SOE 10: Focus Session: Large Language Models, Social Dynamics, and Assessment of Complex Systems

Recent Large Language Models (LLMs) learn the statistics of human communication, and hence they present a promising way of simulating societal interactions and opinion formation. Moreover, their flexible architecture makes them versatile to also learn the statistics of other complex systems. - We will discuss how to make best use of these advancements to model social interactions and dynamics across different scales and settings, how to use them to assess sentiments, narratives and opinions, and also how one can use LLMs to also understand other complex systems. Our aim is to offer an evaluative perspective on both the potential and challenges that are intrinsic to this rapidly evolving research area.

Organized by Vincent Brockers, Johannes Zierenberg, and Viola Priesemann

Time: Thursday 15:00–18:15 Location: H45

Invited Talk SOE 10.1 Thu 15:00 H45

Emergent Behaviors in LLMs-Populated Societies — •Giordano De Marzo[1,2,3], Claudio Castellano[4,2], Luciano Pietronero[2], and David Garcia[1,3] — [1]Konstanz University — [2]CREF Enrico Fermi Research Center — [3]CSH Vienna — [4]CNR-ISC Institute for Complex Systems

Applications of Large Language Models (LLMs) increasingly involve collaborative tasks where multiple agents interact, forming "LLM societies." In this context, we explore whether large groups of LLMs exhibit emergent group behaviors similar to those in human societies. First, by simulating social network formation, we observe that LLMs spontaneously form scale-free networks. Agents connect through linear preferential attachment, mirroring the Barabasi-Albert model and real-world social networks. Second, we investigate the ability of LLMs to reach consensus on arbitrary norms without external preferences, thereby self-regulating their behavior. In human societies, consensus without institutions is limited by cognitive capacities. Similarly, we find that LLMs can reach consensus, with the opinion dynamics described by a majority force coefficient that determines the likelihood of consensus. This majority force strengthens with higher language understanding but decreases with larger group sizes, resulting in a critical group size beyond which consensus becomes unattainable. For more advanced LLMs, this critical size grows exponentially with language capabilities, exceeding the typical size of informal human groups.

SOE 10.2 Thu 15:30 H45

A Framework for Multi-Step Discussions of LLM-Based Agents — •Vincent Brockers[1,2], David Ehrlich[2,3], and Viola Priesemann[1,2] — [1]Max Planck Institute for Dynamics and Self-Organization — [2]University of Göttingen — [3]Göttingen Campus Institute for Dynamics of Biological Networks

Recent studies have demonstrated the utility of Large Language Models (LLMs) across various settings and domains, particularly in their ability to approximate human behavior for simulating and studying societal phenomena. A critical component of this capability is their proficiency in generating realistic dialogues. We are interested in the collective dynamics that arise from their dialogues, in particular in the context of understanding collective opinion dynamics.

We developed a framework for simulating multi-step discussions among LLM-based agents, which is crucial for these studies. We conduct an in-depth analysis of the components of such simulations, highlighting the challenges and specifically investigating potential biases.

We find that these opinion dynamics are primarily influenced by model-inherent topic biases, which exhibit attractors in the opinion space. Additionally, the framing of opinion evaluation plays a crucial role in analyzing these dynamics. Our findings offer valuable insights for future research aimed at optimizing the design of LLM agent-based simulations and enhancing our modeling capabilities of complex systems, such as human networks.

SOE 10.3 Thu 15:45 H45

Collective Turing Tests on LLMs — •Azza Bouleimen[1,2], Giordano de Marzo[3], Taehee Kim[3], Silvia Giordano[2], and David Garcia[3] — [1]University of Zurich, Zurich, Switzerland — [2]SUPSI, Lugano, Switzerland — [3]University of Konstanz, Konstanz, Germany

In this project, we investigate whether social media conversations generated by independent LLMs are indistinguishable from those of humans, i.e., whether LLMs used to generate social media content can pass the Turing test. Formally, we conduct an experiment in which we prepare a series of English Reddit submissions to which we attach two conversations. One of the conversations is an authentic human conversation ,while the other is generated artificially using an LLM. We generate conversations using GPT-4o and llama 3 70B. We vary the temperatures of the models used and the length of the conversations. We recruit participants from Prolific. These subjects are asked to select the conversation they believe is generated by AI. Our preliminary results suggest that, overall, participants are fooled by LLM 40% of the time. Llama 3 70B conversations appear to fool users more often than GPT-4o ones. Through this study, we investigate to what extent and with which configurations LLM could be best used to simulate user conversations on social media. To the best of our knowledge, this is the first attempt to evaluate the performance of LLMs in mimicking a social media conversation between a group of individuals.

