The Limit of Active Self-Representation

Julian De Freitas¹, Lance J. Rips², George A. Alvarez¹

1- Harvard University

2- Northwestern University

Header: Limit of Self-Representation

Version: 4/17/2020

Word Count: 7260 (excluding abstract, figures, and references)

Address for correspondence:

Julian De Freitas (dfreitas@g.harvard.edu)

Department of Psychology, Harvard University

33 Kirkland Street

Cambridge, MA 02139

Abstract (142 words)

A cornerstone of cognitive science is that mental systems are limited: There is a maximum amount of information they can process or store, beyond which performance breaks down. Yet so far the study of such limits has been focused on core systems like attention and memory. Here we explore self-representation, the ability to represent someone or something as being *you*. We find a limit on a well-known index of self-representation known as the self-reference effect, whereby people perform better for items that they associate with themselves than with others. This effect does not scale additively when people actively associate stimuli with more than one self-concept at a time (e.g., their younger and older selves, or two hypothetical versions of themselves). This suggests that even though we can self-represent various self-concepts independently, we can never do so for more than one self-concept simultaneously.

Keywords: self-relevance, self-representation, self-concepts, capacity limits

A distinctive feature of human intelligence is that we are not bound by our immediate sensory experience, but can also travel within our own minds. We can imagine past and future versions of ourselves, hypothetical alternative futures (e.g., what if I train for a marathon, or apply for that job?), different “parts” of ourselves (e.g., who am I at work versus at home?), and even duplicated selves, as popularized in movies such as *The Matrix* and *Harry Potter* (Bartels & Rips, 2010; Buckner & Carroll, 2007; De Freitas, Cikara, Grossmann, & Schlegel, 2017; McConnell, 2011; Newman & Knobe, 2019; Oyserman & James, 2011; Parfit, 1984).

This commonplace human experience has long been appreciated by psychologists dating back to William James (Higgins, 1987; James, 1890; Kurzban & Athena Aktipis, 2007; Markus & Nurius, 1986; Newman, De Freitas, & Knobe, 2015; Roberts & Donahue, 1994). For instance, even if we choose to limit relevant research to the study of future selves, we see references to possible selves (Markus & Nurius, 1986), ideal and ought selves (Havighurst, Robinson, & Dorr, 1946), desired and undesired selves (Ogilvie, 1987), fantasy selves (Bybee, Luthar, Zigler, & Merisca, 1997), and selves associated with goals (Austin & Vancouver, 1996) and personal projects (Little, 1987). This preponderance of selves has led some researchers to propose frameworks that abandon the notion of a single, unitary self-construct in favor of a dynamic concept that includes different kinds of self-concepts that are continuously created, revised, or discarded (McConnell, 2011; Oyserman & James, 2011). For instance, under the Multiple Self-Aspects Framework (MSF), the self is instantiated as an associative network, wherein each node represents a different self-concept (e.g., as a girlfriend, as a sorority member) with associated features (e.g., female, shy, giving) that may partially overlap other self-concepts (e.g., “persistent” may belong to both an “athlete” and a “student” self). Under this model, different situations may activate different self “contexts” or sub-networks, leading to spreading activation through the network depending on the inter-relations of features (McConnell, 2011).

Yet although we can easily imagine multiple selves and parts of those selves, we only seem to identify with one of those selves at one time; you intuitively feel that only one self, current or imagined, is *you*. This suggests that people may allocate a particular self-concept to be represented as *the* self at any one time, rather than simultaneously activating multiple self-concepts. We will interchangeably refer to this process as ‘self-representation’ or ‘self-allocation’. The current paper explores whether people can allocate more than one self-representation at a time (aka Multiple Self-Allocation), or only one at a time (Single Self-Allocation).

The Case for Single Self-Allocation

Although people can identify with several self-concepts, they appear to switch the particular concept for representation at any given point. Empirical work suggests that different self-concepts may compete for representation, with some being activated and others suppressed. For instance, if a person has two self-concepts whose features conflict with each other (e.g., “sorority member” and “Northwestern student”), then priming one of these concepts leads them to respond quicker in a lexical decision task to words related to that concept and slower to words related to the opposing concept, suggesting that the opposing concept is actively inhibited (Hugenberg & Bodenhausen, 2004). Further, how difficult it is to switch from one self-concept to another depends on its relative importance to the person. Specifically, it takes longer to switch from an important self-concept to a less important one than vice versa, suggesting that the important concept is more “sticky” (Hugenberg & Bodenhausen, 2004).

Rationally speaking, it also seems plausible that identifying with multiple selves would cause confusion and inner conflict. After all, the most dramatic examples of multiple identities come from split brain patients (Gazzaniga, Bogen, & Sperry, 1962) or patients with dissociative identity disorder (APA, 2013). For instance, careful experimentation with split brain patients reveals that they are

effectively inhabited by two independent selves, and yet these patients tend to otherwise present as single coherent persons (Marinsek, Gazzaniga, & Miller, 2016). Evidently, part of the brain's left hemisphere actively resolves conflict and uncertainty between the hemifields, maintaining a sense of unity and control. Perhaps the very existence of this mechanism suggests that reconciling two selves is too confusing and even computationally costly in human brains.

