


*Patterns of interaction-dominant dynamics
in individual versus collaborative memory
foraging*

**Janelle Szary, Rick Dale, Christopher
T. Kello & Theo Rhodes**

Cognitive Processing
International Quarterly of Cognitive
Science

ISSN 1612-4782

Cogn Process
DOI 10.1007/s10339-015-0731-8



 Springer

Your article is protected by copyright and all rights are held exclusively by Marta Olivetti Belardinelli and Springer-Verlag Berlin Heidelberg. This e-offprint is for personal use only and shall not be self-archived in electronic repositories. If you wish to self-archive your article, please use the accepted manuscript version for posting on your own website. You may further deposit the accepted manuscript version in any repository, provided it is only made publicly available 12 months after official publication or later and provided acknowledgement is given to the original source of publication and a link is inserted to the published article on Springer's website. The link must be accompanied by the following text: "The final publication is available at link.springer.com".

Patterns of interaction-dominant dynamics in individual versus collaborative memory foraging

Janelle Szary¹ · Rick Dale¹ · Christopher T. Kello¹ · Theo Rhodes²

Received: 1 November 2014 / Accepted: 20 July 2015
© Marta Olivetti Belardinelli and Springer-Verlag Berlin Heidelberg 2015

Abstract The extent to which a cognitive system's behavioral dynamics fit a power law distribution is considered indicative of the extent to which that system's behavior is driven by multiplicative, interdependent interactions between its components. Here, we investigate the dynamics of memory processes in individual and collaborating participants. Collaborative dyads showed the characteristic collaborative inhibition effect when compared to nominal groups in terms of the number of items retrieved in a categorical recall task, but they also generate qualitatively different patterns of search behavior. To categorize search behavior, we used multi-model inference to compare the degree to which five candidate models (normal, exponential, gamma, lognormal, and Pareto) described the temporal distribution of each individual and dyad's recall processes. All individual and dyad recall processes were best fit by interaction-dominant distributions (lognormal and Pareto), but a clear difference emerged in that individual behavior is more power law, and collaborative behavior was more lognormal. We discuss these results in terms of the cocktail model (Holden et al. in *Psychol Rev* 116(2):318–342, 2009), which suggests that as a task becomes more constrained (such as through the necessity of collaborating), behavior can shift from power law to lognormal. This shift may reflect a decrease in the dyad's

ability to flexibly shift between perseverative and explorative search patterns. Finally, our results suggest that a fruitful avenue for future research would be to investigate the constraints modulating the shift from power law to lognormal behavior in collaborative memory search.

Keywords Interaction dominance · Power laws · Multi-model inference · Collaborative recall · Collaborative memory · Memory foraging · Lévy processes

Introduction

The category-based recall or semantic fluency task has a long tradition of use in empirical investigations of memory (Bousfield and Sedgewick 1944). Typically, participants in these tasks are asked to recall as many items that belong to a given category as possible, within some time window. The pattern of responses is often clustered: Short bursts of many items are remembered together, interspersed with longer intervals in which nothing can be recalled. Items that are recalled in temporal clusters also tend to be semantically similar (Bousfield and Sedgewick 1944; Troyer et al. 1997; Hills et al. 2012). This retrieval pattern, with frequent, clustered short intervals but occasional long intervals between clusters, is similar to the pattern of path lengths exhibited by animals foraging for food or other resources in physical environments (for example, see Sims et al. 2007; Viswanathan et al. 2011). This led Rhodes and Turvey (2007) to conceptualize memory retrieval as a foraging process across a cognitive or information landscape. They found that the intervals between consecutive recalls are consistent with Lévy flights—random walks in which path lengths (intervals) are power law distributed such that each length L is observed with probability $P(L) \sim L^{-\mu}$, and

This article is part of the Special Issue on 'Complexity in brain and cognition' and has been edited by Cees van Leeuwen.

✉ Janelle Szary
jszary@ucmerced.edu

¹ Cognitive and Information Sciences, University of California, Merced, 5200 North Lake Road, Merced, CA 95343, USA

² Psychology, State University of New York at Oswego, 7060 Route 104, Oswego, NY 13126-3599, USA

$0 < \mu \leq 3$. The Lévy flight has received a considerable amount of attention in the animal foraging literature, where it has been suggested that Lévy flights are an optimal search strategy (Viswanathan et al. 1999; Viswanathan et al. 2011), but it has also been noted that other search processes (strategies) give rise to similar resulting search patterns, and in fact these patterns may not be optimal for all environments (Benhamou 2007; Plank and James 2008).

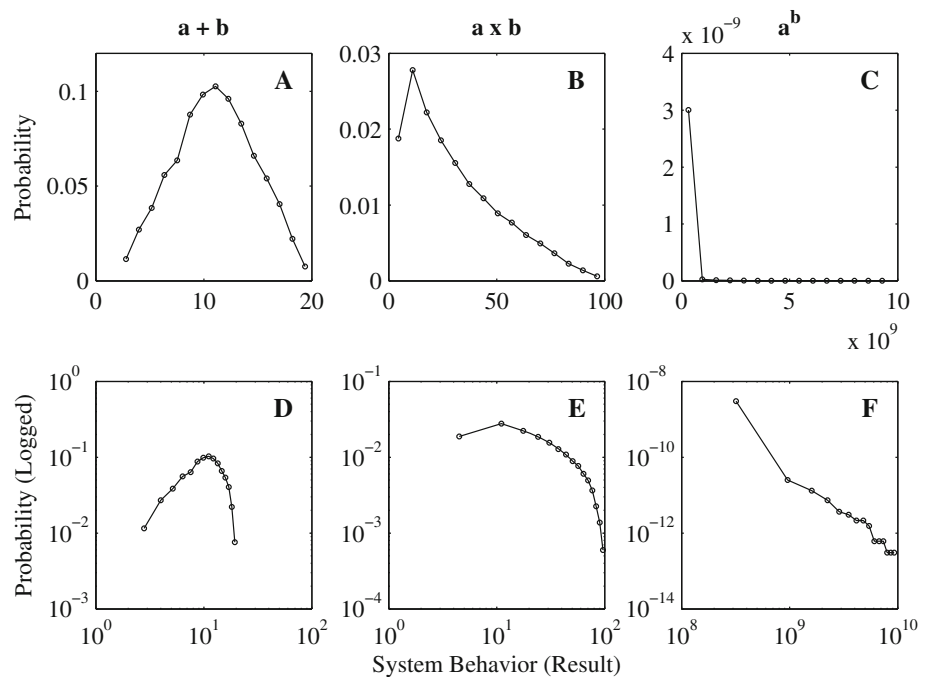
In the cognitive domain, Hills et al. (2012) have used the theoretically different but empirically related framework of area-restricted search modeled as a correlated random walk to describe the same free recall behavior. In the ecological domain, area-restricted search describes a strategy where a forager searches a restricted, or local, area with bountiful resources until the availability of those resources declines. At that point, the forager will move to a new area and begin the process again (Karieva and Odell 1987). Notably, both strategies (Lévy flights and area-restricted search) produce approximately similar behavior: frequent short path lengths and occasional long path lengths (see Reynolds 2010). Using a spatial search task with human participants, Hills et al. (2013) found evidence that both Lévy flights and area-restricted search might be used strategically by the same agent depending on the resource context. Since the cognitive and information landscape of memory is abstract, and presumably different across individuals, the actual foraging processes and the cognitive mechanisms which generate them can only be inferred from the resulting recall patterns. A common theme across these perspectives, though, is that retrieval patterns might be indicative of the underlying cognitive processes, whatever they may be.

