Affective Dynamics and Control in Group Processes

Jesse Hoey
David R. Cheriton School of Computer Science
University of Waterloo
Waterloo, Ontario, Canada
jhoey@cs.uwaterloo.ca

Tobias Schröder
Institute for Urban Futures
Potsdam University of Applied Sciences
Potsdam, Germany
schroeder@fh-potsdam.de

Jonathan H. Morgan
Institute for Urban Futures
Potsdam University of Applied Sciences
Potsdam, Germany

Kimberly B. Rogers
Department of Sociology
Dartmouth College
Hanover, NH, USA

Meiyappan Nagappan
David R. Cheriton School of Computer Science
University of Waterloo
Waterloo, Ontario, Canada

ABSTRACT

The computational modeling of groups requires models that connect micro-level with macro-level processes and outcomes. Recent research in computational social science has started from simple models of human behaviour, and attempted to link to social structures. However, these models make simplifying assumptions about human understanding of culture that are often not realistic and may be limiting in their generality. In this paper, we present work on Bayesian affect control theory as a more comprehensive, yet highly parsimonious model that integrates artificial intelligence, social psychology, and emotions into a single predictive model of human activities in groups. We illustrate these developments with examples from an ongoing research project aimed at computational analysis of virtual software development teams.

CCS CONCEPTS

• **Human-centered computing → Social networks**; **Social network analysis**; • **Computing methodologies → Artificial intelligence**; **Network science**; **Multi-agent systems**; • **Applied computing → Sociology**;

KEYWORDS

group interaction, team dynamics, affect, affect control theory, emotion, collaboration, agents

ACM Reference Format:
Jesse Hoey, Tobias Schröder, Jonathan H. Morgan, Kimberly B. Rogers, and Meiyappan Nagappan. 2018. Affective Dynamics and Control in Group Processes. In *Group Interaction Frontiers in Technology (GIFT'18), October 16, 2018, Boulder, CO, USA.* ACM, New York, NY, USA, 7 pages. https://doi.org/10.1145/3279981.3279990

GIFT'18, October 16, 2018, Boulder, CO, USA

ACM ISBN 978-1-4503-6077-7/18/10…$15.00
https://doi.org/10.1145/3279981.3279990

1 INTRODUCTION

The digital transformation of society has opened up novel opportunities for studying social interaction. People increasingly use digital tools and social-network platforms to communicate with each other, producing large amounts of digital data that can be analyzed using "computational social science" (CSS) methods [23] to answer psychological or sociological questions. One of the biggest hurdles facing CSS is a substantial disconnect between disciplines such as computer science or physics, which produce the methodological expertise for large-scale data analysis and modeling complex social systems, and disciplines such as psychology or sociology, which provide a rich landscape of theories and empirical evidence that enable thoughtful use of these methods.

The purpose of the present paper is to discuss the promise of a cross-disciplinary, computational approach to the study of small-group dynamics, and describe how such an approach might proceed using our own research as an example. Importantly, we focus on a theoretically informed way of using computational simulation for the analysis of large amounts of social media data. We briefly review work in developing Bayesian affect control theory (*BayesAct*), which mathematically integrates widely-accepted psychological and sociological theories of social interaction with decision theoretic reasoning, and thus enables the creation of artificially intelligent agents that are aware of social scientific knowledge [19, 38]. We also show some preliminary results from our ongoing research project to illustrate the logic and feasibility of small-group research enhanced by artificial intelligence (AI).

Before turning to our own work, however, we provide a brief review of some developments in AI and computational social science that are relevant, in our opinion, for group-dynamics research in general. We focus on the issue that small groups are complex systems whose analysis requires more sophisticated mathematical tools than the general linear models usually taught in the social and behavioural sciences. We refer the reader to [17] for a more complete exposition and further details.

2 BACKGROUND

Despite a widespread understanding that dynamic interactions are fundamental to understanding small-group phenomena [11], much of classic group research in sociology and psychology has been surprisingly unaffected by the computational modeling techniques developed in other scientific disciplines that deal with dynamical systems. In contrast, artificial intelligence has attempted to model the behaviour of groups computationally for some time. For example, multi-agent systems (MAS) research aims to build teams of robots that can cooperate in working towards common goals, often by invoking strategic behaviours based on rational utility. This approach typically results in computationally complex strategic models that must account for many agents optimizing their utility functions simultaneously. On the other end of the spectrum, agent-based modeling (ABM) aims to replicate the emergent behaviour of groups using simple models of individual human behaviours.

2.1 Multi-Agent Systems

A classic example of multi-agent systems research is the robot soccer "grand challenge" which has aimed (since 1998) to build a robot soccer team that can compete against the best human soccer teams (robocup.org). As with much artificial intelligence research, most MAS work aims to build *rational* agents who are individual utility maximizers. That is, each agent has a set of preferences encoded in a utility function, which it uses to optimize its behaviour by computing expectations with respect to likely future scenarios. In order to enable group behaviours, such rational agents must model other agents' intelligent behaviour, rapidly leading to computational intractability. However, most empirical and theoretical work from social science doesn't support the notion that people put this much effort into social processing.

