

Conversational Launch Pads

Emma M. Templeton (emma.m.templeton@dartmouth.edu)

Department of Psychological and Brain Sciences, Dartmouth
Hanover, NH 03755 USA

Luke J. Chang (luke.j.chang@dartmouth.edu)

Department of Psychological and Brain Sciences, Dartmouth
Hanover, NH 03755 USA

Thalia Wheatley (thalia.p.wheatley@dartmouth.edu)

Department of Psychological and Brain Sciences, Dartmouth
Hanover, NH 03755 USA
Santa Fe Institute
Santa Fe, NM 87501 USA

Abstract

How do people start conversations with someone they have never met before? In this project, we investigate the hypothesis that good starting topics facilitate transitions to many different topics. To test this, we leverage a dataset of unstructured, 10-minute conversations between pairs of strangers. Using natural language processing (NLP) and network approaches, we show that strangers begin their conversations with topics that are centrally located in a network of topic transitions. These “launch pad” topics are useful starting points because they are well-connected to other topics, potentially increasing the likelihood of finding common ground. These findings underscore the fact that it is not the *semantic meaning* of a topic that makes it an effective starting point, but rather its *transition properties*. This insight paves the way for future research to identify conversational launch pads in different populations, where common starting topics may differ widely but nonetheless hold similar network positions. When people start conversations, they begin the process of trying to understand and connect with another person’s mind. Here, we examine how this important process unfolds.

Keywords: conversation; topics; transitions; NLP; network analyses

Introduction

Imagine it just started snowing. You meet a stranger during your commute and mention the weather. They reply with excitement and tell you about a new cross-country trail that opened nearby, leading to a conversation about winter sports. Later, you mention the weather again, this time to a stranger in your coffee queue. They reply with concern about driving home without snow tires, leading to a conversation about local car mechanics. Both conversations began with the same topic (the weather), but quickly went in distinct directions. As this example illustrates, certain topics may be particularly well-suited to transition to many other topics. In this project, we examined whether strangers tend to use these kinds of topics to launch their conversations.

Prior work provides some insight into how people start their conversations and what types of things they talk about. For example, work in Conversation Analysis examines how

people open their conversations (Pillet-Shore, 2018), with a focus on how people greet each other (Pillet-Shore, 2012) and make introductions (Pillet-Shore, 2011). These openings describe how people *initiate* a conversation, a process that tends to happen before a topic is selected. Other work specifically examines the *topics* that people use throughout their conversations, including how topics differ by gender (Bischoping, 1993; Dunbar, Marriott, & Duncan, 1997), group size (Cooney et al., 2020), and type of conversation partner (Bearman & Parigi, 2004). However, these studies largely ignore how people transition from one topic to another (Maynard, 1986; Yang 2019). How do people decide which topic to talk about first? And how does that choice impact where the conversation goes next? Examining such transition dynamics has proven useful in other domains. For example, quantifying the typical transitions between different emotions (Thornton & Tamir, 2017), mental states (Thornton et al., 2023), and actions (Thornton & Tamir, 2021) elucidates how people make predictions about other’s behavior. Similarly, if we know what topic is being used in a conversation, it may be possible to predict which topic (or set of topics) is likely to follow. Each topic choice sets a conversation on a new trajectory, influenced by the topic choices that came before it. These topic dynamics may have implications for how easily people form connections or find common ground.