The Case for Multiple Self-Allocation

Yet many of the same studies also hint at the possibility of multiple self-allocation. For one, even though self-concepts sometimes conflict, they often have features in common. One study asked participants to list features associated with their different self-concepts, and found that 42% of the traits were associated with more than one concept. On average, at least one trait was common to 69% of their self-concepts (McConnell, 2011). The existence of such overlapping features raises the possibility that more than one self-concept can be allocated as *the* self at once, perhaps due to spreading activation of features associated with several concepts (McConnell, Rydell, & Brown, 2009). For instance, some have speculated that those with greater working memory capacity might be capable of simultaneously activating two self-concepts (Brown, Bailey, Stoll, & McConnell, 2016), yet no work has explored such exciting possibilities.

The notion of multiple self-allocation also seems plausible when we consider the history of capacity limits in cognition more generally, as in the study of attention and memory. The approach of these studies has been to unveil the true max number of objects that people can track or remember at once, by increasing the difficulty of the tasks until performance plateaus (Franconeri, Alvarez, & Cavanagh, 2013). Notably, some of these studies have uncovered situations in which people are capable of tracking or remembering more items than was originally thought possible (Franconeri et al., 2013). For instance, people can remember more items if each to-be-remembered

item contains less information (e.g., simple colors rather than Chinese characters) (Alvarez & Cavanagh, 2004), they can track more objects if the items are evenly spaced out on the display rather than clumped together (Alvarez & Franconeri, 2007), and they can both track and remember more items if the items span the visual hemifields rather than occur within a single visual hemifield (Alvarez & Cavanagh, 2005; Strong & Alvarez, 2019). Inspired by these studies, one wonders whether there are also circumstances in which people can represent multiple selves concurrently.

Rationally speaking, multiple self-allocation could also be behaviorally useful, as when multiple self-concepts are relevant to one's goals. When socializing in a group that consists of both sorority members and other Northwestern students, one could allocate a corresponding self-representation to monitor each respective conversation stream, and update knowledge about each self-concept in parallel. The ability to allocate multiple self-representations would also enable one to achieve various mental feats at once, such as planning out several courses of motor action, or reflecting and learning from several past experiences. An enviable manifestation of this idea is portrayed in the movie *Her*, in which the protagonist's AI lover is simultaneously in conversation with 8,316 people even as she attends to his needs. Although human minds are far less nimble, perhaps they are capable of a feat similar in kind.

A Capacity Limit Approach and the Self-Reference Effect

And so it is that existing work has hinted at both the Single and Multiple Self-allocation theories, but not decidedly adjudicated between them. Partly this is because existing models devoted to explaining multiple self-concepts have been largely based on studies involving self-reported characteristics, and their few behavioral measures (e.g., lexical decisions tasks, or priming) have not selectively targeted simultaneous self-referential processing per se. Here we seek a targeted behavioral measure that directly tests the possibility of multiple self-allocation, while taking a

capacity limits approach, i.e., we want to push people to try and allocate multiple self-representations simultaneously, while determining the point at which performance plateaus.

We propose to take this approach to the phenomenon of self-reference effects in associative learning, wherein people perform better if they associate items with the self than with others, e.g., a familiar other (e.g., a friend, or mother) or a stranger (Rogers, Kuiper, & Kirker, 1977). For instance, a typical self-reference task may ask participants to associate their current self, a friend, and a stranger with simple geometric shapes, such as a triangle, circle, or square. (The reason for using simple shapes is that participants are unlikely to have pre-existing associations between arbitrary shapes and people.) After learning these associations, participants are then serially shown each shape for just 100ms, paired with either its correct label or an incorrect label, and they must rapidly (within < 1s) indicate whether the pairing is a match or mismatch (Sui, He, & Humphreys, 2012). An important feature of this paradigm is that perfect performance requires consistently identifying whether a given shape-label pair is a match, regardless of which person/agent the shape represents. Therefore, if participants happen to perform better for the shape that represents the self (as they do), then this suggests a self-reference effect, given the arbitrary pairing between shapes and people.

Finally, we note that self-reference effects are thought to be due to processing that is specific to associating an item with a self-representation, rather than to encoding of items in just any semantically meaningful way. For instance, one study measured both self-reference and semantic encoding effects within the same paradigm, and found that self-reference effects were greater on several measures (Leshikar, Dulas, & Duarte, 2015). Another study found that a patient suffering from semantic impairments and severe amnesia exhibited a self-reference effect but not a semantic encoding effect (Sui & Humphreys, 2013). In sum, associating an item with the self-representation may confer a self-specific boost in performance.

Arbitrating Between Single and Multiple Self-Allocation

Presumably, participants in previous studies performed best for self-referring labels like ‘you’ because this label referred to whichever self-concept happened to be active in their minds at the time of testing. The current studies will extend this logic by asking people to associate items with various kinds of self-concepts at once — such as their younger self and older self — to test whether the self-reference effect grows as the number of distinct self-concepts associated with stimuli within the paradigm grows, yielding an overall performance boost when there are multiple self-concepts.

If performance does not grow as the number of self-concepts does, then this suggests that there is only one self-representation that can be activated at a time (Single Self-Allocation; Figure 1A). On this picture, incoming perceptual information competes for the attention of a single self-representation/processor (dashed lines). That self-representation then uses the input exclusively — i.e., representing either only Stimulus 1 (e.g., young you) or only Stimulus 2 (e.g., old you) — conferring performance benefits selectively to that stimulus when making perceptual categorizations (green lines). Or, it might serially switch between representing each of the stimuli, conferring some performance benefit to each separate stimulus, but not as much as it would to any single stimulus if it were fully dedicated to representing just that one stimulus.

Alternatively, if performance does grow as the number of self-concepts does, then this suggests that people can spawn multiple self-representations at the same time, each of which can be associated with a different stimulus and can issue its own independent performance benefits (Multiple Self-Allocation; Figure 1B).

