

The complex system of mathematical creativity: Modularity, burstiness, and the network structure of how experts use inscriptions

Tyler Marghetis¹

tyler.marghetis@gmail.com

Kate Samson¹

knsamson1293@gmail.com

David Landy^{1,2}

dhlandy@gmail.com

¹ Indiana University, Bloomington, IN

² Netflix, Los Gatos, CA

Abstract

One of the pinnacles of human cognition is the creative insight of expert mathematics. While its concepts are abstract, the actual practice of mathematics is undeniably material and embodied. Mathematicians draw, sketch, write; having created these inscriptions, they interact with them. This iterated process of inscription is the engine of mathematical discovery. But how does this engine work? Here, using a new video corpus of mathematical experts working on proofs, and deploying tools from network and complexity science, we characterize the structure and temporal dynamics of how mathematical experts create and interact with blackboard inscriptions. We find regularities in the structure of this activity (e.g., emergent ‘communities’ of inscriptions) and its temporal dynamics (e.g., ‘bursty’ shifts in attention). By characterizing this activity, we gain a better understanding of the distributed ecosystem in which mathematical creativity occurs — including the ways that mathematicians actively construct their own notational niches.

Keywords: mathematical cognition; networks; complex systems; inscription; distributed cognition; embodiment

Introduction

One of the pinnacles of human cognition is the creative insight of expert mathematics. Often working alone, sometimes for years, mathematicians generate new knowledge about completely abstract objects, from infinite sets to imaginary numbers. The actual practice of mathematics, on the other hand, is undeniably concrete, material, and embodied. Mathematicians draw. They sketch. They write out derivations, erase them, start again. Having created these inscriptions, mathematicians interact with them: shifting their attention, talking about and gesturing at them, elaborating them further. This iterated process of inscription is the engine of mathematical discovery.

But how does this engine work? While philosophers, historians, and sociologists have argued that notations, diagrams, and the process of inscription are central to mathematical practice (Barany & MacKenzie, 2014; Mialet, 2012; Muntersbjorn, 2003), we know surprisingly little about the details of this process. Here, we use tools from network and complexity science to characterize the structure and temporal dynamics of expert mathematical activity—in particular, the process by which experts create and interact with inscriptions while working on mathematical proofs.

Notations in mathematical cognition

Past work in a range of disciplines has explored the role of notations and inscription in mathematical reasoning. Within mathematics education, for instance, it has long been recognized that choosing the right notation is often half the battle (Polya, 2004). This is true among experts just as much as it is true for schoolchildren (Muntersbjorn, 2003). Indeed, there is now a growing body of qualitative and theoretical research on the centrality of inscription in mathematical reasoning (Barany & MacKenzie, 2014; Greiffenhagen, 2014; Muntersbjorn, 2003; Roth & McGinn, 1998).

More controlled, quantitative studies have established that notations are a critical part of the distributed system of mathematical reasoning. In particular, there are bidirectional influences between, on the one hand, the specific notations used to solve mathematics problems, and, on the other, the psychological processes used to solve problems (Goldstone, Marghetis, Weitnauer, Ottmar, & Landy, 2017). Both undergraduate students and more expert reasoners, for instance, rely on the correspondence between spatial proximity and algebraic precedence in standard algebraic notation; algebraic performance is improved when this correspondence is maintained, harmed when it is violated (Landy & Goldstone, 2007). Conversely, experience with mathematical notations can reshape the psychological processes used to interact with them. Marghetis and colleagues (2016) found that, among adults who had mastered the syntax of algebra, the visual system had learned to perceive syntactically-related elements as unified visual objects. How we think about a mathematical domain shapes how we interact with inscriptions, and interacting with those inscriptions shapes how we see the problem.

Most of this past work, however, has focused on contexts where the notations are *supplied* rather than created by the participant. In real-world mathematical activity, by contrast, the reasoner must often explore multiple approaches to representing a problem—sketching out specific examples, pursuing different algebraic derivations, drawing a variety of different graphs—before settling on the final approach. Focusing only on the end product of this practice hides the dynamic messiness of mathematical reasoning. As mathematician Reuben Hersh put it, this confuses the clear, organized, pristine ‘front stage’ of

published or textbook mathematics for the messy, dynamic ‘backstage’ of real mathematical practice (Hersh, 1991).

