamoto, "Neural mechanisms of social dominance," *Frontiers in neuroscience*, vol. 9, p. 154, 2015.

[29] J. C. Mitani, "Male chimpanzees form enduring and equitable social bonds," *Animal Behaviour*, vol. 77, no. 3, pp. 633–640, 2009.

[30] C. Crockford, R. M. Wittig, K. Langergraber, T. E. Ziegler, K. Zuberbühler, and T. Deschner, "Urinary oxytocin and social bonding in related and unrelated wild chimpanzees," *Proceedings of the Royal Society B: Biological Sciences*, vol. 280, no. 1755, p. 20122765, 2013.

[31] A. Mielke, A. Preis, L. Samuni, J. F. Gogarten, R. M. Wittig, and C. Crockford, "Flexible decision-making in grooming partner choice in sooty mangabeys and chimpanzees," *Royal Society open science*, vol. 5, no. 7, p. 172143, 2018.

[32] A. V. Jaeggi, J. M. Stevens, and C. P. Van Schaik, "Tolerant food sharing and reciprocity is precluded by despotism among bonobos but not chimpanzees," *American Journal of Physical Anthropology*, vol. 143, no. 1, pp. 41–51, 2010.

[33] O. Avila-Garciá, E. Hafner, and L. Cañamero, "Relating behavior selection architectures to environmental complexity," in *Proceedings of the Seventh International Conference on Simulation of Adaptive Behavior (SAB)*, 2002, pp. 127–128.

[34] D. Cañamero, "Modeling motivations and emotions as a basis for intelligent behavior," in *Proceedings of the first international conference on Autonomous agents*. ACM, 1997, pp. 148–155.

[35] L. Kulik, "Development and consequences of social behavior in rhesus macaques (macaca mulatta)," Ph.D. dissertation, Univ. Leipzig, 2015.

[36] U. Wilensky. (1999) Netlogo. [Online]. Available: http://ccl.northwestern.edu/netlogo/

[37] I. Khan. (2019) Social bond simulation supplemental material. [Online]. Available: http://socialbonds.imytk.co.uk