Figure 1. Two models of self-allocation: **A.** Single Self-Allocation, in which stimuli compete (dotted lines) for a single self-representation, but only one stimulus is assigned to the self and prioritized during processing (green line). **B.** Multiple Self-Allocation, in which each stimulus can be assigned to one of many self-representations spawned in parallel, leading all stimuli to be prioritized during processing.

To our knowledge, the only related work looking at self-reference effects and capacity limits tested whether there is a redundancy effect for self-related processing, i.e., whether people exhibit an

enhanced self-reference effect when they associate more than one shape with the exact *same* self label (Sui, Yankouskaya, & Humphreys, 2015). Using mathematical modeling of their behavioral results, the authors found that participants exhibited super capacity performance when multiple stimuli were redundantly associated with their same, present (rather than imagined) self. Those results suggest that associating multiple stimuli with the same self-representation binds those items to that single self-representation, conferring performance benefits to both stimuli.

We want to underscore two ways in which our work both differs from and builds on this previous work. First, those authors did not test performance when participants were asked to associate multiple stimuli with multiple *different* self-concepts, so as to test the possibility of multiple self-allocation. In our case, we took deliberate steps to prevent different stimuli from being bound by a single self-representation, including (i) providing a brief explanation of each separate self-concept label (rather than only including one self label), (ii) associating each of the distinct shapes with separate self-concept labels (rather than redundantly associating them with one self-label), and (iii) presenting shape-label pairs separately in serial (rather than presenting multiple shapes at once paired with a single self label).

Second, the fact that the authors found a performance increase for redundant self-related shapes suggests that the standard self-reference effect typically observed for a single self-concept may not necessarily constitute a ceiling effect, but leaves ‘wiggle room’ to observe a further performance boost. That is, just as the previous work found that it is possible to boost overall performance via redundant self-representation, i.e., by binding multiple stimuli to one self-representation (Sui et al., 2015), here we test whether it is possible to boost overall performance via the hypothesized mechanism of multiple self-allocation, i.e., by allocating separate self-representations to each self-relevant stimulus.

The Current Studies

To adjudicate between the Single and Multiple Self-Allocation views, we employ the associative learning paradigm from previous work, reviewed above (Figure 2) (Sui et al., 2012). The speeded categorization component of this paradigm is suited for our purposes because it requires holding multiple associations active in working memory. As such, we expected that the only way to find a self-reference effect that scaled with the number of selves would be if people were capable of Multiple Self-Allocation. The work has two main aims, to test: (i) whether self-reference effects also extend to imagined selves, rather than just to the present self; and (ii) whether participants perform better overall when the task involves associating stimuli with multiple self-concepts versus just one. Experiments 1-3 test this question for different kinds of self-concepts, including alternative selves (Experiment 1), young and old selves (Experiment 2), and true and surface selves (Experiment 3). Experiments 4-6 then explore the potential role of self-distinctiveness, by conducting a high-powered replication of the experiment that featured the most distinct selves (Experiment 4), and then experimentally making the contrasted self-concepts more distinct (Experiments 5-6).

General Procedure Across Studies

Stimuli were created using a combination of Javascript, html, and css. Participants were recruited from the online crowdsourcing platform, Amazon's mechanical Turk, to take part in a 20-min study for \$2.00 compensation. Such participants are more diverse than typical lab samples, yet provide responses that are equally reliable (Buhrmester, Kwang, & Gosling, 2011). A database kept track of their Turk Ids, to ensure that no participant took part in more than one of the experiments. All studies were approved by the Harvard Internal Review Board.

In all studies except our high-powered replication (Experiment 4), we recruited 60 participants per condition, similar to a previous study which found that forty participants provides

sufficient power to detect a self-reference effect using a simpler version of the paradigm (De Freitas & Alvarez, 2019). We heuristically added 20 participants per condition, since our own paradigm additionally featured a condition with two self-concepts, and we wanted to ensure sufficient statistical power to compare this new condition to the conditions with only one self-concept. Our high-powered replication (Experiment 4) recruited 100 participants per condition (roughly double that of the other experiments).

At the beginning of all experiments, participants answered an attention check (*How many fatal heart attacks have you had?*), and a comprehension question about the images (*Which, if any, of the following images did you observe in the sequence of images you saw? A. Two cars crashing, B. A Building on fire, C. An ambulance, D. Different shapes, E. None of the above.*). Experiments 5 and 6, which manipulated self-distinctiveness, also included a comprehension question about which particular self-concepts participants were asked to imagine; for instance, *Who were you asked to imagine? A. old-you, an enemy named Bob, and a foreigner named Tom. B. old-you, a stranger named John, and a stranger named Bill, C. young-you, old-you, and a stranger named John.* The options were modified based on the experiment. Finally, participants answered basic demographics questions: age, gender, native language, and level of science and philosophy education.

We excluded participants for incorrectly answering any of the attention or comprehension checks. We also excluded them for scoring lower than 55% accuracy overall (50% = chance), or for reacting quicker than 100ms (holding down the key or jumping the gun) for more than 50% of trials. These stringent performance-based checks meant that we excluded more participants than typical self-report studies, which only have attention and comprehension checks. But we stuck with these planned exclusions in order to ensure the highest data quality.