Describing expert inscription activity

In this paper, we zoom in on this messy ‘backstage’ of mathematical reasoning, to try to characterize the dynamic contexts of mathematical creativity. To do so, we draw on tools from network science and complex systems. We describe a video corpus of mathematical experts working on non-trivial mathematical proofs. While this corpus offers endless possibilities for qualitative analysis, here we adopt a quantitative approach that allows us to measure how experts create and interact with mathematical inscriptions by identifying the ‘inscription objects’ that each expert created (e.g., equations, graphs, etc.) and then creating a timeseries of when, exactly, the expert attended to these objects, from their first creation to their final glance. We use this dense timeseries to describe the structure and dynamics of inscription.

To do so, we adapt tools from network science to offer a new methodology for studying situated cognitive activity: representing each expert’s activity as a *directed network*, in which individual inscription objects are represented as nodes, and transitions between objects (e.g., shifting attention from one graph to another) are represented as directed edges (see Methods and Figure 1). This approach is a way to ‘coarse-grain’ the messy, chalk-covered reality of expert inscription, to better reveal the deeper regularities that characterize expert notational practices.

Methods

Corpus

We created a video corpus of experts solving non-trivial mathematics problems in a naturalistic setting (total corpus length: 4 hours and 40 minutes). Doctoral students in mathematics ($N = 7$, 4 men and 3 women) were recruited through the website of the mathematics department at a major research university and compensated \$10/hour.

These experts solved up to three non-trivial problems in a natural setting: either their own office or a nearby seminar room within the mathematics department. They were encouraged to talk out loud as they solved the problems. All participants made ample use of the blackboard.

Videos were recorded with a Sony HDR-CX405 high-definition digital. The camera was positioned such that the board and the participant were visible.

Mathematics problems

Problems were drawn from the William Lowell Putnam Mathematics competition, an annual mathematics

competition for undergraduate students. These problems are typically too difficult for even advanced undergraduate students, but tractable for mathematics experts at the doctoral level or above. Problems were selected to include a range of content areas (i.e., set theory, geometry, analysis):

- (1) Find an uncountable subset, S , of the power set of a countable set, such that the intersection of each pair of elements in S is finite.
- (2) Let $f : \mathbb{R}^2 \rightarrow \mathbb{R}$ be a function such that $f(x, y) + f(y, z) + f(z, x) = 0$ for all real numbers x, y , and z . Prove that there exists a function $g : \mathbb{R} \rightarrow \mathbb{R}$ such that $f(x, y) = g(x) - g(y)$ for all real numbers x and y .
- (3) Let d_1, d_2, \dots, d_{12} be real numbers in the interval $(1, 12)$. Show that there exist distinct indices i, j, k such that d_i, d_j, d_k are the side lengths of an acute triangle.

Each participant worked for approximately an hour on the problems, depending on their availability. Most participants were only able to complete two of the problems in that time.

Video Coding

Each participant created dozens of inscriptions on the blackboard and then interacted with those inscriptions—by talking about them, gesturing towards them, or elaborating them with further inscriptions. We conducted a fine grained coding of the video corpus, at a nearly frame-by-frame resolution, to track the creation of and interaction with ‘inscription objects’ on the blackboard.

Blackboard inscriptions naturally clustered together into objects. For instance, a graph of a function might consist of two axes, labels for those axes (‘ x ’, ‘ y ’), and then a line representing the function. Each of those components, however, naturally cluster together in both meaning (they are all part of the same graph) and in spatial location (they are all located close together, with only minimal blank space in between). We used these two criteria—semantic relatedness and spatial proximity—to identify cohesive ‘inscription objects’ on the blackboard.

A coder viewed each video and annotated the onset and offset of inscription events: either the creation of a new inscription object, or subsequent interactions with that object (via talk, gaze, gesture, further elaboration, or erasing). This generated a timeseries of events for each inscription object, from its initial creation to the final time that the expert attended to it. For instance, if an expert created a graph at the very start of a session, the timeseries would include the onset and offset times for that process of initially drawing the graph; if the expert later looked at the graph, the timeseries also included that event. All coding was conducted in ELAN, software designed for annotating audio and video (Lausberg & Sloetjes, 2009).