In the exploratory project presented here, we consider whether retrieval patterns differ when remembering is performed collaboratively. Remembering is often performed in a collaborative context, but empirically this interactive context is almost always associated with negative performance outcomes (Rajaram and Pereira-Pasarin 2010; but see Harris et al. 2010). Increasingly, researchers are considering the processes by which the joint action of remembering occurs (Tollefsen and Dale 2012) to help explain why collaborative memory performance is worse than (and conversely, why it can sometimes be *better than*) individual memory performance. Here, we consider both performance outcomes and retrieval patterns to investigate whether cognitive search processes change depending on social context. We predict that interactivity will alter the dynamics of individuals in a dyad compared to the dynamics of an individual performing the same task alone. Work by Tollefsen (2006), Schmidt and O'Brien (1998), Yoshimi (2012), and Dale et al. (2013), among others, suggests that the coupling between two agents will result in qualitative change in the distribution of their retrieval intervals.

In order to quantify timing distributions across conditions, we draw on tools from complexity science and statistical mechanics. In this domain, statistical distribution matching can be used to determine the nature of interaction among the component parts of a complex system. Cognitive scientists have begun using these techniques to determine whether cognitive processes are more component dominant, in which case a system's subcomponents act and contribute independently to system-wide behavior, or interaction dominant, in which case the effects of any one component depend on the behavior of another, so that its effects on system-wide behavior are nonlinear (Van Orden et al. 2003; Holden et al. 2009; Stephen and Mirman 2010). To illustrate the distinction between component dominance and interaction dominance, consider a simple system S composed of two random variables a and b , which can each range from 1 to 10. System behavior depends not only on the values of a and b , but also on the relationship between them (see Stephen and Mirman 2010). Mathematically, that relationship might be addition ($S = a + b$), multiplication ($S = a \times b$), or even a power relationship ($S = a^b$). In the additive system, possible values of S range from 2 to 20, and a probability density function will be normally distributed (Fig. 1a). This system is linearly decomposable in the sense that the effects of each component are independent: a will have the same effect on S regardless of the value of b . This is the strictest example of a component-dominant system. In the multiplicative system, $S = a \times b$, resulting values of S range from 1 to 100, and the probability density function shows that the median and mode have shifted to lower values, with a longer tail extending to higher values (Fig. 1b). This skewed distribution results from the multiplicative interaction between the variables. That is, the effect of variable a on S will be different depending on the value b . The effects of each variable are no longer entirely independent, but instead depend on the context of each other's values. Compared to the additive system, this system is more interaction dominant. In the system $S = a^b$, the effect of interaction is further amplified, resulting in the increased skew of the probability density function (Fig. 1c).

These three examples do not exhaust the space of possible systems and interactions, but are meant only to illustrate the variation between component dominance and interaction dominance in systems. Two other important concepts in this literature, which we will return to in the conclusion, are interdependence and feedback. In the examples above, the values of a and b are selected independently from one another. That is, a does not actually change the value of b , but may change the effects of b on system-wide behavior. However, a considerable amount of literature has noted the pervasiveness of scaling laws in cognitive systems, whereby a system exhibits self-

Fig. 1 Example probability distributions for three types of systems, where system behavior $S = a + b$ (left column), $S = a \times b$ (middle column), and $S = a^b$ (right column). The top row shows normalized histograms for 10,000 actual trials, where a and b are random numbers between 1 and 10. The bottom row shows the same data plotted on log-log coordinates



similarity at different scales (Kello et al. 2010) and the relationship between component variables becomes more complex. Kello and colleagues have described, more specifically, the pervasiveness of $1/f$ scaling (Kello et al. 2007, 2008), which signifies long-range correlations in the temporal dynamics of a system, so that effects at any period of time have nonnegligible effects into the future. One mechanism by which long-range correlations can be generated is through multiplicative feedback effects, where the perturbations to one component perturb other components, which in turn may come to perturb the initial component once again. That is, in these more complex systems (as compared to the simple, independent-component system described earlier), changes to a may cause changes to b , indicating *interdependence* of the two variables, and these changes may reverberate back to a , illustrating an example of a feedback effect.

Stephen and Mirman (2010) have used the method of statistical distribution matching to fit distributions of eye movements to five candidate statistical distributions which vary from component dominant to interaction dominant. They found that their empirical data were best fit by and described as an interaction-dominant system. Rhodes (2013) used statistical distribution matching to identify patterns of interaction dominance in individual memory retrieval during a category-based free recall task. Here, we expand upon Rhodes' study by using multi-model inference to investigate whether interaction-dominant patterns persist during collaborative recall. Because two interacting agents perform approximately as one coherent group, are individual-level dynamics constrained by higher-order

constraints of the dyad, as a unit of analysis unto itself? Will the pattern of results from a cohesive dyad unit be comparable to that of an individual? We show that multi-model inference can tease apart changes in the dynamics of individual and collaborative recall behavior, and show how these patterns might relate to the success of group memory behavior, in terms of number of items retrieved. Finally, we use latent semantic analysis (LSA) to provide insights into the semantic similarity of items retrieved by the members in each group.