Techniques from network analysis can be used to model topic transitions in conversation. Network analysis is a powerful tool for describing and understanding complex systems. By representing distinct elements as nodes and the connections between those elements as edges in a network, it is possible to (i) visualize the structure of the system at once and (ii) identify central nodes that may have an outsized influence on that system. For example, social network analysis—which represents people as nodes and relationships between them as edges—has shown that phenomena like obesity (Christakis & Fowler, 2007), cooperation (Fowler & Christakis, 2010), happiness (Fowler & Christakis, 2008), and ideas (Singh, 2005) spread from person to person. Neuroimaging studies have found that people encode and

spontaneously represent the social networks in which they are embedded (Parkinson, Kleinbaum, & Wheatley, 2017; Schwyck et al., 2023). Network analyses have also been used to better understand memory for events in a narrative. A recent paper created networks where nodes represented events in a story and the edges between them represented their semantic similarity (Lee & Chen, 2022). They found that people were more likely to remember events that were centrally located in these semantic networks. As these examples highlight, network analyses can be used to efficiently represent complex systems and specific features of the network structure can be used to characterize processes relevant for cognition.

Here, we employ network analyses to characterize how strangers tend to transition from topic to topic over the course of their conversations. We represent this transition matrix as a weighted, directed network with individual topics as nodes. We find that topics that are centrally located in this network tend to get used early in the conversation and decrease in prevalence as conversations persist.

We introduce the term “conversational launch pad” to define topics that have the tendency to branch into many different topics, just as the weather did in our opening example. It may be no accident that people choose these particular topics to start their conversations. Their transitional properties are well suited to allow conversation partners to find their own path to more interesting places, perhaps increasing the likelihood of building rapport and finding common ground (Cassell et al., 2007; Jucker & Smith, 2022; Tickle-Degen & Rosenthal, 1990).

Methods

Dataset

We used a previously collected dataset of conversations between pairs of undergraduate students (Templeton et al., 2022). In each conversation session, two participants were video recorded as they had a 10-minute conversation. Participants could talk about whatever they wanted and were not provided with any conversation prompts to start their conversation. This allowed us to examine how people naturally start their conversations in the absence of any external instruction.

This dataset leveraged a round-robin design, with every round consisting of 11 same-gender participants. All participants were scheduled to complete 10 conversation sessions, one with each member of the round-robin. Participants never had more than three conversations in a single day. The dataset included six round-robin groups, with 66 participants (33 female).

The video recordings of each conversation were transcribed by an external transcription company. Each transcript contained a start and end timestamp for each speech turn as well as the text of what was said.

Because the focus of this project is on how *strangers* start their conversations, we excluded dyads where both dyad members reported a response greater than 0 to the question,

“How well did you know your study partner before today?” (0 = Not well at all, 50 = Moderately well, and 100 = Extremely well) in a survey following their conversation. The analyses reported in this paper comes from 261 stranger dyads (123 female dyads, 138 male dyads).

Defining Topics

Binning text Although all conversations in this dataset were exactly 10-minutes long, the unstructured nature of the task meant that conversations moved at different paces with different turn-taking dynamics. To facilitate comparisons across conversations, conversation transcripts were binned into 30-second increments (20 bins per conversation). Each bin contained the text of the speech turns occurring in that 30-second window, with a mean word count of 98.81 ($SD = 33.90$). For example, bin 1 contained the text of all the turns that occurred in the first 30 seconds, bin 2 contained the text of all the turns occurred in the second 30 seconds, and so on. This approach ignores speaker identity; if both participants spoke in a 30-second window, both of their turns would be included in that bin. For the purposes of this project, the *conversation* is the unit of analysis, not individual speakers.

Transforming text into language embeddings We next used the pretrained Universal Sentence Encoder (Cer et al., 2018) to embed the semantic meaning of the text in each bin as a 512-dimensional numeric representation. The Universal Sentence Encoder uses a dual encoder framework that combines transformer and deep averaging network architectures. As a result, text is represented as a single point in a high-dimensional semantic space, where points that are closer together are more semantically similar.

Clustering embeddings to reveal topics Language models like the Universal Sentence Encoder are well-suited to describe semantic similarity between text inputs, however the meaning of individual features in the embedding space are not interpretable on their own. Our next step was to cluster the embeddings to reveal interpretable topics describing what strangers in this dataset talked about during their conversations.