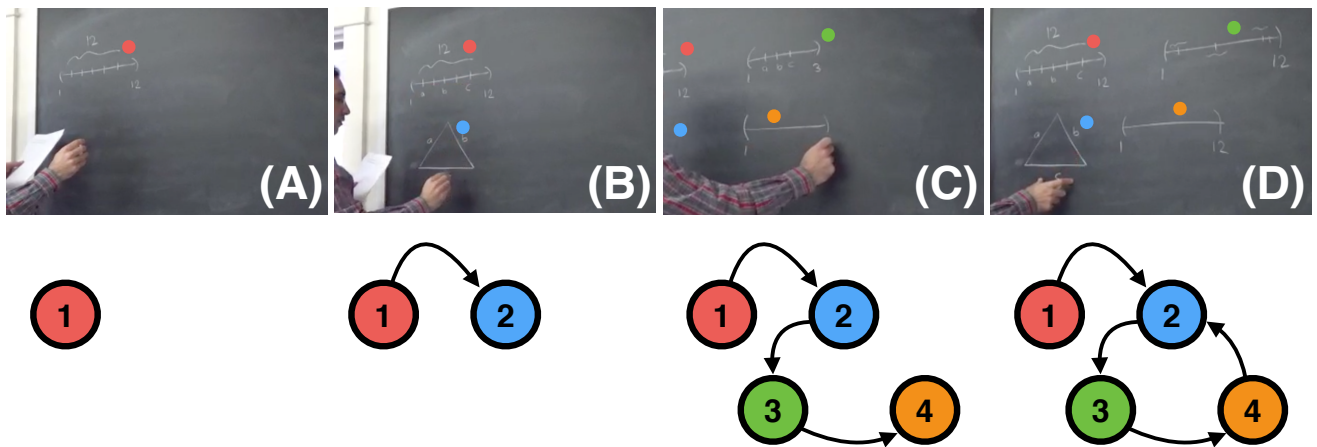


Figure 1. Illustration of a network representation of inscription activity. Blackboard images (top row) capture four consecutive stages in the process of developing a mathematical proof. The network representation of this process (bottom row) includes a node for each inscription object and an edge for transitions in attention from one object to another. We have added colored dots to the blackboard to indicate the location of each inscription object. (Node locations do not correspond to the objects' spatial locations.)

Network representation of inscription activity

To characterize the structure and temporal dynamics of experts' inscription activity, we used tools from network science. For each attempt to solve a problem, we used the timeseries of inscription events to generate a directed network, in which nodes represent inscription objects and directed edges represent transitions between objects.

As a simplified illustration, consider a scenario where an expert begins to solve a problem by creating and interacting with four inscription objects: three number-lines and one triangle (Fig. 1). The final network representation of this inscription activity would consist of four nodes, one for each inscription object, with a directed edge between two nodes whenever the expert attended first to one object and then to the other. For instance, if the expert started by creating a number-line (Fig. 1A), before abandoning that number-line to draw a triangle (Fig. 1B), the network representation of their activity up to that point would consist of two nodes — one for each inscription object — and a single directed edge from the first object to the second. As the expert creates new objects on the blackboard, the network grows (Fig. 1C, D), with their shifts in attention represented by directed edges between nodes 1 and 2, then from 2 to 3, then from 3 to 4, and then back to 2 again as they return their attention to an earlier inscription. This abstract graphical representation thus captures how the expert created and shifted their attention between inscription objects over the course of solving the problem.

Results

We first describe the network *structure* that emerged from the experts' inscription activity, then the *temporal dynamics* of their inscription activity, and finally the relations between the structure and dynamics of their notational activity.

Network structure of inscription activity

Experts created 360 distinct inscription objects, which they interacted with 4718 times. On average, solving an individual problem involved creating 24 inscription objects ($SD = 16$) and interacting with them 315 times ($SD = 191$).

Despite working on the same set of problems, experts in the corpus exhibited considerable variability in how they created and then shifted their attention among inscription objects. Figure 2, for instance, illustrates two different approaches to solving the same problem (problem #3, quoted above). For one individual (left), edges between nodes are distributed more or less randomly; nodes do not group together into interconnected clusters. By contrast (right), another individual interacted with inscriptions in interconnected clusters, and were much more likely to transition from one object to another within these clusters. This reflects a strategy where attention is likely to move to another inscription within the same cluster, creating pockets of activity wherein attention jumps between the same subset of inscriptions.

To identify these “communities” of inscriptions, we used the Girvan–Newman algorithm for community detection (Girvan & Newman, 2002), which identifies highly interconnected clusters of nodes using “edge betweenness” — the number of shortest paths between pairs of nodes that go through the edge — to identify highly central edges. In Figure 2, a node's community is identified by its color.