Rational choice in group behaviour has also been the subject of much investigation in economics and game theory [21]. However, rationality leads to inconsistencies when considering simple games with social interdependence (e.g., social dilemmas). Behavioural economists have tackled the apparent irrationality of human behaviour in social dilemmas with a variety of heuristics, typically resulting in adjustments to utility functions but remaining within the rational model [2, 9, 27, 35]. Cooperative behaviour has been linked to altruism through factors like kinship, direct reciprocity, or indirect reciprocity via reputation [30]. Nevertheless, the fundamental problem persists in that an agent needs to optimize its behaviour by considering all possible strategic behaviours of other agents in order to compute a rational solution. These models lead to shallow (in time) and broad (in number of options considered) solutions due to limited processing power, and fail to provide convincing accounts of human social behaviour at a large scale. Further, it appears that fairness or inequity adjustments may not be comprehensive enough to account for human behaviour across all situations (as games), and a morality concept that is not based on outcomes provides a more parsimonious account [5]. In line with this idea, we argue that by considering identity and group membership as a shared cultural and affective quantity, the forces of the resulting relational commitments bear heavy weight upon the actions of group members [22], and this can be used to account for human behaviour in social dilemmas [20], and may be the type of social intuitionist [13] model of moral reasoning that will explain some of these paradoxes.

2.2 Agent-Based Models and Social Simulation

Agent-based models (ABMs) also consist of autonomous computational agents that interact with each other and thus generate an emergent, group-level outcome. However, in social science applications of ABMs, the goal is to understand and explain the complex behaviours of a social system that are often not trivially reducible to the properties of individual agents [for reviews, see 8, 40]. Many ABMs in social simulation have studied processes of attitude formation and diffusion in groups or societies, emphasizing the importance of social influence between agents [e.g., 7]. The agents in these models are usually very simple; e.g., an ABM might represent an agent's "opinion" as a number on a single dimension that changes according to a simple algebraic rule when subject to "influence" from another agent that is either a "neighbour" on a spatial grid or connected to the agent in a more or less realistic social network. For example, a homophilous agent may tend to change its behaviour to be more similar to the agents it interacts with most often (its "friends" and "co-workers"). The virtue of such models is that they show how even very simple mechanisms can produce complex group-level phenomena that are poorly understood. However, most ABMs lack even the most basic ingredients of intelligence (whether human or artificial), namely the ability to reason, plan, act, and cooperate with other agents.

The psychological simplicity of agents in many social simulations has been the subject of much debate among modelers. It is not helpful to replace the dubious assumption of the full rationality of agents encountered in much AI research with the equally dubious assumption of full stupidity encountered in many ABMs. As the field matures and as computational resources become more ubiquitous, many social simulation researchers have moved to build more psychologically realistic ABMs [28, 31, 41]. One such model is GroupSimulator, developed by [16]. In this model, exchanges of behaviour among a group of computational agents are organized according to the structure of Bales' Interaction Process Analysis (IPA) [4]. The choice of actions across IPA categories at each time step is computed according to the dynamic principles of affect control theory, which we review in more detail below as a possible starting point for a fruitful synthesis of AI and more traditional group-dynamics research.

3 INTEGRATING SOCIAL PSYCHOLOGICAL THEORY AND AI

The work reviewed in the previous section lies at two modeling extremes. Multi-agent systems approaches attempt to build highly complex models of agent behaviour based on strategic analyses that are theoretically elegant and individually sensible, but fail to capture both the simplicity and emergent complexity of human group behaviour. Classic agent-based models, on the other hand, have primarily been simple descriptive models that are able to explain aggregate statistics of emergent human behaviour data, but often fail to account for individual interactions in specific settings. Our recent work bridges the gap between these two extremes, by proposing a dual-systems approach to artificial intelligence that combines rational reasoning with emotional motivations and attentional mechanisms. In this section, we discuss our work in the development of psychologically grounded computational agents, and show how they can be used to account for human behaviour in small groups. This work builds on affect control theory, an early attempt at a mathematically formalized general-purpose theory of social interaction, which is rooted in decade-long theorizing in psychology and sociology. Our research on affectively-motivated artificial intelligence is fundamentally different than models in MAS or ABM, in that it assumes the agent's reward is primarily extrinsic, but that attentional mechanisms based on affect control are used to focus on action choices that are aligned with the prevailing social order. The resulting solutions are therefore narrow (more focused on socially-aligned solutions) and deep, giving longer-term strategies of cooperation that are more predictive of human behaviour.

3.1 Affect Control Theory

Affect control theory (ACT) links social perception with identity, behaviour, and emotion in social interactions [15]. The theory draws on symbolic interactionism [26] as well as theories of psychological consistency [14] and cybernetic control [34], proposing that people rely on linguistic representations with culturally-shared meanings to efficiently orient themselves within social interactions and anticipate the behavioural and emotional responses of others. Their motivation to maintain the cultural meanings associated with their own identities and the identities of others directly governs their interpersonal behaviours and emotions.