We first reduced the dimensionality of the embedding space using Uniform Manifold Approximation and Projection (UMAP), a dimensionality reduction technique that aims to preserve distances between observations. This step helps mitigate the “curse of dimensionality” and potential multicollinearity by using a non-linear compression approach (Assent, 2012). We used the UMAP package in python (McInnes, Healy, & Melville, 2018) with the parameters $n_neighbors=15$, $min_dist=0.1$, $metric='cosine'$. UMAP requires that we specify the number of dimensions of the reduced feature space ($n_components$). To help make this decision, we first examined the pattern of pairwise cosine similarity between the embeddings in the original feature space. Our goal was to choose a reduced feature space that preserved the between-dyad similarity structure. We inspected how this pattern changed with different numbers of




Figure 1: Proportion of dyads in each topic for each 30-second time bin. The proportions are annotated in each cell. Colormap is capped at 0.3 to visually preserve differences between cells by reducing the influence of the bin 1 ‘introductions’ topic.

UMAP components (e.g., 200, 100, 50, 10, 2) and observed that the pattern of pairwise cosine similarity values was quite consistent across many of these versions (e.g., the similarity structure using a UMAP with 200 components looked quite like the similarity structure using a UMAP with 10 components). However, when the number of UMAP components dropped significantly (i.e., to 2) the pattern of similarity values became much coarser. We opted to use a UMAP solution with 10 components to take advantage of this more granular representation, without including too many components that might adversely impact the clustering algorithm.

We then used k-means clustering to divide this reduced space into distinct topic clusters. Although we do not know the “true” number of topics in our datasets (assuming such a thing exists), performing k-means clustering requires selecting a k , or the number of clusters the algorithm will find. One method of doing this is the “Elbow Method”, where k-means clustering is performed over a range of k values and the within-cluster sum of square values are computed for each k . Plotting all this information together should reveal an “elbow” where an increase in k does not dramatically reduce the within-cluster sum of square value. This approach did not reveal a clear elbow for our data, though it suggested that a reasonable k might fall in the range of 10-30 clusters. To inspect this range, we performed k-means clustering for 10, 15, 20, 25, and 30 clusters. For each of those clustering solutions, we computed Silhouette scores. The Silhouette scores for all clustering solutions were quite similar (~0.3).

Finally, for each clustering solution, we generated and inspected word clouds based on the word frequency of the text assigned to each cluster. We were ultimately interested in a clustering solution that resulted in topics that seemed (i) interpretable and (ii) varied, without being redundant. This led us to choose a clustering solution with 25 topics. The selection of these parameters requires a degree of subjectivity. It will be important in future work to investigate the robustness of these results based on choice of language model and subsequent clustering decisions.

Assigning topics to each bin We assigned each bin of text in each conversation to a single topic. We labeled each topic to reflect the themes that emerged when inspecting the text of the bins assigned to each topic. Figure 1 depicts the proportion of dyads in each topic across each 30-second time bin. By looking at the color intensity in each row in Figure 1, it is possible to get a sense of the average timecourse where each topic is likely to emerge within a conversation.

Characterizing topic transitions

Across all conversations, we counted the number of times each dyad transitioned from one topic to another. These counts were represented by a topic transition matrix (Fig 2A). This approach is agnostic to the *timing* of different bins (e.g., early vs late in a conversation), it merely records topic transitions between consecutive bins. We then represented this matrix as a weighted, directed graph with each topic acting as a node in the network and the edges representing the transitions between them (Fig 2B).