One way of describing the structure of inscription activity, therefore, is by how strongly shifts of attention defined communities of highly interconnected nodes — that is, the *modularity* of the network of inscription activity (Clauset, Newman, & Moore, 2004). Overall, inscription activity was significantly modular ($M = 0.18$, $t_{14} = 4.1$, $p = .001$; positive values indicate modularity, while 0 indicates

no modularity). Modularity exhibited both diversity and regularity. The participant illustrated on the left in Figure 2, for instance, generated a network with below average modularity compared to other experts who solved the same problem (modularity = 0.14), while the participant illustrated on the right had the second highest (modularity = .47). The modularity of individuals' activity varied considerably between problems (correlation in modularity between problems: $r = 0.31$); the individual who had the most modular activity on one problem, for instance, had the second-lowest modularity on another. By contrast, the two problems completed by most of the experts elicited reliably different modularity in inscription activity ($M_{\text{triangle}} = 0.26$ vs. $M_{\text{function}} = 0.08$, $t_{13} = 2.4$, $p = .03$). While modular clustering of inscriptions seems to be a recurring pattern in inscription activity, therefore, the precise amount of modularity likely reflects both the demands of the particular problem and stochastic, situated decisions.

In addition to the reliably modular structure of inscription activity overall, we also found finer grained regularities in the structure of communities themselves. Among communities, we observed two recurring 'motifs' or subgraph structures. One such motif was the 'cluster' motif, in which most nodes within a community were connected to each other (Fig. 3, right). These 'clusters' captured cases where a subset of inscription objects were all 'in conversation' with each other, with the expert shifting their attention among all inscriptions within that community. In contrast to these clusters, other 'loop' communities consisted entirely of a single, recurring route from one node, to another, to another, etc. in a straight, non-branching path (Fig. 3, left). These loops reflect inscriptions with a canonical pathway of attention—such as an algebraic derivation, where the experts attention would typically flow from the first expression to the last, in a set order.

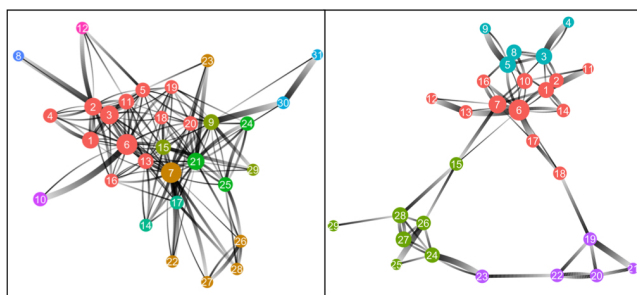


Figure 2. Different approaches to solving the same problem. Two different experts (left and right) solved the same problem, using approximately the same number of inscription objects (nodes). However, they interacted with those inscriptions in different ways, producing networks with different topological properties (see text). (Edge thickness indicates transition probabilities. Node color indicates community membership, as detected using the Girvan–Newman algorithm.)




Figure 3. Recurring motifs in network communities. (left) Multiple communities involved only a single, recurring route from one node, to another, to another, etc. in a straight, non-branching path. (right) Other communities were highly interconnected, with most nodes connected to most other nodes.

Temporal dynamics of shifts in attention

We next characterized the temporal dynamics of inscription activity. To do so, we focused on the sequence of inter-event intervals — that is, the amount of time between the onset of attention towards one inscription object and the onset of attention towards the next object. A similar approach has been used to study the temporal dynamics of other human and human-technical systems (Barabasi, 2005; Goh & Barabási, 2008), such as email wait-times and dynamics of phone calls.

Here, we use a measure that have been used previously to characterize complex systems of social and cognitive activity, and to distinguish those systems from natural (e.g., earthquakes) and autonomous physiological activity (e.g., heartbeats): the 'burstiness' of the activity (Goh & Barabási, 2008). Past work has established that human activity systems often exhibit heavy-tailed dynamics — for instance, bursts of high activity followed by long periods of inactivity, with longer periods of inactivity than expected (e.g., assuming a Poisson or Gaussian distribution). This 'burstiness' may reflect interaction-dominant dynamics, with multiple processes combining in non-additive ways, or by an underlying process that involves priority queuing (Barabasi, 2005). The burstiness of a distribution of inter-event intervals $\{t\}$ is typically measured by:

$$B = \frac{\sigma - M_t}{\sigma + M_t}$$

where σ is the standard deviation and M_t is the mean inter-event interval. More recently, this measure has been found to be sensitive to the size of finite samples, and the following elaboration has been adopted:

$$B = \frac{\sqrt{n+1}r - \sqrt{n-1}}{(\sqrt{n+1}-2)r - \sqrt{n-1}}$$

where r is the coefficient of variation, σ/M_t , and n is the sample size. Both these measures are designed to equal 0 for random, Poisson distributions; -1 for regular, periodic distributions; and +1 for bursty distributions. Past work has found that human activity systems typically exhibit significant burstiness (Goh & Barabási, 2008).

Overall, inscription activity was significantly bursty ($B = 0.17$, $p < .0001$, bootstrapped with $n = 1000$ samples). The burstiness was even more pronounced when we considered the distribution of times spent within a community of inscriptions — that is, the time between attending to one inscription object that belonged to a new community, and attending to a new object that belongs to a new community. This timeseries of inter-community dynamics was extremely bursty ($B = .47$, $p < .0001$, bootstrapped with $n = 1000$ samples), comparable to the most bursty human systems (Goh & Barabási, 2008). Inscription activity, therefore, was marked by with long periods of time spent within a community of inscriptions, followed by ‘bursty’ periods with rapid transitions between communities.

Relationship between structure and dynamics

Finally, we sought to characterize the relationship between the topological structure of experts’ inscription activity (e.g., community structure and modularity) and the temporal dynamics of that activity.

First, we examined when, exactly, experts transitioned from one community of inscriptions to another. To do so, we took our timeseries of inscription object attention and determined whether the new object of attention belonged to the same or a different community— that is, a community transition. We then tried to predict the transition to a new community, using a generalized linear mixed-effects model of whether the new object belonged to a different community. We included as fixed effects the cumulative time spent on the problem; the amount of time spent attending to the current object; and, critically, the amount of time spent in the current community since most recently beginning to attend to that community (‘sticking time’). We included random intercepts and slopes by participants, and random intercepts by problem.

There was no reliable relationship between the cumulative amount of time spent on the problem and the probability of transitioning to a new community of inscriptions ($b = 0.22 \pm 0.30$ SEM, $p = .46$). By far the strongest predictor, however, was the amount of time spent within the current community, which had a large and negative relationship to the probability of transitioning to a new community ($b = -2.29 \pm 0.40$ SEM, $p < .0001$). In other words, communities of inscriptions were themselves ‘sticky,’ so that the longer an expert spent within a community of inscriptions, the more likely they were to stay there going forward. This thus helps explain the highly

bursty dynamics of transitions between communities, reported above: experts become fascinated with a particular cluster of inscriptions and spend considerable time, before suddenly transitioning to different inscription, and perhaps then undergoing a ‘bursty’ period of rapid transition between communities.

Finally, we looked at the relationship modular structure and between bursty dynamics. We used a linear mixed-effects model of the burstiness of the inscription activity used to solve each problem, and included predictors for the problem, the total number of events, the total number of inscription objects, the mean duration of an inscription event, a measure of the ‘memory’ of the activity dynamics (Goh & Barabási, 2008), and, crucially, our measure of modularity. The predictor with the largest relationship to burstiness, and the only one that was statistically significant was modularity ($b = 1.29 \pm 0.52$ SEM, $p < .04$). More modular inscription activity—with communities of densely interconnected inscriptions—was associated with more bursty temporal dynamics (Fig. 4).

Discussion

Drawing on a corpus of mathematical experts working on non-trivial problems, and deploying tools from network and complexity science, we set out to characterize the ‘manual labor’ of mathematics (Marghetis, Edwards, & Núñez, 2014). We found that expert mathematical practice involved actively creating dozens of inscriptions and navigating between them, shifting attention from one to another. These shifts in attention were not random, however, but exhibited systematic modularity; inscriptions clustered together into ‘communities,’ subgroups of inscriptions that were likely to follow each other in a cascade of attention. This *structure* of inscription activity was related to the *temporal dynamics* of inscription, with a systematic relationship between inscription modularity and temporal burstiness (a hallmark of complex human activity). Overall, our network analysis of mathematical activity revealed both diversity and regularity in the inscription activity of experts.