ACT uses the cultural meanings associated with labels for identities, behaviour, and emotions to model how humans interpret and respond to social events. Based on classic work by Osgood and colleagues [32], three universal semantic dimensions measure cultural meanings for various concepts: 1) evaluation (good vs. bad), 2) potency (weak vs. strong), and 3) activity (calm vs. excited). Evaluation is associated with perceptions of warmth, likeability, and approachability. Potency is associated with perceptions of competence, dominance, and submission. Activity is associated with perceptions of social agency and action readiness [37]. Shared cultural knowledge, expressed on these dimensions (referred to collectively as EPA), describes and differentiates social concepts, with each concept possessing a specific pattern of affective meanings known as fundamental sentiments. Fundamental sentiments reflect how the members of a given culture view elements of the social world; they characterize how good, powerful, and active particular identities, behaviours, or emotions seem in general, outside of the context of social events. For example, we tend to see heroes as good, powerful, and active (2.6, 2.3, 2.1), senior citizens as good, powerless, and inactive (1.2, -0.0, -1.8), and dropouts as bad, powerless, and inactive (-1.7, -1.8, -1.5)[1].

Fundamental sentiments for identities, behaviours, and emotions shift when they appear together in the context of social events. For example, a hero seems much more good, powerful, and active when he rescues a child (3.8, 2.0, 1.7) than when he compromises with a villain (0.1, 0.9, -0.1). These event-contextualized EPA meanings, known as transient impressions, capture the group's interpretation of actors, behaviours, and other elements of the situation and help to predict their behavioural and emotional responses to unfolding events. Affect control theory postulates that we can derive a group member's likely behavioural and emotional responses to a given situation from their transient impressions of that situation because human beings seek mental consistency between cultural expectations and social action. In other words, people act in ways that maintain the affective meanings associated with the group's interpretation of the situation, and expect others to do the same.

When our expectations about the identities and behaviours involved in an event are violated, we experience *deflection*, a tension about the situation which signals that our experiences are out of alignment with cultural expectations. Affect control theorists calculate deflection as the sum of the squared Euclidean distances between transient impressions of the identities and behaviours emerging from the situation and fundamental sentiments for these event elements. Thus, the lower the deflection, the greater the alignment between cultural expectations and situational circumstances. Deflection is much lower, for example, when a hero rescues a child than when they compromise with a villain. People seek to minimize deflection by acting in ways that maintain the group's interpretation of the situation; this is known as the affect control principle. Social actions are planned and carried out to either maintain situational meanings or to bring them back into alignment with cultural expectations. Reinterpretation and re-labeling can also be used to reduce deflection, and emotional signaling can be used to enhance coordination and agreement between agents.

Affect control theory was extended to model small-group interactions by developing an agent-based simulation platform called GroupSimulator [16]. Like the classic ACT model of dyadic interactions on which it is based, GroupSimulator rests on the affect control principle, according to which agents strive to maintain the shared meanings of all identities involved in the interaction. The model capitalizes on the many strengths of ACT, such as its capacity to efficiently model the creative human interpretive process in a diversity of social situations, using a parsimonious dimensional structure to represent cultural meanings, and small set of inputs to characterize events. GroupSimulator also provides a mechanism by which group behaviours, coded into a set of basic categories as defined in Interaction Process Analysis [4], can be mapped into the ACT framework and used to drive a simulation. IPA categories classify acts directed from one member of a group to another into twelve different functional categories. These categories cluster together into two types of task-oriented behaviours (giving vs. soliciting information or guidance) and two types of expressive behaviours (positive vs. negative) aimed at socio-emotional regulation. IPA and similar categorical systems have been employed in numerous studies reviewed by [4] to pursue questions such as status emergence in groups, over-time phases in group dynamics, and the effectiveness of collective problem-solving. IPA categories are mapped to EPA space in [16], and used to validate GroupSimulator using mock jury deliberations previously recorded, transcribed and manually classified into IPA categories. The distribution of behaviours exhibited by the jurors in this study was successfully reproduced using GroupSimulator [16].

Nevertheless, the model is not without its shortcomings. Although GroupSimulator is able to reproduce the behaviours of task groups, many of the parameters associated with social sensemaking and turn-taking are external to the model rather than theoretically integrated components. Consequently, studies conducted with GroupSimulator are vulnerable to overfitting (creating an

[1] for historical reasons, EPA measurements are scaled to lie between -4.3 and +4.3

overly complex model to explain idiosyncrasies in the data). For instance, two new theoretical assumptions were introduced in constructing this model to address uncertainty in turn-taking: 1) the actor with the greatest deflection will act next when self-selection is possible; and 2) actors will choose to interact with the group member that will most effectively minimize their deflection. In addition, GroupSimulator features only one utility function, the minimization of deflection. In real-world small-group interactions, group members must balance identity maintenance with task-related priorities. The recent development of Bayesian affect control theory (*BayesAct*), provides a means to address many of these limitations.