Identifying topics with similar transition properties

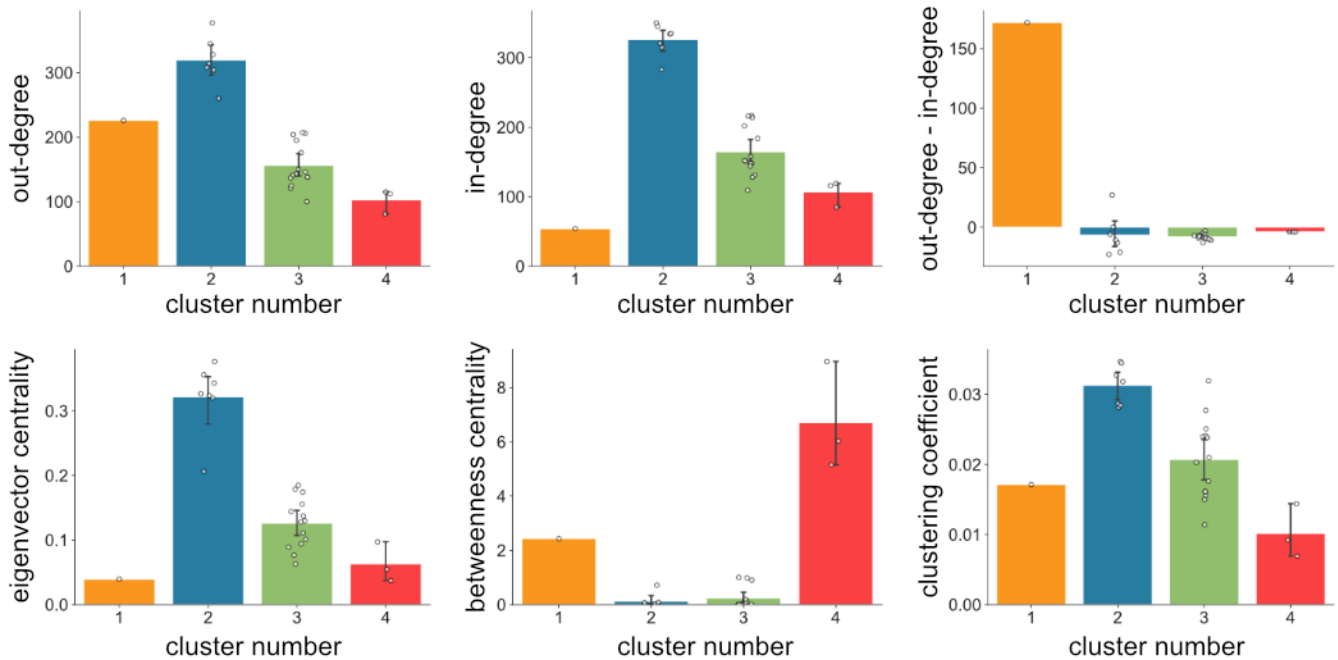
We computed six, weighted node metrics for each topic in the network: in-degree, out-degree, out-degree minus in-degree, eigenvector centrality, betweenness centrality, and clustering coefficient.

In-degree is the sum of the edge weights for edges pointing to a particular node. In our framework, in-degree represents the total number of times that people transitioned into a given topic. *Out-degree* is the sum of the edge weights for edges pointing out of a particular node. In our framework, out-degree represents the total number of times that people transitioned out of a given topic. *Out-degree minus in-degree* is the in-degree score for a particular node subtracted from the out-degree score for a particular node. A high value would indicate that people transitioned out of a given topic more than in, whereas a low value would indicate that people transitioned into a given topic more than out. *Eigenvector centrality* is a measure of influence in a network. Higher scores indicate that a topic is connected to other well-connected topics. *Betweenness centrality* is a measure of centrality based on shortest paths. Higher scores indicate more connections across groups of topics. *Clustering coefficient* is a measure of how much a particular node clusters together with different nodes. Higher scores indicate that a topic is connected to topics that are all connected to

A Cluster assignment

Cluster 1	Cluster 2	Cluster 3	Cluster 4
- Introductions	- Hometowns - Professors - Greek life - Classes: psych / neuro	- Languages - Recap convo - Disclosure - College choice - Mutuals friends	- Technology - Future plans - Recap of week - Upcoming plans - Living situation
	- Playing sports - Classes: science - Exams	- Outdoor activities - Weather experiences - What classes to take - Professional plans	- Watching sports - Food - Classes: economics