All experiments included three conditions: The Self #1 condition asked participants to imagine one particular self and two strangers (a stranger named John and another named Bill), the

Self #2 asked participants to imagine another particular self and the same two strangers, and the Two Selves condition asked participants to imagine both of the selves featured in each of the Self #1 and Self #2 conditions in addition to one stranger. Thus, each condition asked participants to associate three agents with three shapes, while varying how many of those agents were selves. Since there were always three labels associated with three shapes, this gave 6 possible associations per each of three self conditions, yielding 18 unique combinations total; we counterbalanced these combinations between-subjects. Remaining details of the paradigm are summarized in the introduction and in Figure 2.

Figure 2. Task structure of associative learning paradigm, using the self-concepts from Experiment 2 as an example.

Experiments 1-3: Testing for Multiple Self-Allocation

Experiments 1-3 tested whether people exhibit an enhanced self-reference effect when they have to associate items with two different self-concepts rather than one. The different experiments contrasted different self-concepts, which we explain below together with the specific participant numbers recruited in each experiment.

Experiment 1: Alternative Selves. We recruited 179 participants, and excluded 107 for failing any of the strict attention, comprehension, and performance checks, yielding a final sample of 72 participants ($M_{\text{age}} = 34$, 40% female).

The Self #1 condition asked participants to imagine one future self (*Imagine an activity that you could do in 2 hours from now*, labeled as *future-you1*) and two strangers (a stranger named John and a stranger named Bill, labeled *stranger-john* and *stranger-bill*), and to associate each of these labels with a separate shape (Figure 2). The Self #2 condition asked participants to imagine two self-concepts (*...now imagine a different activity that you could do instead in 2 hours from now*, labeled *future-you2*), but to only associate the second label with a shape, such that Future-you1 played no role in the rest of the procedure. And, as for Self #1, participants were also asked to associate the same two strangers with the remaining shapes. Finally, the Two Selves condition asked participants to associate future-you1, future-you2, and stranger-john with their respective shapes.

Experiment 2: Past and Future Selves. We recruited 179 participants, and excluded 98 for failing any of the checks, yielding a final sample of 81 participants ($M_{\text{age}} = 35$, 53% female).

Participants imagined a past self (*the younger you 20 years ago*) and a future self (*the older you 20 years from now*).

Experiment 3: True and Surface Selves. We recruited 180 participants, and excluded 113 for failing any checks, yielding a final sample of 67 participants ($M_{\text{age}} = 34$, 52% female). Participants

imagined their true self (*who you truly are deep down*) and their surface self (*those things that you picked up from society rather than who you truly are deep down*).

Results, Experiments 1-3

Self-reference effects. We combined performance levels for the matching and mismatching conditions of the associative learning task, in order to form an unbiased measure of overall accuracy, d' (Green & Swets, 1966). We replicated the self-reference effect in each of the Self #1 and Self #2 conditions, finding in both cases that participants performed better for the self-concept than for either of the strangers (Figure 3). To our knowledge, this is the first demonstration of a self-reference effect for various self-concepts/ imagined selves other than just one's present self. Other work has tested such extensions but failed to observe self-reference effects, for reasons we discuss in the general discussion (Golubickis et al., 2017; Kim, Stolte, & Humphreys, 2019).

Second, we found that participants in the Two Selves condition performed indistinguishably for the two self-concepts. So, did they also perform better overall in the Two Selves condition than in the Self #1 and Self #2 conditions? To answer this question, we tested whether the summed performance across each of the agent categories was higher in the Two Selves condition than the Self #1 and Self #2 conditions. It was not: Across Experiments 1-3, people performed no better for two selves than one, even though they exhibited a self-reference effect in each of the Self #1 and Self #2 conditions (see the gray bars in Figure 3). These results support the Single Self-Allocation view.

Figure 3. Performance on the associative learning task in Experiments 1-3. The gray bars represent summed performance across each of the agent categories for the Self #1, Self #2, and Two Selves conditions, from left to right. Error bars depict standard error of the mean. *ns* = not-significant, $\dagger < .10$, $* p < .05$, $** p < .01$, $*** p < .001$.

If participants in the Two Selves condition were only able to allocate one self-representation at a time, then this representation might have either (i) selectively represented only one of the self-concepts, or (ii) serially switched between representing each of the self-concepts, leading to some confusions, given the challenging nature of the task. To explore these possibilities, we also compared the sizes of the self-reference effects in the Two Selves and Self #1 and Self #2 conditions. For each participant in the Two Selves condition, we calculated the difference between the *best* performing self and stranger-John, i.e., $[(\max(\text{self1}, \text{self2}) - \text{stranger-John})]$, then compared this difference to that between the self and stranger-John in each of the Self #1 and Self #2 conditions, i.e., the standard self-reference effects. If only (i) occurred, then we should see no difference between the Two Selves and Self #1 and Self #2 conditions, since in all conditions there should only be a standard self-reference effect for the single self-concept that was being represented as the self. If (ii) occurred, then we should see signs of worse performance in the Two Selves condition than in the Self #1 and Self #2 conditions, since serial switching between self-concepts should lead to some confusions, which in turn manifests as an overall weakened self-reference effect for the best performing self-concept.

Again, we found no evidence for a performance boost in the Two Selves condition relative to the Self #1 and Self #2 conditions. Additionally, we found that numerical performance in the single self conditions (Self #1 and Self #2) was actually higher than in the Two Selves condition for 6/6 of the comparisons (as shown in Table 1), with several comparisons trending toward significance. These results are consistent with the idea that participants in the Two Selves conditions were to some extent switching between representing the different self-concepts as the self.