3.2 Bayesian Generalization of ACT

The Bayesian generalization of ACT, called *BayesAct*, overcomes many of the limitations mentioned in the previous section [19, 38]. *BayesAct* adds three new elements to ACT, which can also be viewed as removing limiting assumptions of the theory.

(1) *BayesAct* models all sentiments as probability distributions, thereby accounting for population-level differences in affective meanings for identities and behaviour that are likely replicated in personal uncertainties in social perception. Sentiment distributions can also be multi-modal, meaning that different viewpoints and multiple simultaneous identities and emotions of social agents are accounted for.
(2) *BayesAct* includes a denotative state space that can represent other semantically meaningful elements of an interaction. Using this, utility can be defined beyond deflection to include other aspects of individual preference that are likely to affect agents' interpretations of and responses to events. *BayesAct* can therefore account for the tension involved in a social dilemma where individual and social gains are at odds with each other. *BayesAct* can also account for task-related aspects of group behaviour in addition to the regulation of social-emotional relations among agents.
(3) *BayesAct* allows for the simultaneous optimisation of all elements of an interaction, including identities, behaviours, turn-taking. The model thus overcomes the above-mentioned overfitting problems of the GroupSimulator.

With these additions, *BayesAct* is constructed as a suitable basis model for task-oriented group interactions. *BayesAct* uses a probabilistic and decision theoretic model of stochastic control that arises in operations research called a partially observable Markov decision process (POMDP) [1].

BayesAct has been extended to take into account notions of the self [18], paralleling recent work on the affect control theory of self [24]. Self-sentiments can be represented as distributions over the same affective space as identities and behaviour, and reflect persons' autobiographical memories about themselves as they really are. The affect control theory of self builds on the key insight of ACT, the affect control principle, showing that people are motivated to seek out situations that help them maintain their self-sentiments. The Bayesian affect control theory of self therefore includes a mechanism for selection of interactants into social situations (i.e., the alignment of self-sentiments with situational identity enactments), providing a theoretical justification for some of the seemingly ad hoc mechanisms used in GroupSimulator. A *BayesAct* version of GroupSimulator is under construction in order to allow such simulations to be crried out. Some early examples can be found in [18]. In the following section, we describe simulations done with the basic ACT GroupSimulator, leaving simulations with *BayesAct* for future work.

4 ILLUSTRATION: GROUP DYNAMICS IN VIRTUAL TEAMS

In this section, we describe an ongoing study [2] and some preliminary results from it, aimed at understanding group dynamics in online software-development teams with the help of artificial-intelligence tools capitalizing on ACT. First, we review the GitHub platform, which is widely used for virtual collaboration and makes the resulting digital data traces available to researchers. Second, we use GroupSimulator to investigate some simple collaborative dynamics on GitHub in simulation.

4.1 Online Collaborative Networks

GitHub is a social coding platform where software developers from around the world come together to collaborate on software projects of common interest. The site enables software developers to work on the same software project (and even the same file in a project) simultaneously, and to merge their contributions without overwriting one another. The history of their contributions is saved, and one can always revert to an older version. Discussion about contributions can focus on the actual function of the code, or on the alignment of the function with the overall approach taken on the project. From a social psychological point of view, this process of discussion and revision is a crucial part of creating a relational meaning for the group of developers, which may later become a strong motivating force behind the group collaboration. Below, we show how we can make use of the data generated in such discussions to make inferences about the group process with machine-learning techniques informed by social psychological theories.

4.2 Affective Dynamics on GitHub

Emotions and interaction processes play an important role in software collaborations. For example, both positive and negative emotions have been shown to affect task quality, productivity, creativity, group rapport, and job satisfaction [3, 6, 10]. While large-scale digital data traces for discussions of software projects are openly available on GitHub, sentiment and emotional analysis can be challenging as affective content is embedded in technical discussions and punctuated with segments of code. Previous attempts include: [29], who perform a feasibility study of emotions mining using *Parrott's framework* on *Apache* issue reports; [12], who use lexical sentiment analysis to study emotions expressed in commit comments of open source projects; and [33], who use a Natural Language Text Processing tool (Natural Language Toolkit) to conduct a sentiment analysis of security-related discussions on GitHub. While studies like these yield interesting results, they are primarily descriptive in nature and therefore contribute little to the kind of theoretical explanation of group dynamics usually of interest to social and behavioural scientists. We hope that our present theory-driven approach helps

[2] themis-cog.ca

to overcome such gaps between description and explanation, which are common in computational social science.

4.3 Simulating Interactions on GitHub

We focus on two example simulations as an illustration of how GroupSimulator works, and the types of insights it can provide: 1) peer interactions occurring in a group of developers, and 2) interactions consisting of a leader and two newcomers. This allows us not only to compare a non-hierarchical group to a hierarchical one, but also to address a common and important type of interaction in online communities, the integration of newcomers [25]. We use GroupSimulator to examine the behaviours that are produced in each group, as well as who is enacting these behaviours and who is the target. We also examine how experiences of deflection are distributed across members of the group.