Materials and methods

For the current study, either an individual or a dyad completed two 20-min semantic fluency tasks. One semantic category was the set of all animals, and the other was cities and towns in California. The two categories were presented in counterbalanced order, with a break between sessions. Here, we discuss the results from only the animals category, so that our results can be considered in the context of other work (e.g., Bousfield and Sedgewick; Rhodes and Turvey 2007; Rhodes 2013; Thompson and Kello 2013) on semantic fluency.¹ In addition to comparing the number of items recalled, we employ multi-model inference to characterize the distributional properties of the time series of recall events in order to investigate the search processes employed in the different conditions, and we use LSA to

¹ We refer the reader to Szary et al. (2015) for a description of and results from the California locations task.

explore semantic differences between members of each group condition.

Participants

Participants were recruited from a subject pool of University of California, Merced undergraduate students who participated for course credit. All participants reported fluency in English. The individual condition included 12 participants (4 males, 8 females), while the dyadic condition included 10 dyads composed of 20 participants (2 males, 18 females). None of the dyads included partners who knew one another prior to the experiment. Across conditions, the average age was 19.75 (SD = 1.55 years). Informed consent was obtained from all participants included in the study. All procedures performed in studies involving human participants were in accordance with the ethical standards of the institutional and/or national research committee and with the 1964 Helsinki Declaration and its later amendments or comparable ethical standards.

Procedure

Before the study began, dyads were given 5 min to introduce themselves while the experimenter left the room. All participants were seated comfortably at a table in a small experiment room, with dyads facing each other across the table. Participants wore Shure microphone headsets, connected to a computer in another room via a USB preamp, to record verbal responses. Participants completed two 20-min recall tasks, with a short break of roughly 3 min in between. For each recall task, participants recalled as many items as possible from either the “animal” category or the “cities and towns in California” category, but as mentioned previously, only results from the animal category are presented here. Before receiving their category assignment, participants were given the following instructions: “In a moment, I’ll give you the name of a category. Your goal is to (work together to) think of as many items from that category as you can. When you think of an item, just say it out loud. You can be as specific or as general as you wish. For example, if the category were ‘food’, you could say ‘fruit’, and you could also say ‘orange’ or ‘mandarin orange’.” In order to minimize awkwardness or discomfort, dyads were invited to converse freely with their partners during each session, but encouraged to stay focused on the task of coming up with as many items as possible.

Datasets

Participants’ responses were transcribed using the Praat audio analysis software. The utterance for each retrieved item was isolated, and its onset time determined, but

extraneous conversation and noncategory utterances were ignored.² Two variables of interest were collected for each dataset: (1) score (the number of unique items retrieved) and (2) the interretrieval intervals (IRIs; the milliseconds between the onsets of consecutive recalls).

In order to investigate the effects of collaboration on the recall task, we consider performance at both the individual level (one participant) and the group level (two participants). Group-level datasets are created by merging sequences of recalled items from individuals and treating them as one. The two original datasets we collected were single participants working alone (*independents*, $n = 12$) and pairs of two participants working collaboratively (*dyads*, $n = 10$). For the purposes of comparison, we also generated three artificial datasets: *single partners* ($n = 20$), *nominal pairs* ($n = 66$), and *mismatched pairs* ($n = 180$). The datasets, as illustrated in Table 1, vary on three dimensions: the number of participants included in the analysis (one or two), the interaction condition (worked individually or worked collaboratively), as well as whether the datasets were real or artificially generated. The single partners dataset is created by isolating individual-level behavior from each participant who worked in a dyad, which allows us to see how individuals behave during collaboration. The mismatched dataset is all possible pairings of these collaborating individuals (all unique combinations of single partners), but matched with those with whom they had not actually participated. The nominal dataset is all possible pairings of individuals who worked alone (all unique combinations of independents). The mismatched and nominal groups allow us to tease apart whether any differences between individual and group performance are a product of having a different number of participants, or whether the collaboration itself is having an effect.

Data analysis

As noted, we compare the recall performance of the different conditions by considering both the outcome (score) and the process of retrieving items from memory (IRIs). IRI distributions are characterized using multi-model inference, in which the relative fit of a set of candidate models is determined and a best fitting model can be selected (see Rhodes 2013). Here, we consider the normal, exponential, gamma, lognormal, and Pareto (power-law) distributions. The normal distribution is indicative of a system with additive dynamics, while the exponential and

² While obvious non-category utterances were removed, some items were retained such as imaginary, extinct, or incorrectly named animals. A second analysis was performed where these items were removed, but the relative between-condition scores were unaffected.

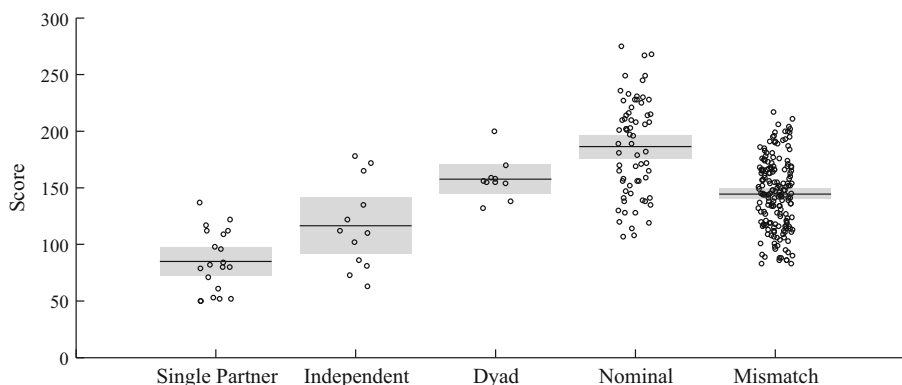
Table 1 Datasets

	Interaction condition	
	None (independent)	Collaborative
Level of analysis		
One participant	Independent ($n = 12$)	Single partner ($n = 20$)
Two participants	Nominal ($n = 66$)	Dyad ($n = 10$) Mismatch ($n = 180$)

The composition of the five datasets used in analyses is illustrated. The two bolded datasets (independent, dyad) are those originally collected. The other three (single partner, nominal, and mismatch) are artificial datasets generated for comparison

gamma distributions are indicative of additive systems with slight multiplicative interactions. The two models that have received the most attention in the domain of search processes are the lognormal and Pareto distributions, both of which are indicative of systems with multiplicative, interaction-dominant dynamics. The lognormal distribution is considered a special case of multiplicative interaction in which feedback effects are sufficiently constrained (Farmer 1990), while the Pareto distribution is thought to be indicative of a multiplicative system with stronger feedback effects and interdependence. In multi-model inference, each candidate model is fit to each IRI distribution using maximum likelihood estimation (MLE). At each model's best fitting parameters, negative log-likelihoods are used to determine Akaike's Information Criterion (AIC), which is used as a measure of the relative fit of the candidate models (see Burnham and Anderson 2002). In addition to determining the best fitting of the distributions using AIC, the log-likelihoods from MLE are used to compute log ratios for conditions in which two models appear to fit. The log ratio between two models is simply the difference in their log-likelihoods and can be used to directly quantify the goodness-of-fit of one model over the other (Stephen and Mirman 2010; Singer and Willett 2003).

