Northeastern University London

Information Sharing: Strategies

A PolyGraphs Public Engagement Workshop

As part of the PolyGraphs research project,[1] Northeastern University London hosted the second of two public engagement workshops concerned with *Information Sharing* on 30 June 2023.

The workshop brought together the PolyGraphs research team, including colleagues at the Center for Design, with panellists and participants from a range of stakeholder groups in the information sharing ecosystem: news media organizations; oversight groups; and the general public.

Our panel consisted of: Francesca Panetta (AKO Storytelling Institute, UAL); Kayleen Devlin (BBC Verify); and Omri Preiss (Alliance for Europe & the DISARM Foundation).

Here we present a brief (executive) summary of the research that was discussed at the workshop.

Executive Summary

The informational environment has changed dramatically since the advent of the internet at the end of the last century, giving rise to various concerns about attitudes and opinions within contemporary societies, and the behaviours they may bring about.

The PolyGraphs project investigates the influences on public opinion of social network *structures* and information consumption *strategies*: and whereas the first workshop was concerned primarily with the former, the second was concerned with the latter.

Our research team uses computer simulations of communities of inquiring agents, who learn from their own observations, and from the testimony of their network neighbours.

We make a number of idealizations: we assume the opinions concern factual matters, and so are either *correct* (true) or *incorrect* (false); our agents are *rational*, in the sense that they are appropriately responsive to evidence; and the *evidence* is stochastic, or chancy – it can be thought of as the results of coin tosses, to determine whether there is a bias towards heads or tails.

Others have found that, even in an environment comprising only accurate information, distrust of those with divergent opinions can lead to polarization (O'Connor & Weatherall, *The Misinformation Age*). We showed that, in their (homophily-based) models, *more trust led to more knowledge*.

[1] PolyGraphs is supported by the British Academy, the Royal Society, the Royal Academy of Engineering, and the Leverhulme Trust under the APEX scheme.

In our own (higher-order evidence) models, we simulate the effects of introducing mis- and disinformation into the environment when agents pursue various information processing strategies.

We distinguish mere *misinformants* from *disinformants*. The former provide 'evidence' that is neutral overall with respect to the underlying question, which may therefore be thought of as 'noise'; whereas the latter present testimony that is biased away from the truth.[2]

We also distinguish a trusting, or '*gullible*', strategy for processing the information available, from a more sceptical strategy in which the level of trust is '*aligned*' with the level of reliability of the informants in the networked community.

We find that, whether agents pursue the gullible or aligned strategies, the more misinformants are present in the network, the less likely it is that a correct consensus will emerge in the community, and when it does, it takes longer to arrive at this opinion (i.e. the truth).

We also find that, for a given level of misinformation, the aligned strategy is more likely to achieve a correct consensus than the gullible one, but it takes longer to arrive at that consensus (when there is a significant difference in the number of simulation steps required). In short, when we compare the two strategies, there is a trade-off between *accuracy* and *efficiency*.

In the presence of disinformation, the ability of gullible agents to discern the truth plummets, collapsing almost entirely when levels of disinformation are high. Agents pursuing the aligned strategy do better in this regard, but are nevertheless significantly delayed in arriving at the truth.

In simulations involving large, real-world networks, we find that gullible agents' average degree of belief in the correct opinion reduces in the presence of misinformation; and, intringuingly, that a very small number influential misinformants can have an effect that is comparable to that of a much larger number of misinformants chosen (uniformly) at random.

In our ongoing research, we explore 'confessionals' models in which neighbours share information about their own beliefs, rather than about their observations. This poses different (though related) challenges of information processing strategies.

The computational framework we employ has been built to be efficient, customizable, and suitable for integration with machine learning. It is available open source on GitHub.

Our data can also be experienced – see here for a guide to interpreting our visualizations.

The panel discussion of our research findings, and of the current political, legislative, and media landscapes to which they are applicable, revealed fruitful avenues for further investigation, and we received a great deal of valuable input from our workshop participants.

We look forward to collaborating in our ongoing research with external partners to address the societal challenges posed by mis- and disinformation.

[2] At the workshop we distinguished misinformants who draw their reported observations from a *binomial* vs a *uniform* distribution. The findings reported concern the former; data from the latter requires further analysis.