Table 1. Difference in the size of the self-reference effect for the Two Selves vs. Self #1 and Self #2 conditions, i.e., [(Two Selves: max(self1, self2) – stranger-John) - (Self #1: self 1 – stranger-John)] and [(Two Selves: max(self1, self2) – stranger-John) - (Self #2: self 2 – stranger-John)].

	<i>M (SD) vs. M (SD)</i>	<i>t</i>	<i>d</i>
E1: Alternate future selves			
Self #1 vs Two Selves	0.91 (1.05) v 0.45 (0.59)	1.87 [†]	0.52
Self #2 vs Two Selves	0.75 (0.81) v 0.45 (0.59)	1.37	0.42
E2: Past and future selves			
Self #1 vs Two Selves	0.64 (1.16) v 0.25 (0.59)	1.69 [†]	0.42
Self #2 vs Two Selves	0.38 (0.62) v 0.25 (0.59)	0.75	0.23
E3: True & surface selves			
Self #1 vs Two Selves	0.80 (0.75) v 0.37 (0.72)	1.88 [†]	0.58
Self #2 vs Two Selves	0.58 (0.74) v 0.37 (0.72)	0.95	0.29
E4: Powered replication of E6			
Self #1 vs Two Selves	0.84 (0.74) v 0.68 (0.81)	0.77	0.21
Self #2 vs Two Selves	0.43 (0.75) v 0.68 (0.81)	-1.31	-0.33
E5: Distinct past & future selves			
Self #1 vs Two Selves	0.99 (0.79) v 0.23 (0.77)	3.35**	0.96
Self #2 vs Two Selves	0.73 (0.77) v 0.23 (0.77)	2.27*	0.65
E6: Distinct alternate selves			
Self #1 vs Two Selves	0.85 (0.68) v 0.50 (0.60)	1.84 [†]	0.54
Self #2 vs Two Selves	0.84 (0.73) v 0.50 (0.60)	1.61	0.52

Confusion patterns. If such representation switching occurred, we would also expect to see greater confusions between the two self-concepts than between either self-concept and the stranger, since the self-representation should specifically target the two self-concepts. To test this possibility, we looked at the subset of mismatched trials in the Two Selves condition, asking whether participants false alarmed more when the mismatched label referred to the other self than to the stranger, e.g., were participants more likely to false alarm when the shape represented *healthy-you*, but the label read *unhealthy-you* rather than *stranger-john*?

Table 2 summarizes the mean difference in the proportion of false alarms for these two cases of mismatched labels (label refers to other self – label refers to stranger), doing so for each of the self-concepts presented in the Two Selves Condition. As indicated by the positive differences in Table 2, we found greater confusion between the two selves than between either of these two selves

and the stranger. Further, the size of these effects did not tend to differ depending on which of the particular selves the shape represented— except in Experiment 3, for which the confusion between selves vs. strangers was greater when the shape represented the true self than the surface self. This confusion between selves further suggests that people were incapable of Multiple Self-Allocation, but had to switch between representing each of the self-concepts with a single self-representation.

At the same time, these results raise the possibility that Multiple Self-Allocation might be tenable when the self-concepts are more distinct, making them less confusable. Experiments 4-6 explore this possibility directly.¹

Table 2. Were participants in the associative learning experiments more likely to false alarm if the incorrect label referred to the other self than to the stranger?

	Mean diff. in false alarms for other self - stranger		Was difference of these differences significant?
	Self 1	Self 2	Self 1 – Self 2
E1: Alternate future selves	.08 (.10)**	.10 (.09)***	future1 = future2
E2: Past & future selves	.09 (.06)***	.10 (.08)***	young = old
E3: True & surface selves	.15 (.12)***	.07 (.11)**	true > surface**
E4: Powered replication of E6	.09 (.11)***	.05 (.10)*	young = old [†]
E5: Distinct past & future selves	.07 (.09)***	.08 (.11)**	young = old
E6: Distinct alternate selves	.05 (.10)*	.06 (.06)***	healthy = unhealthy

Note: Columns 1 and 2 present the difference in the proportion of false alarms when the label that accompanied a self-related shape referred to the other self-concept versus a stranger, for each of the two self-concepts used in the Two Selves condition. Column 3 indicates whether the difference between the first two columns is significant. All significance symbols reflect paired t-test comparisons. [†] < .01, * p < .05, ** p < .01, *** p < .001.

¹ It seems unlikely that the lack of a difference between self-concepts with more similar labels (“future-you1” “future-you2”) can simply be written off due to the similarity of the labels themselves, since (to preview our results) we continued to find no evidence for Multiple Self-Allocation even when the self-concepts were directly manipulated to be maximally distinct.

Experiments 4-6: The Role of Self Distinctiveness

Given the confusion results presented in Table 2, we wondered if it is possible to facilitate Multiple Self-Allocation by making the paired self-concepts more distinct. We suspected — and later confirmed in a norming study that we describe further below — that the most distinct pair of self-concepts across Experiments 1-3 were the young vs. old self. So, we started by running a high-powered replication of this study (Experiment 4), to ensure that the previous lack of evidence for Multiple Self-Allocation was not due to underpowered samples. We then increased the chances of observing Multiple Self-Allocation by experimentally making some of the self-concepts from the previous experiments more distinct (Experiments 5-6). Here we summarize the main differences across experiments, together with their recruited participant numbers.

Experiment 4: High-powered Replication of Experiment 2. We recruited 298 participants (100 per condition, roughly double that of the previous experiments), and excluded 198 for failing any of the checks, yielding a final sample of 100 participants ($M_{\text{age}} = 35$, 35% female). Aside from the larger number of participants, the design was the same as Experiment 2.