Fig. 2 Data points show scores by condition (plus x -jitter). Lines show means, boxes show 95 % confidence intervals for the means



For those distributions that are best fit by the Pareto distribution, we consider the slope parameter $-\mu$, as generated by MLE, in the context of the literature on Lévy foraging. As noted above, $1 < \mu \leq 3$ is taken as indicative of Lévy foraging, and $\mu = 2$ is considered optimal (Viswanathan et al. 1999). For each dataset best fit by the Pareto, we calculated the magnitude of the deviation between the observed μ and the theoretically optimal $\mu = 2$. We then test whether these μ -deviation values are correlated with task performance, as in previous research (e.g., Rhodes and Turvey 2007).

Finally, we look deeper into the effects of collaboration on group performance using LSA. LSA is a statistical tool that can quantify the similarity between two words or groups of words in some semantic space. Here, we use the online LSA tool (lsa.colorado.edu) described in Landauer et al. (1998), specifying the semantic space as general reading (up to the first year of college). For the set of retrieved items from each participant, a term-document matrix is created which notes the occurrence of each word from the semantic space, without log entropy weighting. This matrix is then reduced using singular value decomposition. Finally, the cosine between the matrices for two datasets represents their similarity on a scale of -1 to 1 , where 1 would reflect identical datasets.

Results

Number of items recalled

Mean retrieval scores for each condition and category are shown in Table 2. A one-way analysis of variance showed that the effect of condition was significant, $F(4, 283) = 41.56, p < 0.0001$. More items were retrieved by dyads ($M = 157.7, SD = 18.32$) than by independents ($M = 116.58, SD = 39.09$), a difference found to be statistically significant in a two-sample t test, $t(20) = 3.04, p < 0.01$. Among both individual-level datasets,

Table 2 Mean number of unique items retrieved and the number of items repeated (standard deviation of the means in parentheses)

	Retrieved	(SD)	Repeated	(SD)
Single partner ($n = 20$)	84.85	(27.09)	7.15	(9.84)
Independent ($n = 12$)	116.58	(39.09)	27.67	(81.72)
Dyad ($n = 10$)	157.70	(18.32)	26.30	(20.30)
Nominal ($n = 66$)	186.38	(42.81)	102.12	(109.18)
Mismatch ($n = 180$)	144.59	(31.08)	39.41	(18.63)

independent participants found significantly more items than single partners isolated from their collaborators ($M = 84.85$, $SD = 42.81$), $t(30) = 2.71$, $p < 0.05$. Among group-level datasets, actual dyads were outperformed by nominal groups ($M = 186.38$, $SD = 42.81$), with marginal significance $t(74) = 2.08$, $p < 0.05$, and did slightly better than mismatched groups ($M = 144.59$, $SD = 31.08$), although this difference did not reach significance. Table 2 summarizes the number of unique items retrieved in each condition, as well as the number of repeats that occurred. Figure 2 shows scores for all conditions.

Distribution of recall events

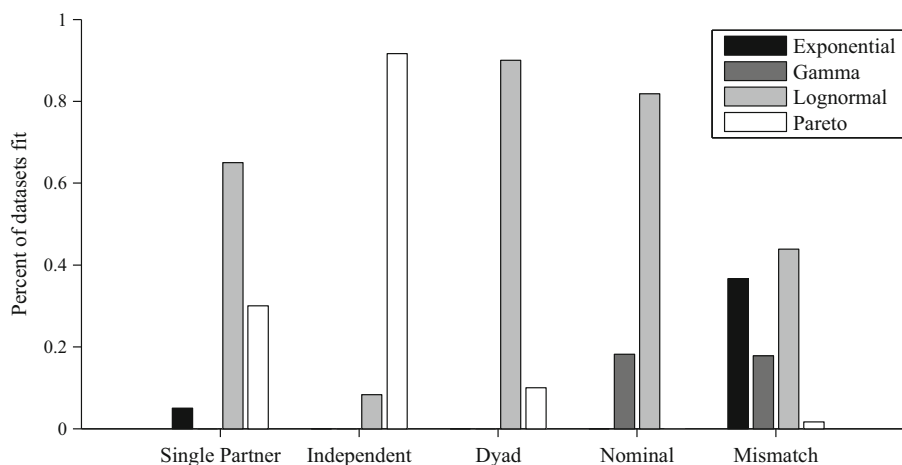
Table 3 shows the number of datasets from each condition that were best fit by each model, according to multi-model

inference. None of the 288 datasets were best fit by the normal distribution. Sixty-seven were best fit by the exponential, 44 by the gamma, 156 by the lognormal, and 21 by the Pareto. Figure 3 plots the percentage of datasets from each condition that is best fit by each of these four models. Notably, the exponential and gamma distributions only appear as fits to the artificial datasets, but not to either of the two experimental conditions (independent and dyad). All datasets from the experimental conditions are best fit by either the lognormal or Pareto distribution, which are the two models considered to be indicative of interaction dominance. However, the majority of dyad datasets (9/10) were best fit by the lognormal distribution, while the majority of independent datasets (10/11) were best fit by the Pareto distribution. Observed distributions for dyad and independent datasets, along with the best fitting distribution at their estimated parameter values, are shown in Fig. 4. When considering the individual-level behavior of single partners extracted from the collaborating dyads, some single partners (6/20) maintained the Pareto distribution that characterizes individuals working independently, but the majority (13/20) switched to behavior fit by the lognormal distribution. A Chi-square test of independence showed a significant relationship between condition and best model fit, $\chi^2 (12, N = 288) = 209.07$, $p < 0.0001$.