In order to develop generative models of self-organized collaborations on GitHub, we first recruited a sample of 503 GitHub users and asked them to provide evaluation, potency, and activity ratings of 587 identities, behaviours, and other concepts. We oversampled with respect to both gender (50% female) and race (30% non-white). Participants ranged from eighteen to seventy-nine years of age, with most being in their thirties. While the majority of respondents had some college (17%) or a bachelor's (38%) or advanced degree (20%), others reported having some high school education (3%), a high school education (14%), or vocational training (5%). This study provides a basis for the further evaluation of group behaviour on GitHub and is used as a resource throughout our ongoing work.

Our simulation of a non-hierarchical group of developers consists of three good, powerful, and lively agents. The agents' identity sentiments are drawn randomly from a multivariate normal distribution, centered at 1.61, 1.91, and 1.76 in E, P, and A dimensions, respectively. In contrast, the leader and newcomer simulation consists of one very good, potent, and active identity (the group's leader), and two good but less potent and active identities (the two newcomers). Identity sentiments for the leader are drawn from a multivariate normal distribution centered at 2.67, 2.37, and 2.27, while the sentiments of the two newcomers are drawn from multivariate distributions centered at 1.78, .77, and .62. These values come from the survey described above.

GroupSimulator works by first selecting a group member to act. Each group member is examined in turn, and the one with the action that will lead to the greatest reduction in deflection is selected next. This group member is then simulated as taking that deflection minimizing action, and we refer to the acted-upon group member as the *object* of the interaction. The process repeats, but has a number of heuristic methods that allow for a more reasonable simulation. First, group members that have acted are more likely to act next (regardless of deflection reduction), and turn-taking is also more likely than switching to a new interactant. The behaviors from the simulations (E,P,A values) are converted to IPA categories according to the mapping defined in [16].

Figure 1 displays the behaviour distributions predicted by GroupSimulator for a group of developers and a leader and two newcomers in the left and right panels, respectively. The x-axis indicates the IPA categories to which the behaviours were assigned. IPA categories refer to four clusters of behaviours: positive expressive behaviours

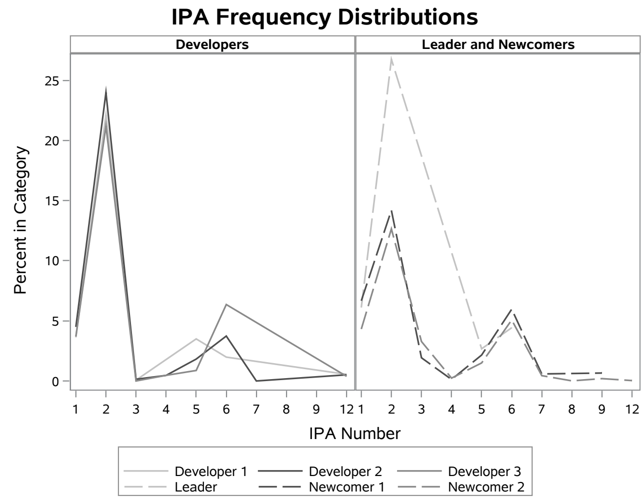

Figure 1: IPA Frequency Distributions for a Group of Three Developers and a Leader and Two Newcomers

(categories 1-3), behaviours associated with providing information or advice (categories 4-6), behaviours associated with soliciting information or advice (categories 7-9), and negative expressive behaviours (categories 10-12). The y-axis indicates the percentage of behaviours by IPA category each group member enacted, with the lines indicating the identities of each group member. The points indicate the frequency of behaviours in each category, with the absence of points indicating that the agent did not engage in a behaviour associated with that category. For example, developer 3 never agreed with other group members over the course of 1,500 turns across 220 simulations.

The simulation is able to capture the difference in the power dynamics implied by the identity labels. Although group members in both groups most frequently laughed or joked with others (category 2), the leader had many more opportunities than either developers or newcomers to engage in these behaviours. The leader also more frequently provided information than newcomers, and solicited for information or advice less often than newcomers. As expected, newcomers solicited for information and advice more often than either leaders or developers. Leaders also never engaged in negative expressive behaviours, most likely because the leader had numerous opportunities to affirm its identity. The developers also fell into a superior/subordinate pattern, with developer 1 most often giving advice and suggestions while developers 2 and 3 most often solicited information and advice. The higher frequency of antagonism (category 12) among the developers, however, suggests that this was not always a happy arrangement.