Experiment 5: Distinct Past & Future Selves. We recruited 184 participants, and excluded 107 for failing any of the checks, yielding a final sample of 77 participants ($M_{\text{age}} = 33$, 42% female). As in Experiment 2, participants were asked to imagine younger and older selves, but this time the selves were more distinct in age (18 year old and 60 year old, rather than a 20-year younger or older self). We present norming results below verifying that people did indeed rate these selves as numerically more distinct.

Experiment 6: Distinct Alternate Selves. We recruited 183 participants, and excluded 119 for failing any of the checks, yielding a final sample of 64 participants ($M_{\text{age}} = 33$, 35% female). As in Experiment 1, participants were asked to imagine alternative selves, but this time the selves had more distinct preferences: one self was healthy-minded (*imagine a time in your life in which you have been*

very health conscious, wanting to eat only healthy foods), while the other was unhealthy-minded (*Now imagine a time in your life in which you have been much less health conscious, eating whatever you want*). Although this manipulation varies several factors other than just distinctiveness, our main goal was to make the selves as distinct as possible.

Results, Experiments 4-6

Self-reference effects. Across experiments, we did not find any evidence that self-distinctiveness influenced how participants performed for two selves vs. one (Figure 3). Our high-powered Experiment 4 replicated results in Experiments 1-3, and making the selves maximally distinct seemed to have no effect either. Across all studies, there was no overall performance boost for two selves versus one.

As in Experiments 1-3, we also tested for Multiple Self-Allocation by comparing the size of the self-reference effect between the Two Selves condition and each of the Self #1 and Self #2 conditions (Table 1). Once again, there was no evidence for a performance boost in the Two Selves condition. If anything, performance in the single self conditions (Self #1 and Self #2) was again numerically higher than that in the Two Selves condition for 5/6 of the comparisons. Finally, in the Two Selves condition we again saw higher levels of confusion between the two self-concepts than between either of those self-concepts and the stranger (Table 2).

Fig. 4. Performance on the associative learning tasks in Experiments 4-6. The gray bars represent summed performance across each of the agent categories for the Self #1, Self #2, and Two Selves conditions, from left to right. Error bars depict standard error of the mean. *ns* = not-significant, $\dagger < .01$, $* p < .05$, $** p < .01$, $*** p < .001$.

Distinctiveness Analysis. Justifying our high-powered test of the distinctiveness hypothesis in Experiment 4, a separate norming experiment (reported in the Supplemental Material) confirmed

that the most distinct pair of self-concepts across Experiments 1-3 was indeed the young and old selves. The norming study also confirmed that our manipulations of self-distinctiveness in Experiments 5-6 were effective in eliciting higher impressions of distinctiveness. Finally, we tested whether the degree of self-distinctiveness across all six experiments predicted the extent to which participants performed better for two selves vs. one. Consistent with the Single Self-Allocation view, it did not ($r = -0.04, p = .938$). Evidently, even though two self-concepts were more confusable with each other than with a stranger (as indicated by the analyses of Experiments 1-3), making those self-concepts more distinct was not enough to elicit a performance benefit for two selves vs. one. This may be because only one self-concept can ever be activated at a time.

General Discussion

How many self-concepts can we represent as *the* self or selves at a time? We determined the limit on a well-known index of self-representation known as the self-reference effect, whereby people are better at processing information related to the self than others (Rogers et al., 1977). we found that this effect occurred not only for the present self, as demonstrated in previous work, but also for various imagined self-concepts, including past and future selves, alternative selves, and true and surface selves. Yet the effect did not grow additively as the number of self-concepts grew from one to two, suggesting that the mind can only exhibit a self-reference effect for one actively represented self-concept at a time. In line with this view, people tended to confuse multiple self-concepts with each other more than they did either of those self-concepts with a stranger, and making the two self-concepts more distinct did not amplify their performance either. Participants never performed better for two selves than one, even though they always exhibited a self-reference effect for each of the self-concepts when presented in isolation.

Self-reference effects for various imagined self-concepts

The current work is the first to demonstrate that self-reference effects do not only occur for one's present self, but also for various other self-concepts with which one may personally identify. Two recent studies explored whether self-reference effects extend to future and past selves, but only found a self-reference effect for participants' present selves (Golubickis et al., 2017; Kim et al., 2019). We suspect that this is because they contrasted the present self with more distant imagined selves, whereas the current studies contrasted distant selves with strangers. Crucially, as we found evidence for across multiple studies, people may only be able to allocate one self representation at a time. So, the fact that participants in previous studies identified exclusively with the present self may simply indicate that the present self was the most salient of the presented options, but not that they could not identify with more distant, imagined self-concepts at all. So long as an imagined self is paired against non-self contenders, such as strangers, it should win out for representation, yielding the usual self-reference effect.

Relatedly, the current results indicate that imagined selves (e.g., alternative selves) can influence basic associative learning and categorization mechanisms, rather than only operate within high-level cognition and language. Apparently, what gets represented as the self in perception and memory can be flexibly designated by our moment-to-moment thoughts about ourselves. Future work should explore the generality of such effects, as well as their boundary conditions. For instance, we expect that self-reference effects might also extend to owned objects, which may be viewed as extensions of the self (DeScioli, Karpoff, & De Freitas, 2017; DeScioli & Wilson, 2011; Nancekivell, Friedman, & Gelman, 2018). Another open question is whether self-reference effects only occur for agents that people were or could personally be, or whether they also occur for agents that people have imagined being, such as when taking the perspective of the first person narrator of a novel, or character in a video game.