Table 3 Counts of the number of the best fitting distributions from each condition, from multi-model inference

	Normal	Exponential	Gamma	Lognormal	Pareto
Single partner ($n = 20$)	0	1	0	13	6
Independent ($n = 12$)	0	0	0	1	11
Dyad ($n = 10$)	0	0	0	9	1
Nominal ($n = 66$)	0	0	12	54	0
Mismatch ($n = 180$)	0	66	32	79	3

Fig. 3 Normalized counts of the best fitting distributions for each condition



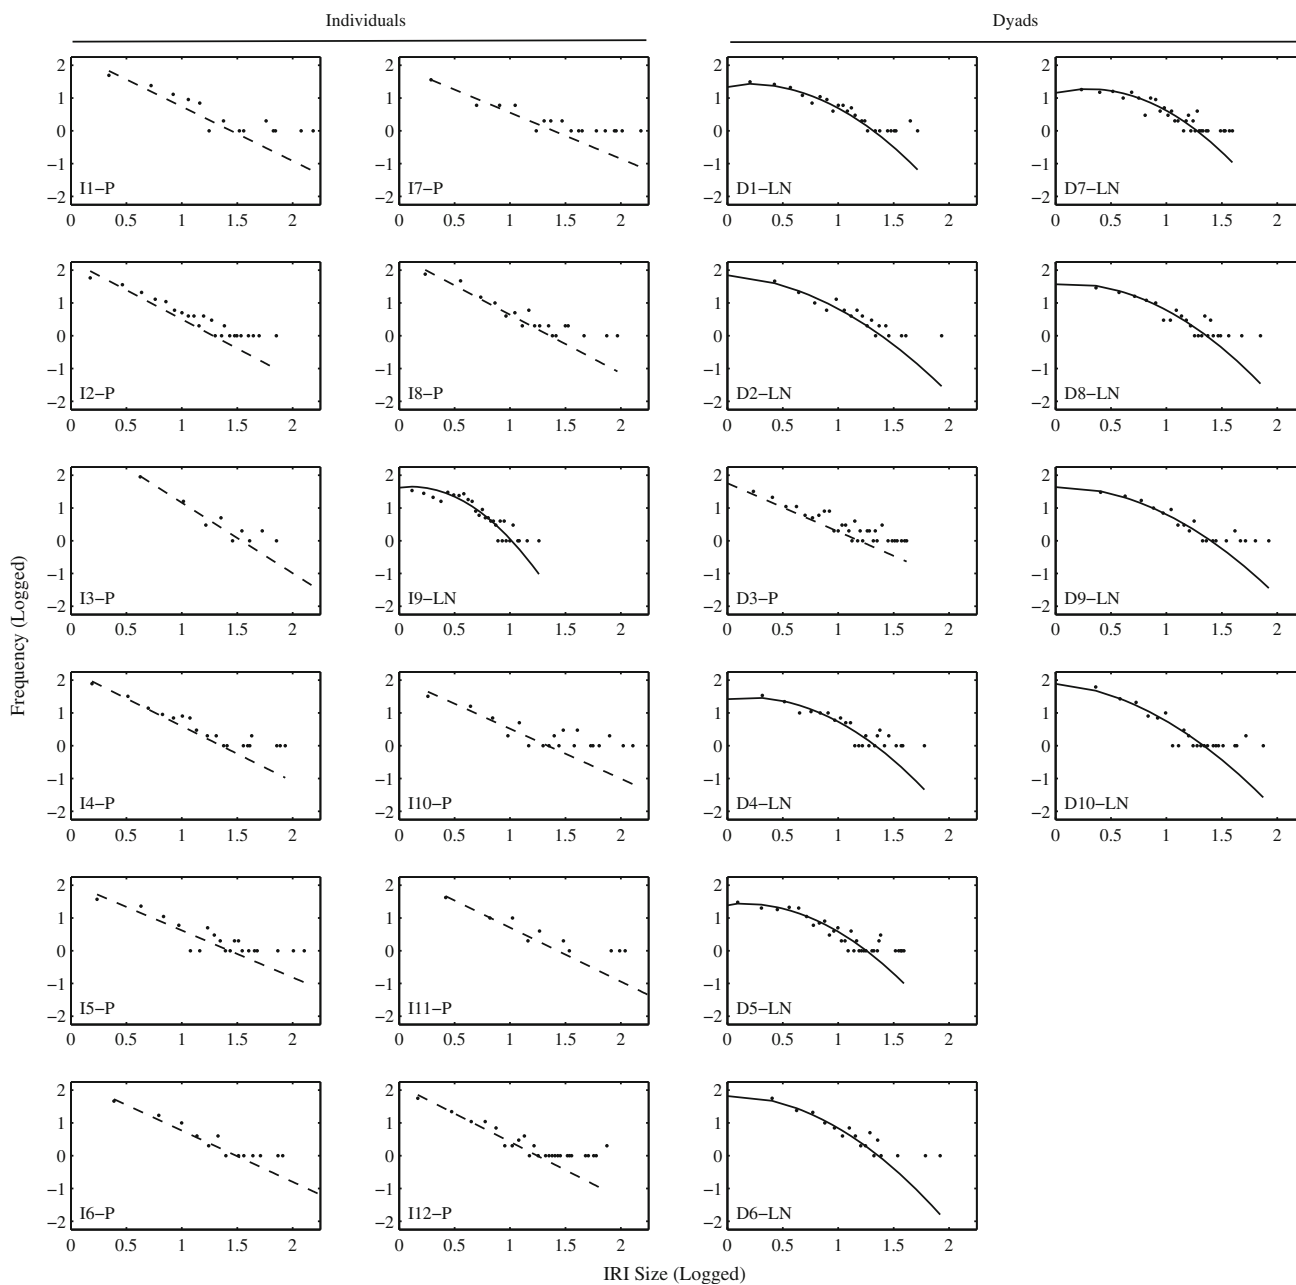


Fig. 4 IRI distributions for each individual (*two left columns, I1–12*) and dyad (*two right columns, D1–10*), along with the best fitting distributions for each. *Dotted lines* (with ‘-P’ suffix) are Pareto fits; *solid lines* (with ‘-LN’ suffix) are lognormal fits

Log ratios

To quantify the relative degree to which the lognormal and Pareto distributions fit our data, we computed the log ratio, which is the difference in log-likelihoods from MLE. The log ratios for independents favored the Pareto ($M = 36.13$, $SD = 41.47$), while the log ratios for dyads favored the lognormal ($M = -332.76$, $SD = 277.87$). Log ratios for the artificial single partner, nominal, and mismatched datasets also all favored the lognormal distribution

($M = -18.98$, -522.55 , -691.99 , $SD = 36.48$, 327.23 , 636.02). The difference in mean log ratios was significantly different between independents and dyads, $t(20) = 4.56$, $p < 0.001$, and between independents and single partners, $t(30) = 3.93$, $p < 0.001$. Although both single partners and dyads are fit by the lognormal distribution most often, their log ratios are significantly different, $t(28) = 5.05$, $p < 0.001$. Dyads are not significantly different from nominal or mismatched groups, though. Log ratios are plotted in Fig. 5.