SOE 10.4 Thu 16:00 H45

Dialogue React: Enhancing Conversation Synthesis for Social Science Simulations — •Ruggero Marino Lazzaroni — RWTH Aachen

Recent advancements in Large Language Models (LLMs) have spurred significant progress in conversation synthesis and social dynamics modeling. However, existing approaches often fall short in generating human-like dialogues suitable for complex social science simulations and opinion formation studies. This research introduces DialogueReact, a novel framework that builds upon previous work in conversation synthesis by incorporating react prompting, dialogue acts and agentic behaviour across different scales. By leveraging LLMs' ability to learn communication statistics, DialogueReact generates qualitatively better conversations compared to previous state-of-the-art methods. Our findings demonstrate the potential of DialogueReact to improve social science research by providing an extra tool for simulating complex social interactions and understanding collective behavior in social systems.

SOE 10.5 Thu 16:15 H45

Computational modeling of LLM powered personalized recommendations — •Alessandro Bellina[1,2,4], Giordano De Marzo[1,3,4], David Garcia[3], and Vittorio Loreto[1,2,4] — [1]Centro Ricerche Enrico Fermi, Piazza del Viminale, 1, I-00184 Rome, Italy — [2]Sony Computer Science Laboratories Rome, Joint Initiative CREF-SONY, Piazza del Viminale, 1, 00184, Rome, Italy — [3]University of Konstanz, Universitaetstrasse 10, 78457 Konstanz, Germany — [4]Dipartimento di Fisica Università La Sapienza, P.le A. Moro, 2, I-00185 Rome, Italy

Large language models (LLMs) are transforming recommendation systems by tailoring content to user preferences, but they also risk reinforcing filter bubbles and driving polarization. This study investigates the dual impact of LLM-based recommendations, analyzing their inherent biases and exploring how prompt engineering can mitigate these effects. By combining synthetic simulations with real-world Twitter data, we assess how LLMs influence user behavior and the extent to which recommendations amplify or reduce polarization. Preliminary results suggest that prompt engineering enables greater control over recommendations, fostering diversity and creativity. This research provides insights into the risks and opportunities of LLM-powered systems, offering a framework for designing more inclusive and balanced recommendation algorithms.

SOE 10.6 Thu 16:30 H45

Statistical Modelling of Physics Classroom Interactions — •Niklas Stausberg, Karina Avila, Steffen Steinert, Jochen

Kuhn, and Stefan Küchemann — Lehrstuhl für Didaktik der Physik, Ludwig-Maximilians-Universität München, Geschwister-Scholl-Platz 1, 80539 München

Teacher-student and student-student dialogues are one of the central means to mediate learning in classroom settings, as they foster active participation, promote cognitive engagement, scaffold conceptual understanding, and enhance metacognitive skills. However, strategies for effective classroom dialogues have primarily been studied at an organizational level, and the nuanced mechanisms through which real-time teacher-student and student-student interactions influence individual learning outcomes within the classroom remain largely unexplored.

New advances in analysis of verbal interactions now offer opportunities to perform granular analyses of classroom dialogues. Using a combination of response times and speed, voice emotion recognition, and self-regulation state detection enables clustering of patterns in students' cognitive, emotional, and self-regulatory states and map them to specific types of teacher and student responses to gain insights into the dynamics of effective classroom interactions. In this talk, we discuss how statistical modeling of classroom interactions, e.g. via percolation theory, can be used to provide additional perspectives in the understanding of social learning processes.

15 min. break

SOE 10.7 Thu 17:00 H45

Disentangling individual vs. collective contributions to polarization in political voting — Gavin Rees and •Edward Lee — Complexity Science Hub, Vienna, Austria

Politics around the world exhibit increasing polarization, most visibly demonstrated by rigid voting configurations in legislatures. This impacts their ability to seek effective compromise and pass impactful legislation. The crux of polarization is the emergence of a unidimensional ideological axis that primarily determines voting. Yet, legislative bills often negotiate multiple issues, whose effects may add up or compete for any individual vote. We develop a model inspired from statistical physics that accounts for voters with multi-dimensional preferences that can be linearly combined to generate coalitions. This allows us to capture higher-order correlations beyond those captured in direct spin-spin interaction models. As a result, the model is dramatically simpler, fits the data better, and is more interpretable than canonical maximum entropy or utility-based models from political science for large voting bodies. We study roll-call voting in US Congress, and we find coexistence of non-polarized with polarized modes. We show how increased partisan polarization is the contribution of two terms: increased party-line voting by the individual plus depletion of votes that would elicit complex, bipartisan response. Both decrease over time, but the latter falls much faster. Thus, we propose a more principled way of tracking the emergence of polarization at the level of the voter and the collective issues the many vote on.