Figure 2 clarifies these dynamics by showing the proportion of actions directed at each group member, and at the group as a whole. The left and right panels display the interaction networks of the developers, and of the leader and newcomers respectively. The nodes correspond to each group member and the group, and are sized by the number of behaviours directed at them. The arrows indicate behaviours directed by one group member at another

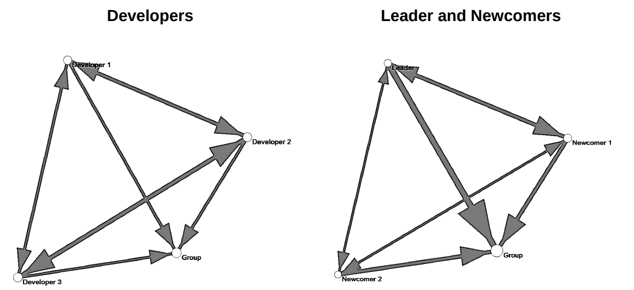

Figure 2: Interaction Networks of the Group of Three Developers and of the Leader and Two Newcomers.

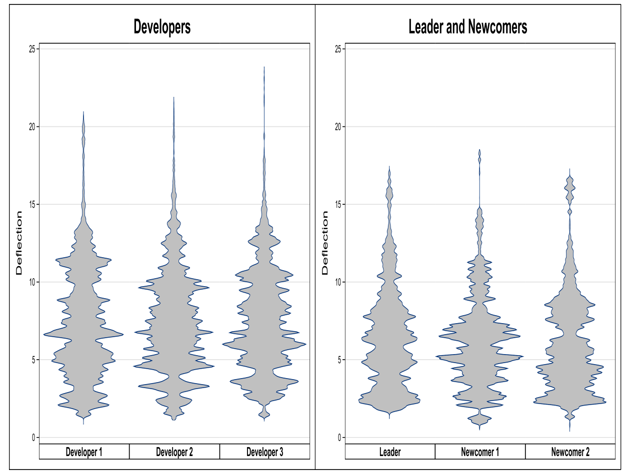

Figure 3: Deflection Experienced in a Group of Three Developers Compared to a Leader and Two Newcomers

member; the thickness of the arrows indicates the relative proportion of the total behaviours that occurred between each pair of actors. For example, the relatively equal weighting of the arrows directed by each developer towards the group indicates that each developer had essentially an equal number of opportunities to address the group, while the thicker arrows between developer 1 and developer 2 and between developer 2 and developer 3 indicate that interactions between these pairs of group members were more frequent than between the other group members. The roughly equal sizes of the nodes, however, indicate that each group member and the group were addressed by other group members at nearly the same frequency.

The high proportion of behaviours enacted by the leader compared to newcomers and developers in Figure 1 emerges from the interaction patterns in Figure 2. Leaders address the group, and newcomers tend solicit information and advice. The leader and newcomers tend to address the group more often than each other, with the exception of newcomer 1 and the leader. Nevertheless, the dominant interaction is the leader addressing the group, with far fewer pairwise interactions occurring than in the developers group.

Finally, Figure 3 displays the distribution of deflecting events experienced by each group member across the 220 runs, with the left and right panels referring to the developers and to the leader and newcomers respectively. The mean deflection and confidence interval of the developers is 7.41 (7.39-7.43). In contrast, the mean deflection and confidence interval of the leader and two newcomers is 6.39 (6.38-6.40) and 5.5 (5.41-5.62) respectively. The difference in the levels of deflection experienced by the developers compared to the leader and newcomers emerges from an interaction dynamic referred to by affect control theorists as the object diminishment effect. Being an object of an interaction results in a loss of perceived potency, and thus is a source of deflection for potent identities such as developers [39]. Increased deflection results in group members directing compensatory behaviours meant to restore their perceived loss of potency to other group members which in turn leads to greater deflection, resulting in the pattern of peer-to-peer interactions shown in Figure 2 and the antagonism shown in Figure 1. In contrast, the lower relative potency of newcomers compared to the leader allows them to endure the leader's jokes and accept direction, and to direct most interactions towards the group rather than towards each other. By directing actions towards the group, the loss of perceived potency is distributed across the group reducing the tendency towards compensatory behaviours and thus reducing the overall level of deflection experienced by the group.

5 FUTURE WORK

The long term aim of our research is to understand the social forces behind group interactions on social collaborative networks like GitHub. By using well-grounded social-psychological models, we aim to better understand how these forces shape group outputs, including implicit biases and relations, and how these outputs then translate into products that affect the wider social structure. We also aim to use predictions to help groups be more effective by focussing interpersonal interactions and group-level ties. These predictions could take the form of automated suggestions for structural changes, on-the-fly interaction recommendations (e.g. for how to communicate), or could take the form of new, automated group members ("bots") whose function is to serve as a social lubricant.

Our next steps are to automate the extraction of relational information from GitHub interactions. Initial efforts in this direction have shown the problem to be challenging, but not insurmountable [36]. Analysis of identity from GitHub user profiles and overall interactions, framed in terms of social interaction systems [4], will complement the survey described above, and will allow the creation of simulated groups (as presented here) that reflect real groups. Using *BayesAct*, we can then simulate group dynamics and compare statistics of interactions on dimensions of interaction processes, following the same model as in [16]. This will help validate the *BayesAct* models, and will allow us to make wider and longer term predictions about the social impacts of collaboration online. Further, we plan to use the data and simulations to learn the parameters of the *BayesAct* model, allowing it to adapt and change with the changing culture on the platform. Finally, we are also interested in moving beyond online (text-based) groups and to in-person groups with verbal and gestural communication.