Fig. 5 Log ratios for Pareto compared to lognormal for each condition (plus x -jitter). Positive log ratios (above the dotted line, $y = 0$) indicate better relative fit for the Pareto distribution; negative log ratios indicate a better relative fit for the lognormal. Data points below $y = -1000$ (appearing only in the mismatch condition) are not shown. Lines and boxes show means and the 95 % confidence intervals for the means

Lévy foraging

MLE was used to calculate the best fitting slope parameter $-\mu$ for each of the datasets best described by the Pareto distribution. All Pareto-distributed datasets had μ consistent with Lévy foraging, $1 < \mu \leq 3$ ($M = 1.49$, $SD = 0.10$). Deviations between each estimated μ and the theoretically optimal $\mu = 2$ were calculated. Across all datasets, the μ -deviation was negatively correlated with score $r(19) = -0.64$, $p < 0.005$. That is, the closer the slope parameter of the search distribution is to 2, the higher the score. Correlations are shown in Fig. 6. Among just the independent participants, this correlation persists, $r(9) = -0.70$, $p < 0.05$, but it disappears for all other datasets.

LSA

Similarity ratings between the two participants from each group-level dataset did not differ between dyad groups ($M = 0.772$, $SD = 0.04$) and nominal groups ($M = 0.772$, $SD = 0.05$), but were slightly lower for mismatched groups ($M = 0.757$, $SD = 0.05$). Similarity measures are plotted in Fig. 7. A one-way analysis of variance did not show a main effect of group condition. Across all groups, the LSA similarity measure was positively correlated with score $r(254) = 0.233$, $p < 0.001$, but the pattern changes within groups. That is, both mismatched and nominal groups showed positive correlations between LSA similarity and score, but this only achieved significance for the mismatched group, $r(178) = 0.248$, $p < 0.001$, which




Fig. 6 Correlation of score and deviation of μ -parameter estimate. Line shows least-squares fit




Fig. 7 Data points show LSA similarity scores for group conditions (plus x -jitter). Lines show means; boxes show 95 % confidence intervals for the means

likely drove the significant effect across condition. Within just the dyads, however, there is no significant correlation between LSA similarity and score, but there is a slight trend in the reverse direction. That is, for collaborating dyads, there is a subtle (though nonsignificant) negative relationship: Increased LSA similarity may be related to decreased score. Figure 8 plots LSA similarity measures against score for all groups.

Discussion

Consistent with previous work (Barnier et al. 2008; Szary and Dale 2013, 2014), we found that dyads retrieved significantly more unique items than individuals, but significantly less than a nominal grouping of those individuals,

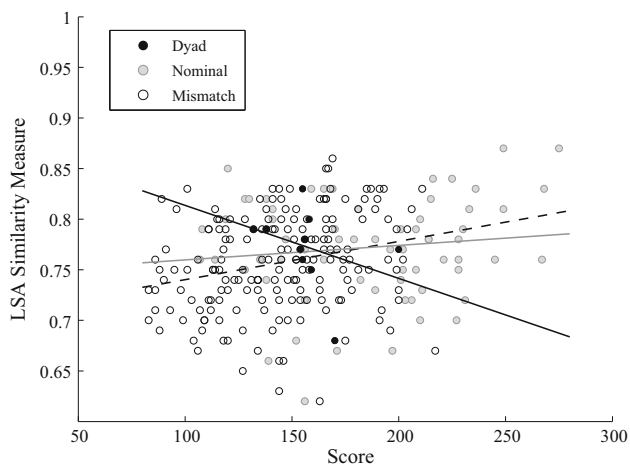


Fig. 8 Correlation of LSA similarity rating and score, by group. Lines show least-squares fits for dyads (black line $y = -0.00072 \times x + 0.89$), nominal groups (gray line $y = 0.00014 \times x + 0.75$), and mismatched groups (dotted line $y = 0.00038 \times x + 0.7$)

demonstrating the classic collaborative inhibition effect (Rajaram and Pereira-Pasarin 2010). Our task differs from the larger body of collaborative inhibition tasks, though, in that it tests the retrieval of existing knowledge using a categorical free recall task, instead of requiring lists to be learned during the task. We also find that correctly paired dyads retrieved more items than incorrectly paired (mismatched) groups, but this trend did not achieve significance.

In terms of the distribution matching results, we find that only the artificially generated datasets are ever best fit by component-dominant models (exponential and gamma). This is unsurprising, given that these datasets actually contain two independent components that are only combined afterward. Of the two models indicative of interaction dominance, we replicate Rhodes' (2013) finding that individuals' retrieval patterns are almost entirely best fit by the Pareto distribution, which is considered to be the strongest indicator of an interaction-dominant system. When two people are working collaboratively in a dyad, the dyad as a unit unto itself generally produces retrieval patterns that are lognormally distributed, indicative of an interaction-dominant system with relatively stronger constraints (such as the need to work collaboratively with another person). Interestingly, although the Pareto distribution is prevalent in retrieval patterns of individuals working alone, it is comparatively rare in the retrieval patterns of individuals working as members of a collaborating dyad. The majority of these collaborating individuals produced lognormally distributed retrieval patterns.

We looked deeper into this result by computing the log ratio for Pareto over lognormal distribution fit. Given the counts of best fitting distributions for each condition, it is unsurprising that the independent condition is the only one

where retrieval patterns are reliably best defined by the Pareto distribution (indicated by the positive mean in Fig. 5). By illustrating the relative fit of Pareto versus lognormal distributions, the log ratio shows that while individuals working collaboratively have retrieval patterns that are most often fit by the lognormal, they are still somewhat “Pareto like,” as indicated by the small negative value of their mean log ratio. Of the other lognormally inclined retrieval distributions, those from dyads are the next most Pareto like, followed by nominal groups. Mismatched group datasets show the strongest fit to lognormal as compared to Pareto, indicating that system behavior is increasingly constrained.