B Node metrics



C Temporal trends



Figure 3: Topics are clustered based on six different node metrics. (A) List of topics assigned to each cluster. Cluster 2 contains the candidate “launch pad” topics. (B) How topics in different clusters vary across the six different node metrics. (C) Average proportion of dyads discussing topics clustered by network features for each time bin. These subplots emphasize the temporal trends for each cluster of topics. Note that different subplots have different y-axes, indicating frequency differences between clusters.

in-degree, eigenvector centrality, and clustering coefficients. We think of these topics as “conversational launch pads”.

Launch pads topics are connected to many other topics, easily allowing conversation to flow from one topic to another. Although launch pad topics are most useful at the

start of a conversation, their high in-degree suggests that these topics are also easily returned to as the conversation progresses. The temporal trends reflect this. Though reliance on launch pad topics is highest at the start of a conversation, these topics are often used at other times as well. Launch pad

topics may be particularly well-positioned to help people find common ground, throughout the course of a conversation.

Although we focus our results and discussion on launch pad topics (Cluster 2 in this project), our clustering analysis revealed three other types of topics. Cluster 1, solely comprised of the ‘introductions’ topic, appears at the very beginning of a conversation and never again. It is characterized by having a much higher out-degree than in-degree, meaning that dyads tend to transition out of (rather than into) this topic. This makes intuitive sense. Once we introduce ourselves, there is rarely a need to introduce ourselves again (unless a new person joins the conversation). Although introductions are also used at the start of a conversation, we do not think of them as conversational launch pads. They are more akin to a conversational “preamble.” After people exchange names, they still need to think of something to say (except in cases where someone’s name prompts more discussion). Cluster 3 shows a temporal trajectory that increases as the conversation progresses. Many topics in that cluster seem to indicate a “deeper” level of conversation (e.g., mutual friends, disclosure). It is likely that launch pad topics aid a transition into these deeper topics. Cluster 4 has the least interpretable temporal trend, which may be because it is comprised of topics that were used less often overall in this dataset. The way that conversations *start* likely constrains how they *end up*. Future work will more comprehensively examine how these conversational dynamics unfold.

Conversational launch pads are defined at the level of the population. Here, we examined topic transitions in a group of unacquainted students attending the same college. Fittingly, many of the launch pad topics that emerged centered around school, something they all have in common. We expect that different populations will utilize different topics as conversational launch pads and plan to formally investigate this in future work. We think of good launch pads as being the “lowest common denominator” of topics in a given population. If people can assume they have something specific in common, it makes sense to start conversations there. We predict that the less information people have about each other, the more general their starting topic should be. Certain topics are always available (e.g., weather) and others rise and fall based on world events, like talking about a virus during a global pandemic (Reece et al., 2023).

Launch pad topics are defined by their transition properties, not their semantic content. Conversations about “Greek life” are likely quite different than conversations about “Exams”, yet our analyses reveal that both topics serve similar functions in this particular dataset. Similarly, the topics that scientists use to start conversations with strangers at an academic conference will likely be quite different than the topics music-lovers use to start their conversations with strangers at a concert. Although the semantic content may vary between groups, we expect that common starting topics will tend to have similar transition properties within each group.

Future work will also investigate the cognitive and social benefits of using launch pad topics in conversation. When we meet someone new, we need to solve a critical theory-of-mind task: What should we say to someone when we do not know anything about them? A major goal of conversation is to find common ground (Brennan, Galati, & Kuhlen, 2010; Clark & Brennan, 1991) and establish a sense of shared reality—the feeling that you have a shared understanding of some aspect of the world (Echterhoff et al., 2009; Hardin & Higgins, 1996; Rossignac-Milon et al., 2020). The use of conversational launch pads may be an efficient first step in this process. By starting with topics that can easily launch into more topics, strangers can increase the odds of hitting common ground quickly, leading to a sense of shared reality, and increased feelings of connection.