6 CONCLUSION

This paper described our initial work towards integrating artificial intelligence models with classic social psychological approaches to the study of groups from Bales' Interaction Process Analysis to affect control theory. We illustrated here how we use sophisticated, psychologically grounded agent-based modeling to explain dynamical patterns of behaviour observed in such groups.

The use of artificial intelligence (AI) in group process research has, to date, been somewhat limited. The limitations have stemmed in part from the inherent complexity of modeling human behaviours in groups. Artificial intelligence has largely been concerned with building artificial agents based on a principle of rationality in the decision theoretic sense. These approaches fail to yield sufficiently rich or detailed models of human behaviour, especially within groups. In contrast, our work in building emotionally aligned AI is built upon a foundation of social-psychological theorizing about the role of emotion in group behaviour. Its fundamental tenet is that relational attachments between individuals and between individuals and groups define social orders that are strong, long lasting, and cooperative. These attachments form the basis for much human social interaction.

ACKNOWLEDGMENTS

This work was supported by Natural Sciences and Engineering Research Council of Canada (NSERC) and Social Sciences and Humanities Research Council of Canada (SSHRC), the Deutsche Forschungsgemeinschaft (DFG; Germany; Grant No. SCHR1282/3-1) and the National Science Foundation (NSF; United States, Grant Number 1723608) through the Trans-Atlantic Platform's Digging into Data Program.