Of all datasets fit by the Pareto distribution, we computed the difference between the estimated slope parameter (from MLE) with the theoretically optimal value for Lévy flights. We found a significant correlation between this difference and performance—slope parameters closer to the theoretically optimal $\mu = 2$ were associated with higher retrieval scores. Within conditions, this correlation only remains significant for independent participants, but we note that the other conditions have significantly less power. Accurately determining whether the effect is driven by independents, or whether the correlation remains for all conditions, will require additional data.

Although more detailed and rigorous investigation of the content of retrieved items is necessary, LSA provides initial insights about the semantic similarity between the sets of items retrieved by the two members of each group. Although LSA similarity measures did not differ drastically across the three group conditions, our results suggest that semantic similarity may play different roles for different group compositions. That is, we find an overall positive correlation between semantic similarity and score—but this effect is largely driven by the artificial (mismatched and nominal) groups. The fact that this correlation is only present for the artificial groups may reflect that individuals who remembered more items tended to remember items that were more or less generic than the less successful individuals, which would cause consistent but noninteresting (for the present purposes) patterns in similarity measures. What is more interesting, here, is that for actual collaborating dyads, we see a trend toward the opposite relationship: Semantic similarity may be associated with *decreased* performance (but more work is necessary to see whether this trend is retained with increased data points). If this result does hold, one explanation worth investigating is that when collaborative dyads employ a “divide-and-conquer” strategy across information space, they are able to recall more items than when exploring the space together. In the collaborative memory literature, there is also precedent for this pattern (e.g., Hollingshead

2001). This suggests that optimal collaboration may require semantic divergence at either a strategic level (choosing different areas of focus) or even a component level (where different participants have different areas of expertise). LSA is a relatively coarse measure of the semantic recall trajectories in our dataset, so further research is needed to build a fuller picture of how the content of retrieved items may change in itself, or may cause changes to the actual search processes in different collaborative conditions.

Conclusion

The temporal distributions of memory retrieval differ depending on social-interaction contexts. Specifically, when participants work collaboratively on a free recall task, their recall behavior becomes less Pareto and more lognormal, on both the individual and group levels. This indicates that collaborative memory behavior should be regarded as an interaction-dominant process, but with feedback constraints. We note, however, that our results do not indicate a perfect fit to either lognormal or Pareto distributions for any dataset. The log ratio results, which show that interacting dyads have distributions that are closer to being fit by the Pareto as compared to distributions from the artificial, noninteracting groups, suggest that collaborating groups may still exhibit some interdependent feedback processes. In fact, Holden et al. (2009) propose a “cocktail model” in which a system’s behavior is a mixture of lognormal and power law, the proportions of which are determined by the existing constraints on component interactions. Constraints, which are decreases in degrees of freedom for the behavior, could come from the nature of the task, the collaborative context, or simply the current state of the mind and body.

A cocktail model makes sense given the perspective of “active internalism” described by Yoshimi (2012), in which two relatively independent systems (two people) may have some intrinsic internal dynamics, but those dynamics can become coupled through interaction. The partial independence of their cognitive processes, along with the constraints introduced by their interaction, may bring about stable behavior that flows fluidly into a performance that “stands between the two extremes because it combines independent, random variables with multiplicative interactions” (Holden et al. 2009, p. 321). This is not to say that each separate system (each person), unto itself, would exhibit Gaussian behavior. Instead, it suggests that there are processes in each system that are *not* direct participants in the coupling. Aspects of each system that do not become coupled may act as sources of instability in the combined behavior that renders lognormal distributions.

Considering our findings from this perspective, we suggest that a common process may underlie both individual

and collaborative memory search, but that the constraints imposed by collaboration transition behavior (on the individual level) from power law dominant to lognormal dominant. This transition reflects a decrease in the feedback effects which allow for small perturbations to be amplified over time and space (scaling laws). In the complex systems literature, this amplification property leads power law distributions to be considered symptomatic of (although not sufficiently indicative of) self-organizing systems operating near critical points or phase transitions (Bak et al. 1988; Bak 1996; also see Kello et al. 2010; Van Rooij et al. 2013). In these systems, components flexibly and adaptively organize themselves to achieve a context-appropriate balance between independence and interdependence (Kelso 1995; Van Orden et al. 2003), which is considered to be optimal for information transmission through networks (Kello 2013). It is from this high-information, memory-laden metastability that emergent properties of cognition can be exhibited (Kello et al. 2007).

In our task, this means that constraints introduced by interaction are driving the composite memory systems further from their metastable critical points. From the Lévy foraging perspective, this might suggest that effective search behavior is lost as a system moves away from its critical point. The members of a composite search system are moving more independently than as a cohesive search system. From the area-restricted search perspective, the loss of power law behavior might reflect a loss of the ability to move flexibly and adaptively between perseverative search paradigms, indicated by short paths within a local area, and exploratory search paradigms, indicated by longer jumps to different areas. Regardless of one’s theoretical inclinations, our results suggest that collaboration brings about important, qualitative changes to the dynamics of composite search systems.

Future work will need to develop these ideas further, but we conclude the present discussion by reiterating that our findings are consistent with the notion that memory is an interaction-dominant cognitive process, and remains so even in the case of collaborative memory. Still, there are measurable differences in the interaction-dominant dynamics of independent versus collaborative memory, reflected even at the individual level. Understanding the interaction processes in the collaborative memory system, and not just the isolable components (individuals) comprising them, will be important in understanding how people remember in social contexts. Understanding what it is about collaboration that causes the shift from power law to lognormal behavior, and especially the individual differences whereby this happens for some dyads more than others, may explain why some groups just “click,” while others are “out of synch,” and may help in the designing of more cohesive group units.

Acknowledgments We would like to thank Jacqueline Pagobo and Maxine Varela for their assistance with data collection and coding, Nick Duran for his help with Praat annotations, and Drew Abney for helpful discussion.

Compliance with ethical standards

Conflict of interest The authors declare that they have no conflict of interest.