By the same token, the current results build on existing frameworks devoted to explaining the existence of multiple self-concepts and their mental management, by underscoring a distinction between the notion of various self-concepts and the more limited process of self-allocation. This previous work has speculated about whether people can activate multiple self-concepts simultaneously. In the current studies we deliberately tried to push participants to achieve this feat, but to no avail. This suggests that previous evidence suggesting co-activation of multiple self-concepts may be picking up on lower-level neural activations that precede the behaviorally-consequential allocation of one of these self-concepts as *the* self.

What do these limits tell us about the self?

Are the current effects specific to the self? Although performance benefits also occur for certain other stimuli, such as those that are rewarding (Sui & Humphreys, 2015) or familiar (Sui, Sun, Peng, & Humphreys, 2014), self-related performance benefits are consistently larger than these other effects (Leshikar et al., 2015), are not merely due to deep elaborative/semantic encoding (Sui & Humphreys, 2013), and appear to arise from a distinctive neural network that is implicated in various other self-referential tasks too (Sui & Gu, 2017; Sui, Rotshtein, & Humphreys, 2013). So, even if the self-reference paradigm taps into general properties of cognition, the specific behavioral effects of associating stimuli with the self are thought to reflect self-specific processing..

A related question is whether people performed worse in the Two Selves condition for the general reason that the two self-concepts were more similar to each other, and thus more confusable. This view is supported by the finding that participants were more likely to confuse these self-concepts with each other than with a stranger (although they did not perform better for the stranger overall). Yet it is important to note that previous work finds that participants perform especially well if the two self-related stimuli are even more *similar* — two shapes that redundantly refer to the very same self-label (Sui et al., 2015). In those studies, participants were instructed to

associate multiple shapes with a single self-label. Then, during the test trials, they were shown either one or both of the shapes presented with a single self label or friend label. It is thought that this procedure led to redundant binding of the shapes to a single self-representation, producing the overall performance boost (Sui et al., 2015). That finding is consistent with previous work showing more efficient processing of redundant vs. isolated stimuli (Miller, 1982), such as when emotion and face identity information are presented as parts of a single face configuration (Yankouskaya, Booth, & Humphreys, 2012) or when co-occurring objects are treated as a single gestalt (Mordkoff & Yantis, 1991).

In contrast to this previous work, our paradigm deliberately used different labels for different self-concepts, with each shape-label pair presented in serial, so as to ensure that participants had to discriminate between the two self-labels throughout the task. Under such challenging task constraints, we expected that the only way for the usual self-reference effect to grow with the number of self-concepts would be if participants could allocate multiple self-representations in parallel— each devoted to actively representing its own self-concept. If participants were capable of such active multiple self-allocation, then the added processing resources should have minimized confusions between the self-concepts, ensuring better performance overall. The fact that confusions did occur suggests that no such additional resources were available, but that participants had only a single self-representation to work with.

With all that said, our results cannot rule out the possibility of simultaneous self-representation altogether. Future work could explore whether multiple self-allocation might be elicited depending on other factors, including (i) cultural differences in *independent versus interdependent self-orientation* (Geertz, 1975; Grossmann & Na, 2014); (ii) individual differences in *self-complexity* (Linville, 1987; Rafaeli-Mor, Gotlib, & Revelle, 1999), a composite measure of the number of selves one identifies with and the features they share in common; (iii) individual difference in *relative self-*

importance, the relative importance one places on different self-concepts (Brown et al., 2016; Sanz-Vergel, Demerouti, Bakker, & Moreno-Jiménez, 2011); and (iv) *intersectional identities*, whether one typically manages multiple self-concepts that are especially salient, as is common in transgender individuals and black women (De Vries, 2012; Reid & Comas-Díaz, 1990). Future work could also try splitting the self-relevance task across the visual hemifields, given that performance on other high-level tasks such as attentive tracking improves when items are divided across the hemifields, suggesting that each brain hemisphere has separate cognitive resources to devote to the items in its corresponding hemifield (Alvarez & Cavanagh, 2005; Sperry, Gazzaniga, & Bogen, 1969). Perhaps there are also hemifield-specific mechanisms for self-representation that can be leveraged to represent multiple self-concepts at once.

Finally, we believe that the current findings provide a promising case study of how searching for limits on higher-level thought can illuminate cognitive architecture. In the current case, this search enabled us to (1) isolate a new limited system, i.e., that for self-reference; (2) characterize its limits, finding that the mind cannot allocate more than one self-representation concurrently; (3) make conceptual progress, illuminating a psychological distinction between being able to imagine multiple self-concepts and allocating one of them as *the* self; (4) characterize its flexibility, i.e., people exhibit self-reference effects for various imagined self-concepts, not just for the present self; and (5) find evidence pertaining to whether the self-reference effect is parallel versus serial, finding evidence for seriality (although future studies can also employ other methods to target this specific question; Townsend, 1990).

Origins of the self-reference limit

The limit on the self-reference effect is different from previous findings of capacity limits on attention and memory in that it is especially severe. Why? Existing neuroscientific work on the self-reference effect finds that it relies on a large neural network linking brain regions associated with

self-representation (the ventral medial prefrontal cortex) and sensory-driven regions associated with social attention (the left posterior superior temporal sulcus) (Sui et al., 2013). Therefore, one speculation is that the energy costs of replicating such a large network are too high to outweigh the benefits that would be conferred from doing so, leading items to compete for representation by this single network.