Conversational launch pads are topics that likely help strangers more than friends. When people know each other well, they can begin their conversations in idiosyncratic topics that make sense for their particular friendship. Thus, a decreasing reliance on launch pad topics could indicate the deepening of a relationship—moving away from feeling like strangers and towards feeling like friends. Future work is needed to test the implications of launch pads for these important conversational consequences.

When we meet someone new, we do not know how much our minds are aligned. Launch pad topics may help jump-start the process of finding common ground. By reducing the space of possible things to talk about to well-connected topics, people can begin the process of converging on something meaningful, specific, and shared.

Acknowledgments

We thank the members of the Dartmouth Social Systems Lab and the Computational Social Affective Neuroscience Lab for helpful discussions and feedback. This work was funded in part from an award to LJC from the National Institute on Drug Abuse (R01DA053311).

References

- Assent, I. (2012). Clustering high dimensional data. *Wiley Interdisciplinary Reviews. Data Mining and Knowledge Discovery*, 2(4), 340–350.
- Bearman, P., & Parigi, P. (2004). Cloning headless frogs and other important matters: Conversation topics and network structure. *Social forces*, 83(2), 535-557.
- Bischoping, K. (1993). Gender differences in conversation topics, 1922–1990. *Sex roles*, 28, 1-18.
- Brennan, S. E., Galati, A., & Kuhlen, A. K. (2010). Chapter 8 - Two Minds, One Dialog: Coordinating Speaking and Understanding. In B. H. Ross (Ed.), *Psychology of Learning and Motivation* (Vol. 53, pp. 301–344). Academic Press.
- Cassell, J., Gill, A., & Tepper, P. (2007). Coordination in Conversation and Rapport. *Proceedings of the Workshop on Embodied Language Processing*, 41–50.
- Cer, D., Yang, Y., Kong, S.-Y., Hua, N., Limtiaco, N., St. John, R., Constant, N., Guajardo-Cespedes, M., Yuan, S.,