The complex ‘ecosystem’ of cognition

By transforming raw video of situated problem solving into a directed network of inscription activity, we created a tractable representation of an otherwise prohibitively nuanced practice. This allowed us to adopt a quantitative approach without sacrificing a systems-level analysis. This approach shifts the focus away from individuals and skull-confined brain, and toward the ecosystem of mathematical practice, spanning brains, bodies, and blackboards.

From this perspective, the engine of mathematics is not the mathematicians’ brain, locked away inside their skull. The brain is undeniably part of that engine. But equally important is the system of notations to which the mathematician has recourse, the particular inscriptions she creates in the moment, and the way her body allows her to

bring all those parts into coordination—by looking, pointing, sketching. The mathematician’s creative insights are the product, not of solitary brains, but of a socially and materially distributed cognitive system (Hutchins, 1995). Indeed, this was true even of the physicist Stephen Hawking, who was famously confined to a wheelchair; instead of creating his own inscriptions, he worked closely with his able-bodied students, who created inscriptions on his behalf (Mialet, 2012).

This shift away from traditional intracranial processes to the larger complex system involved in creative mathematics puts a new emphasis on the material context of mathematical discovery. How should we characterize the endless inscriptions that mathematicians produce daily? How do those inscriptions change over the course of their mathematical training? How important are the inscriptions that a mathematician produces for herself, compared to those produced by her colleagues?

These questions suggest an analogy with another context of insight and learning: The early development of infants’ visual and linguistic systems. Recent work has begun to characterize the rich visual and linguistic input that is received by the developing child—including how the child actively shapes that input to facilitate learning (e.g., Smith, Jayaraman, Clerkin, & Yu, 2018). Understanding the larger ecosystem in which learning occurs—whether by a pre-verbal child or a highly trained mathematician—will be critical to understanding how, exactly, that learning occurs.

Indeed, this analogy with early child learning highlights another critical component of situated mathematical practice: Mathematicians do not receive carefully formed representations of their problems. They must figure out how to represent their ideas. In this way, the mathematician is like the child who actively shapes their visual input. Children shape their visual context to facilitate learning. Mathematicians transform their material context to facilitate creative insight. By sketching, drawing, graphing, and writing various algebraic expressions, they engage in a form of niche construction: ‘notational niche construction.’

Limitations

Our analysis has a number of limitations.

For one, naturally occurring inscriptions need not necessarily cluster into objects defined by semantic relatedness and spatial proximity. In our corpus, however, the inscriptions did typically fall into unambiguous clusters, and the few unclear cases were resolved through discussion among the authors (e.g., deciding whether a vertical stack of equations should count as one object or multiple, distinct objects). Second, one modality by which experts could engage with an object was through gaze; however, since our data consisted only of a single camera, it was not always possible to determine where a participant was looking. Gaze toward an object was only coded when there was unambiguous evidence that the participant had shifted their gaze toward an object, such as when they turned their entire

head to look at an inscription that was relatively isolated on the blackboard. As a result of this conservative approach, we may have underestimated the number of transitions between objects. To address both these issues, future work will need to establish the reliability of the coding scheme by using multiple coders and calculating inter-coder reliability.

Third, the methodology introduced here is very time-consuming, both when initially collecting the data (which requires recruiting highly trained experts) but especially when coding the video data afterwards. As a result, the current corpus consists of hundreds of inscriptions and thousands of interactions, but these were drawn from the activity of only seven experts. We are currently working to expand our corpus in order to investigate the generality of the current findings.

Future Directions and Conclusions

We have not even begun to look at how the structure and dynamics of inscription might change over the course of a problem solving episode. For instance, are there distinct phases of activity—perhaps early exploration of different inscriptions, followed by later exploitation of successful ones?

Relatedly, we have yet to investigate the association between the structure and dynamics of inscription and various other outcome measures. For instance, does the network structure of inscription activity predict the creativity or completeness of the final proof? On a more granular level, what happens immediately before the expert has a sudden insight—can we predict the onset of a critical transition in understanding (e.g., Setzler, Marghetis, & Kim, 2018; Stephen, Boncoddio, Magnuson, & Dixon, 2009)?

Third, future analyses will look in more detail at the *kinds* of inscriptions that experts are creating. Does inscription activity differ between, say, algebraic equations versus Cartesian plots? Might the structure and dynamics of inscription offer insights into how individuals tend to use different kinds of inscriptions—or perhaps reveal that superficially dissimilar inscriptions are actually treated similarly by experts?