SOE 10.8 Thu 17:15 H45

Understanding Information Spread in Social Networks through the Lense of Self-Organized Criticality: A Study on Telegram — •Roman Ventzke[1,2], Anastasia Golovin[1,2], Sebastian Bernd Mohr[1,2], Andreas Schneider[1,2], and Viola Priesemann[1,2] — [1]Max-Planck-Institut für Dynamik und Selbstorganisation, Göttingen — [2]Georg-August-Universität Göttingen

To effectively address the proliferation of misinformation in social media, one fundamentally needs to understand how information spreads in online social networks. We investigate the dynamics of information spread using a large dataset from the messenger platform Telegram, showing that information spread in the networks happens in bursty avalanches with scale-free statistics that resembles critical behavior of physical systems.

We find that the critcal exponents of the system can be well described by a mean-field Random Field Ising Model (RFIM), alluding to an important role of complex contagion and peer-pressure effects in the propagation of information. By coarse-graining dynamics in the topic space we show additional evidence that the process indeed belongs to the RFIM class. We demonstrate through simulations that the spreading process on the network can be well-described by mean-field models and discuss how self-regulation of the network gives rise to criticality.

SOE 10.9 Thu 17:30 H45

Dynamics of information spread: the phase space of topics on Telegram — •Anastasia Golovin[1,2], Simon Bauer[1], Sebastian Mohr[1], and Viola Priesemann[1,2] — [1]Max Planck for Dynamics and Self-Organisation, Göttingen, Germany — [2]Institute for the Dynamics of Complex Systems, University of Göttingen, Germany

The information that is spread on digital platforms shapes public discourse and influences decision-making. In this talk, we explore the dynamics of topics discussed on Telegram, a privacy-oriented messenger that became a hub for misinformation spread during the COVID-19 pandemic. Using transformer-based topic modeling techniques, we map each Telegram channel onto a trajectory in a high-dimensional phase space of topics. Treating these trajectories as a stochastic dynamical system, we identify topic areas that act as attractors, pulling trajectories from their neighborhood, and others that emerge only transiently in response to global events. This approach bridges computational social science and dynamical systems theory, offering a novel perspective on understanding discourse dynamics in communication networks.

SOE 10.10 Thu 17:45 H45

Computational approaches to analyzing political narratives — •Armin Pournaki[1,2,3] and Tom Willaert[4] — [1]Max Planck Institute for Mathematics in the Sciences, Leipzig, Germany — [2]Laboratoire Lattice, École Normale Supérieure - PSL - CNRS, Paris, France — [3]médialab, Sciences Po, Paris, France — [4]Brussels School of Governance, Vrije Universiteit Brussel, Brussels, Belgium

Narratives are key interpretative devices by which humans make sense of political reality. As the significance of narratives for understanding current societal issues such as polarization and misinformation becomes increasingly evident, there is a growing demand for methods that support their empirical analysis. To this end, we propose a graph-based formalism and machine-guided method for extracting, representing, and analyzing selected narrative signals from digital textual corpora, based on Abstract Meaning Representation (AMR) [1]. The formalism and method introduced here specifically cater to the study of political narratives that figure in digital trace data. We demonstrate the method on corpora gathered from archived political speeches and social media posts. [1] https://arxiv.org/abs/2411.00702

SOE 10.11 Thu 18:00 H45

Optimizing pandemic mitigation in the presence of seasonality, emerging variants and vaccination — •Laura Müller[1,2], Fabio Sartori[1,3], Jonas Dehning[1,2], Maximilian F. Eggl[4], and Viola Priesemann[1,2] — [1]Max-Planck-Institute for Dynamics and Self-Organization, Göttingen, Germany — [2]Institute for the Dynamics of Complex Systems, University of Göttingen, Göttingen, Germany — [3]Karlsruhe Institute of Technology, Karlsruhe, Germany — [4]Institute of Neuroscience, CSIC-UMH, Alicante, Spain

Mitigation efforts are often necessary to control the spread of infectious diseases during pandemics and epidemics. However, these measures come with societal and economic costs. We present a general optimization framework based on control theory to balance the trade-off between mitigation costs and infection rates, aiming to identify the optimal time-dependent mitigation strategies across a range of scenarios: constant reproduction number, seasonal changes, emerging variants, vaccination rollout, and delayed mitigation onset. Our results show that timely reactions and proactive, preventative measures are crucial in pandemic management to save both lives and costs.