REFERENCES

[1] K. J. Åström. 1965. Optimal Control of Markov Decision Processes with Incomplete State Estimation. *J. Math. Anal. App.* 10 (1965), 174–205.
[2] George A. Akerlof and Rachel E. Kranton. 2000. Economics and Identity. *The Quarterly Journal of Economics* 115, 3 (2000), 715–753.
[3] Scott Ambler. 2002. *Agile modeling: effective practices for extreme programming and the unified process.* John Wiley & Sons.
[4] Robert F Bales. 1999. *Social interaction systems: Theory and measurement.* New Brunswick, NJ: Transaction Publishers.
[5] Valerie Capraro and David G. Rand. 2017. Do the Right Thing: Preferences for Moral Behavior, Rather Than Equity or Efficiency per se, Drive Human Prosociality. (2017). Forthcoming in Judgment and Decision Making. Available at SSRN.
[6] Munmun De Choudhury and Scott Counts. 2013. Understanding Affect in the Workplace via Social Media. In *Proceedings of the 2013 Conference on Computer Supported Cooperative Work (CSCW '13).* ACM, New York, NY, USA, 303–316.
[7] Guillaume Deffuant, David Neau, Frederic Amblard, and Gérard Weisbuch. 2000. Mixing beliefs among interacting agents. *Advances in Complex Systems* 3, 01n04 (2000), 87–98.
[8] Bruce Edmonds and Ruth Meyer. 2017. *Simulating social complexity: A handbook.* Springer.
[9] Ernst Fehr and Klaus M. Schmidt. 1999. A theory of fairness, competition, and cooperation. *The Quarterly Journal of Economics* 114, 3 (1999), 817–868.
[10] BL Fredickson. 2001. The role of positive emotions in positive psychology. *American psychologist* 56, 3 (2001), 218–226.
[11] David R Gibson. 2005. Taking turns and talking ties: Networks and conversational interaction. *American journal of sociology* 110, 6 (2005), 1561–1597.
[12] Emitza Guzman, David Azócar, and Yang Li. 2014. Sentiment Analysis of Commit Comments in GitHub: An Empirical Study. In *Proceedings of the 11th Working Conference on Mining Software Repositories (MSR 2014).* ACM, New York, NY, USA, 352–355.
[13] Jonathan Haidt. 2001. The emotional dog and its rational tail: A social intuitionist approach to moral judgment. *Psychological Review* 108 (2001), 814–834.
[14] Fritz Heider. 1946. Attitudes and cognitive organization. *The Journal of psychology* 21, 1 (1946), 107–112.
[15] David R Heise. 2007. *Expressive order: Confirming sentiments in social actions.* Springer.
[16] David R Heise. 2013. Modeling interactions in small groups. *Social Psychology Quarterly* 76, 1 (2013), 52–72.
[17] Jesse Hoey, Tobias Schröde, Jonathan H. Morgan, Kimberly B. Rogers, Deepak Rishi, and Meiyappan Naggappan. 2018. Artificial Intelligence and Social Simulation: Studying Group Dynamics at Massive Scale. *Small Group Research* (2018). forthcoming.
[18] Jesse Hoey and Tobias Schröder. 2015. Bayesian Affect Control Theory of Self. In *Proceedings of the AAAI Conference on Artificial Intelligence.* 529–536.
[19] Jesse Hoey, Tobias Schröder, and Areej Alhothali. 2016. Affect Control Processes: Intelligent Affective Interaction using a Partially Observable Markov Decision Process. *Artificial Intelligence* 230 (2016), 134–172.
[20] Joshua D.A. Jung and Jesse Hoey. 2016. Grounding Social Interaction with Affective Intelligence. In *Proceedings of the Canadian Conference on AI.* Victoria, BC.
[21] Peter Kollock. 1998. Social dilemmas: The anatomy of cooperation. *Annual review of sociology* 24, 1 (1998), 183–214.
[22] Edward J. Lawler, Shane R. Thye, and Jeongkoo Yoon. 2009. *Social Commitments in a Depersonalized World.* Russell Sage Foundation.
[23] David Lazer, Alex Sandy Pentland, Lada Adamic, Sinan Aral, Albert Laszlo Barabasi, Devon Brewer, Nicholas Christakis, Noshir Contractor, James Fowler, Myron Gutmann, et al. 2009. Life in the network: the coming age of computational social science. *Science (New York, NY)* 323, 5915 (2009), 721–723.
[24] Neil J MacKinnon and David R Heise. 2010. *Self, identity and social institutions.* Palgrave and Macmillan, New York, NY.
[25] Jennifer Marlow, Laura Dabbish, and Jim Herbsleb. 2013. Impression formation in online peer production: activity traces and personal profiles in github. In *Proceedings of the 2013 conference on Computer supported cooperative work.* ACM, 117–128.
[26] George Herbert Mead. 1934. *Mind, self and society.* Vol. 111. Chicago University of Chicago Press.
[27] David M. Messick and Charles G. McClintock. 1968. Motivational Bases of Choice in Experimental Games. *Journal of Experimental Social Psychology* 4 (1968), 1–25.
[28] Jonathan H Morgan, Geoffrey P Morgan, and Frank E Ritter. 2010. A preliminary model of participation for small groups. *Computational and Mathematical Organization Theory* 16, 3 (2010), 246–270.
[29] Alessandro Murgia, Parastou Tourani, Bram Adams, and Marco Ortu. 2014. Do Developers Feel Emotions? An Exploratory Analysis of Emotions in Software Artifacts. In *Proceedings of the 11th Working Conference on Mining Software Repositories (MSR 2014).* ACM, New York, NY, USA, 262–271.
[30] M.A. Nowak. 2006. Five rules for the evolution of cooperation. *Science* 314 (2006), 1560–1563.
[31] Mark G Orr, Christian Lebiere, Andrea Stocco, Peter Pirolli, Bianica Pires, and William G Kennedy. 2018. Multi-scale Resolution of Cognitive Architectures: A Paradigm for Simulating Minds and Society. In *International Conference on Social Computing, Behavioral-Cultural Modeling and Prediction and Behavior Representation in Modeling and Simulation.* Springer, 3–15.
[32] Charles E Osgood, William H May, and Murray S Miron. 1975. *Cross-cultural universals of affective meaning.* University of Illinois Press.
[33] Daniel Pletea, Bogdan Vasilescu, and Alexander Serebrenik. 2014. Security and emotion: sentiment analysis of security discussions on GitHub. In *Proceedings of the 11th working conference on mining software repositories.* ACM, 348–351.
[34] William T Powers. 1973. *Behavior: The control of perception.* Aldine.
[35] Matthew Rabin. 1993. A theory of fairness, competition and cooperation. *The American Economic Review* 83, 5 (1993), 1281–1302.
[36] Deepak Rishi, Jesse Hoey, Meiyappan Naggappan, Kimberly B. Rogers, and Tobias Schröder. 2018. Emotion and Interaction Processes in a Collaborative Online Network. In *International Conference on Computational Social Science.* Chicago,IL.
[37] Kimberly B Rogers, Tobias Schröder, and Wolfgang Scholl. 2013. The affective structure of stereotype content: Behavior and emotion in intergroup context. *Social Psychology Quarterly* 76, 2 (2013), 125–150.
[38] Tobias Schröder, Jesse Hoey, and Kimberly B. Rogers. 2016. Modeling dynamic identities and uncertainty in social interaction: Bayesian affect control theory. *American Sociological Review* 81 (2016), 828–855.
[39] Lynn Smith-Lovin and David R Heise. 1988. Affect control theory: Research advances. *New York: Gordon and Breach* (1988).
[40] Flaminio Squazzoni. 2012. *Agent-based computational sociology.* John Wiley & Sons.
[41] Changkun Zhao, Ryan Kaulakis, Jonathan H Morgan, Jeremiah W Hiam, Frank E Ritter, Joesph Sanford, and Geoffrey P Morgan. 2015. Building social networks out of cognitive blocks: factors of interest in agent-based socio-cognitive simulations. *Computational and Mathematical Organization Theory* 21, 2 (2015), 115–149.