Finally, we are curious about which aspects of inscription activity are specific to highly-trained mathematical experts, and which might also occur among novices. Work on other complex systems has found that the temporal dynamics of a complex system can predict the system’s health or resilience (Kleiger, Miller, Bigger Jr, & Moss, 1987). One possibility, for instance, is that bursty dynamics during inscription is diagnostic of mathematical expertise.

Answering these questions will bring us closer to understanding how one of the most *abstract* forms of human understanding is so undeniably *concrete*: Covered in chalk, gesturing emphatically at the blackboard, the thinking mathematician is engaged in manual labor — and then, suddenly, she understands infinity.

Acknowledgments

We are thankful for the generous feedback from members of the Landy Lab and four anonymous reviewers. This work is supported by Indiana University Bloomington through the Emerging Area of Research initiative — Learning: Brains, Machines, and Children.

References

- Barabasi, A.-L. (2005). The origin of bursts and heavy tails in human dynamics. *Nature*, 435(7039), 207.
- Barany, M. J., & MacKenzie, D. (2014). Chalk: Materials and concepts in mathematics research. *Representation in Scientific Practice Revisited*, 107–130.
- Clauset, A., Newman, M. E., & Moore, C. (2004). Finding community structure in very large networks. *Physical Review E*, 70(6), 066111.
- Girvan, M., & Newman, M. E. (2002). Community structure in social and biological networks. *Proceedings of the National Academy of Sciences*, 99(12), 7821–7826.
- Goh, K.-I., & Barabási, A.-L. (2008). Burstiness and memory in complex systems. *EPL (Europhysics Letters)*, 81(4), 48002.
- Goldstone, R. L., Marghetis, T., Weitnauer, E., Ottmar, E. R., & Landy, D. (2017). Adapting Perception, Action, and Technology for Mathematical Reasoning. *Current Directions in Psychological Science*, 26(5), 434–441.
- Greiffenhagen, C. (2014). The materiality of mathematics: Presenting mathematics at the blackboard. *The British Journal of Sociology*, 65(3), 502–528.
- Hersh, R. (1991). Mathematics has a front and a back. *Synthese*, 88(2), 127–133.
- Hutchins, E. (1995). *Cognition in the Wild*. MIT press.
- Kleiger, R. E., Miller, J. P., Bigger Jr, J. T., & Moss, A. J. (1987). Decreased heart rate variability and its association with increased mortality after acute myocardial infarction. *The American Journal of Cardiology*, 59(4), 256–262.
- Landy, D., & Goldstone, R. L. (2007). How abstract is symbolic thought? *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 33(4), 720.
- Lausberg, H., & Sloetjes, H. (2009). Coding gestural behavior with the NEUROGES-ELAN system. *Behavior Research Methods*, 41(3), 841–849.
- Marghetis, T., Edwards, L. R., & Núñez, R. (2014). More than mere handwaving: Gesture and embodiment in expert mathematical proof. *Emerging Perspectives on Gesture and Embodiment in Mathematics*, 227–246.
- Marghetis, Tyler, Landy, D., & Goldstone, R. L. (2016). Mastering algebra retrains the visual system to perceive hierarchical structure in equations. *Cognitive Research: Principles and Implications*, 1(1), 25.
- Mialet, H. (2012). *Hawking incorporated: Stephen Hawking and the anthropology of the knowing subject*. University of Chicago Press.
- Muntersbjorn, M. M. (2003). Representational innovation and mathematical ontology. *Synthese*, 134(1–2), 159–180.
- Polya, G. (2004). *How to solve it: A new aspect of mathematical method*. Princeton university press.
- Roth, W.-M., & McGinn, M. K. (1998). Inscriptions: Toward a theory of representing as social practice. *Review of Educational Research*, 68(1), 35–59.
- Setzler, M., Marghetis, T., & Kim, M. (2018). Creative leaps in musical ecosystems: early warning signals of critical transitions in professional jazz. *The 40th Annual Meeting of the Cognitive Science Society*.
- Smith, L. B., Jayaraman, S., Clerkin, E., & Yu, C. (2018). The developing infant creates a curriculum for statistical learning. *Trends in Cognitive Sciences*.
- Stephen, D. G., Boncoddio, R. A., Magnuson, J. S., & Dixon, J. A. (2009). The dynamics of insight: Mathematical discovery as a phase transition. *Memory & Cognition*, 37(8), 1132–1149.