- Tar, C., Sung, Y.-H., Strobe, B., & Kurzweil, R. (2018). Universal Sentence Encoder. In *arXiv [cs.CL]*. arXiv. <http://arxiv.org/abs/1803.11175>
- Christakis, N. A., & Fowler, J. H. (2007). The spread of obesity in a large social network over 32 years. *New England journal of medicine*, 357(4), 370-379.
- Clark, H. H., & Brennan, S. E. (1991). Grounding in communication. In L. B. Resnick (Ed.), *Perspectives on socially shared cognition*, (pp (Vol. 429, pp. 127–149). American Psychological Association, xiii.
- Cooney, G., Mastroianni, A. M., Abi-Esber, N., & Brooks, A. W. (2020). The many minds problem: disclosure in dyadic versus group conversation. *Current Opinion in Psychology*, 31, 22-27.
- Cui, M. (2020). Introduction to the K-means clustering algorithm based on the elbow method. *Accounting, Auditing and Finance*, 1(1), 5–8.
- Dunbar, R. I., Marriott, A., & Duncan, N. D. (1997). Human conversational behavior. *Human nature*, 8, 231-246.
- Echterhoff, G., Higgins, E. T., & Levine, J. M. (2009). Shared Reality: Experiencing Commonality With Others' Inner States About the World. *Perspectives on Psychological Science: A Journal of the Association for Psychological Science*, 4(5), 496–521.
- Fowler, J. H., & Christakis, N. A. (2008). Dynamic spread of happiness in a large social network: longitudinal analysis over 20 years in the Framingham Heart Study. *Bmj*, 337.
- Fowler, J. H., & Christakis, N. A. (2010). Cooperative behavior cascades in human social networks. *Proceedings of the National Academy of Sciences*, 107(12), 5334-5338.
- Hagberg, A., Swart, P., & S Chult, D. (2008). *Exploring network structure, dynamics, and function using NetworkX* (No. LA-UR-08-05495; LA-UR-08-5495). Los Alamos National Lab.(LANL), Los Alamos, NM (United States).
- Hardin, C. D., & Higgins, E. T. (1996). Shared reality: How social verification makes the subjective objective. In R. M. Sorrentino (Ed.), *Handbook of motivation and cognition*, Vol (Vol. 3, pp. 28–84). The Guilford Press, xxvi.
- Jucker, A. H., & Smith, S. W. (2022). Explicit and implicit ways of enhancing common ground in conversations. *Pragmatics*, 6(1), 1–18.
- Lee, H., & Chen, J. (2022). Predicting memory from the network structure of naturalistic events. *Nature Communications*, 13(1), 4235.
- Maynard, S. K. (1986). Interactional aspects of thematic progression in English casual conversation. *Text-Interdisciplinary Journal for the Study of Discourse*, 6(1), 73-106.
- McInnes, L., Healy, J., & Melville, J. (2018). Umap: Uniform manifold approximation and projection for dimension reduction. *arXiv preprint arXiv:1802.03426*.
- Parkinson, C., Kleinbaum, A. M., & Wheatley, T. (2017). Spontaneous neural encoding of social network position. *Nature Human Behaviour*, 1(5), 0072.
- Pillet-Shore, D. (2011). Doing introductions: The work involved in meeting someone new. *Communication Monographs*, 78(1), 73-95.
- Pillet-Shore, D. (2012). Greeting: Displaying stance through prosodic recipient design. *Research on Language & Social Interaction*, 45(4), 375-398.
- Pillet-Shore, D. (2018). How to begin. *Research on Language and Social Interaction*, 51(3), 213-231.
- Reece, A., Cooney, G., Bull, P., Chung, C., Dawson, B., Fitzpatrick, C., Glazer, T., Knox, D., Liebscher, A., & Marin, S. (2023). The CANDOR corpus: Insights from a large multimodal dataset of naturalistic conversation. *Science Advances*, 9(13), eadf3197.
- Rossignac-Milon, M., Bolger, N., Zee, K. S., Boothby, E. J., & Higgins, E. T. (2020). Merged minds: Generalized shared reality in dyadic relationships. *Journal of Personality and Social Psychology*. <https://doi.org/10.1037/pspi0000266>
- Schwuyck, M. E., Du, M., Natarajan, P., Chwe, J. A., & Parkinson, C. (2023). Neural encoding of novel social networks: evidence that perceivers prioritize others' centrality. *Social Cognitive and Affective Neuroscience*, 18(1), nsac059.
- Singh, J. (2005). Collaborative networks as determinants of knowledge diffusion patterns. *Management science*, 51(5), 756-770.
- Templeton, E. M., Chang, L. J., Reynolds, E. A., Cone LeBeaumont, M. D., & Wheatley, T. (2022). Fast response times signal social connection in conversation. *Proceedings of the National Academy of Sciences of the United States of America*, 119(4).
- Thornton, M. A., & Tamir, D. I. (2017). Mental models accurately predict emotion transitions. *Proceedings of the National Academy of Sciences*, 114(23), 5982-5987.
- Thornton, M. A., & Tamir, D. I. (2021). People accurately predict the transition probabilities between actions. *Science Advances*, 7(9), eabd4995.
- Thornton, M. A., Rmus, M., Vyas, A. D., & Tamir, D. I. (2023). Transition dynamics shape mental state concepts. *Journal of Experimental Psychology: General*.
- Tickle-Degnen, L., & Rosenthal, R. (1990). The Nature of Rapport and Its Nonverbal Correlates. *Psychological Inquiry*, 1(4), 285–293.
- Yang, Y. (2019). *The management of topics in ordinary conversation* (Doctoral dissertation, University of York).