References

- Bak P (1996) How nature works. Springer, New York
- Bak P, Tang C, Wiesenfeld K (1988) Self-organized criticality. *Phys Rev A* 38(1):364–375
- Barnier AJ, Sutton J, Harris C, Wilson RA (2008) A conceptual and empirical framework for the social distribution of cognition: the case of memory. *Cogn Syst Res* 9(1–2):33–51
- Benhamou S (2007) How many animals really do the Lévy walk? *Ecology* 88:1962–1969
- Bousfield WA, Sedgewick CHW (1944) An analysis of restricted associative responses. *J Gen Psychol* 30:149–165
- Burnham KP, Anderson DR (2002) Multimodel inference: understanding AIC and BIC in model selection. *Sociol Methods Res* 33:261–304
- Dale R, Fusaroli R, Duran N, Richardson DC (2013) The self-organization of human interaction. *Psychol Learn Motiv* 59:43–95
- Farmer JD (1990) A Rosetta stone for connectionism. *Physica D* 42:153–187
- Harris C, Keil P, Sutton J, Barnier A (2010) Collaborative remembering: When can remembering with others be beneficial? In: ASCS09: Proceedings of the 9th Conference of the Australasian Society for Cognitive Science, pp 131–134
- Hills TT, Jones MN, Todd PM (2012) Optimal foraging in semantic memory. *Psychol Rev* 119(2):431–440
- Hills TT, Kalff C, Wiener JM (2013) Adaptive Lévy processes and area-restricted search in human foraging. *PLoS ONE* 8(4):e60488
- Holden JG, Van Orden GC, Turvey MT (2009) Dispersion of response times reveals cognitive dynamics. *Psychol Rev* 116(2):318–342
- Hollingshead AB (2001) Cognitive interdependence and convergent expectations in transactive memory. *J Pers Soc Psychol* 81(6):1080
- Karivava P, Odell G (1987) Swarms of predators exhibit “preytaxis” if individual predators use area-restricted search. *Am Nat* 130:233–270
- Kello CT (2013) Critical branching neural networks. *Psychol Rev* 120(1):230–254
- Kello CT, Beltz BC, Holden JG, Van Orden GC (2007) The emergent coordination of cognitive function. *J Exp Psychol Gen* 136:551–568
- Kello CT, Anderson GG, Holden JG, Van Orden GC (2008) The pervasiveness of 1/f scaling in speech reflects the metastability of cognition. *Cogn Sci* 32:741–754
- Kello CT, Brown GDA, Ferrer i Cancho R, Holden JG, Linkenkaer-Hansen K, Rhodes T, Van Orden GC (2010) Scaling laws in cognitive sciences. *Trends Cogn Sci* 14(5):223–232
- Kelso JAS (1995) Dynamic patterns: the self-organization of brain and behavior. MIT Press, Cambridge
- Landauer TK, Foltz PW, Laham D (1998) Introduction to Latent semantic analysis. *Discourse Process* 25:259–284
- Plank M, James A (2008) Optimal foraging: Lévy pattern or process? *J R Soc Interface* 5:1077–1086
- Rajaram S, Pereira-Pasarin LP (2010) Collaborative memory: cognitive research and theory. *Perspect Psychol Sci* 5(6):649–663
- Reynolds A (2010) Bridging the gulf between correlated random walks and Lévy walks: autocorrelation as a source of Lévy walk movement patterns. *J R Soc Interface* 7(53):1753–1758
- Rhodes T (2013) Recalling items from a category for 1 hour: an inquiry into power-law behavior and memory foraging. *Nonlinear Dynamics Psychol Life Sci* 17(3):361–384
- Rhodes T, Turvey MT (2007) Human memory retrieval as Lévy foraging. *Phys A* 385(1):255–260
- Schmidt RC, O’Brien B (1998) Modeling interpersonal coordination dynamics: implications for a dynamical theory of developing systems. In: Molenaar PC, Newell K (eds) Dynamics systems and development: beyond the metaphor. Erlbaum, Hillsdale, pp 221–240
- Sims DW, Southall EJ, Humphries NE, Hays GC, Bradshaw CJA, Pitchford JW, Metcalfe JD (2007) Scaling laws of marine predator search behavior. *Nature* 451:1098–1102
- Singer JD, Willett JB (2003) Applied longitudinal data analysis: modeling change and event occurrence. Oxford University Press, New York
- Stephen DG, Mirman D (2010) Interactions dominate the dynamics of visual cognition. *Cognition* 115:154–165
- Szary J, Dale R (2013) Dyadic cooperation enhances retrieval and recall of crossword solutions. In: Knauff M, Sebanz N, Pauen M, Wachsmuth I (eds) Proceedings of the 35th annual meeting of the cognitive science society. Cognitive Science Society, Austin
- Szary J, Dale R (2014) Familiarity modulates the dynamics of collaborative inhibition in a trivia game. In: Bello P, Guarini M, McShane M, Scassellati B (eds) Proceedings of the 36th annual meeting of the cognitive science society. Cognitive Science Society, Austin
- Szary J, Kello C, Dale R (2015) Memory foraging in a spatial domain. In: Noelle DC, Dale R, Warlaumont AS, Matlock T, Jennings CD, Yoshimi J, Maglio P (eds) Proceedings of the 37th annual meeting of the cognitive science society. Cognitive Science Society, Austin
- Thompson G, Kello CT (2013) Searching semantic memory as a scale-free network: evidence from category recall and a wikipedia model of semantics. In: Knauff M, Sebanz N, Pauen M, Wachsmuth I (eds) Proceedings of the 35th annual meeting of the cognitive science society. Cognitive Science Society, Austin
- Tollefsen DP (2006) From extended mind to collective mind. *Cogn Syst Res* 7(2–3):140–150
- Tollefsen DP, Dale R (2012) Naturalizing joint action: a process-based approach. *Philos Psychol* 25(3):385–407
- Troyer AK, Moscovitch M, Winocur G (1997) Clustering and switching as two components of verbal fluency: evidence from younger and older healthy adults. *Neuropsychology* 11:138–146
- Van Orden GC, Holden JG, Turvey MT (2003) Self-organization of cognitive performance. *J Exp Psychol Gen* 132(3):331–350
- Van Rooij MMJW, Nash BA, Rajaraman S, Holden JG (2013) A fractal approach to dynamic inference and distribution analysis. *Front Physiol* 4:1–16
- Viswanathan GM, Buldyrev SV, Havlin S, da Luz MG, Raposo EP, Stanley HE (1999) Optimizing the success of random searches. *Nature* 401:911–914
- Viswanathan GM, da Luz MGE, Raposo EP, Stanley HE (2011) The physics of foraging: an introduction to random searches and biological encounters. Cambridge University Press, Cambridge
- Yoshimi J (2012) Active internalism and open dynamical systems. *Philos Psychol* 25(1):1–24