Similar competitive effects occur for attention and memory, but here they may arise because items compete for representation on space-limited cortical real-estate, which permits holding active at least a handful of items at a time, depending on their spatial representation on cortex (Franconeri et al., 2013). Thus, what attention, memory and the self-reference effects appear to have in common is that items must compete for representation, yet the mechanisms that give rise to competition may differ for the self compared to attention and memory, explaining the different limits.

A second related, and perhaps more philosophical, possibility is that the severe self-reference limit reflects a deliberate design feature of cognition. It may be advantageous to represent only one self-concept as *the* self at a time, maybe even regardless of the mind's capacity to push beyond this limit. In principle, the empirical limit we found is consistent with either of the above possibilities.

Conclusion

Humans are better at processing items that refer to themselves than others. We found that this effect occurs not only for our present selves but also for various imagined self-concepts. Further, this self-advantage did not grow as the number of imagined selves grew, suggesting that the capacity limit of self-referential encoding is equal to a puny size of one. Evidently, although we can travel within our own minds in many forms, each trip has a very strict occupancy limit.

Acknowledgments

For helpful feedback, thanks to Nathan Couch and Tomer Ullman.

References

- Alvarez, G. A., & Cavanagh, P. (2004). The capacity of visual short-term memory is set both by visual information load and by number of objects. *Psychological science*, 15(2), 106-111.
- Alvarez, G. A., & Cavanagh, P. (2005). Independent resources for attentional tracking in the left and right visual hemifields. *Psychological science*, 16(8), 637-643.
- Alvarez, G. A., & Franconeri, S. L. (2007). How many objects can you track?: Evidence for a resource-limited attentive tracking mechanism. *Journal of vision*, 7(13), 14-14.
- APA. (2013). *Diagnostic and statistical manual of mental disorders (DSM-5®)*: American Psychiatric Pub.
- Austin, J. T., & Vancouver, J. B. (1996). Goal constructs in psychology: Structure, process, and content. *Psychological bulletin*, 120(3), 338-375.
- Bartels, D. M., & Rips, L. J. (2010). Psychological connectedness and intertemporal choice. *Journal of Experimental Psychology: General*, 139(1), 49-69.
- Brown, C. M., Bailey, V. S., Stoll, H., & McConnell, A. R. (2016). Between two selves: Comparing global and local predictors of speed of switching between self-aspects. *Self and Identity*, 15(1), 72-89.
- Buckner, R. L., & Carroll, D. C. (2007). Self-projection and the brain. *Trends in cognitive sciences*, 11(2), 49-57.
- Buhrmester, M., Kwang, T., & Gosling, S. D. (2011). Amazon's Mechanical Turk: A new source of inexpensive, yet high-quality, data? *Perspectives on psychological science*, 6(1), 3-5.
- Bybee, J., Luthar, S., Zigler, E., & Merisca, R. (1997). The fantasy, ideal, and ought selves: Content, relationships to mental health, and functions. *Social Cognition*, 15(1), 37-53.
- De Freitas, J., & Alvarez, G. (2019). Personal identity without mind. *PsyArXiv*.
- De Freitas, J., Cikara, M., Grossmann, I., & Schlegel, R. (2017). Origins of the belief in good true selves. *Trends in cognitive sciences*, 21(9), 634-636.

- De Vries, K. M. (2012). Intersectional identities and conceptions of the self: The experience of transgender people. *Symbolic Interaction, 35*(1), 49-67.
- DeScioli, P., Karpoff, R., & De Freitas, J. (2017). Ownership dilemmas: The case of finders versus landowners. *Cognitive science, 41*, 502-522.
- DeScioli, P., & Wilson, B. J. (2011). The territorial foundations of human property. *Evolution and Human Behavior, 32*(5), 297-304.
- Franconeri, S. L., Alvarez, G. A., & Cavanagh, P. (2013). Flexible cognitive resources: competitive content maps for attention and memory. *Trends in cognitive sciences, 17*(3), 134-141.
- Gazzaniga, M. S., Bogen, J. E., & Sperry, R. W. (1962). Some functional effects of sectioning the cerebral commissures in man. *Proceedings of the National Academy of Sciences, 48*(10), 1765-1769.
- Geertz, C. (1975). On the nature of anthropological understanding: Not extraordinary empathy but readily observable symbolic forms enable the anthropologist to grasp the unarticulated concepts that inform the lives and cultures of other peoples. *American Scientist, 63*(1), 47-53.
- Golubickis, M., Falben, J. K., Sahraie, A., Visokomogilski, A., Cunningham, W. A., Sui, J., & Macrae, C. N. (2017). Self-prioritization and perceptual matching: The effects of temporal construal. *Memory & Cognition, 45*(7), 1223-1239.
- Green, D. M., & Swets, J. A. (1966). *Signal detection theory and psychophysics*. New York: Wiley.
- Grossmann, I., & Na, J. (2014). Research in culture and psychology: Past lessons and future challenges. *Wiley Interdisciplinary Reviews: Cognitive Science, 5*(1), 1-14.
- Havighurst, R. J., Robinson, M. Z. R., & Dorr, M. (1946). The development of the ideal self in childhood and adolescence. *The Journal of Educational Research, 40*(4), 241-257.
- Higgins, E. T. (1987). Self-discrepancy: a theory relating self and affect. *Psychological review, 94*(3), 319-340.

- Hugenberg, K., & Bodenhausen, G. V. (2004). Category membership moderates the inhibition of social identities. *Journal of Experimental Social Psychology*, 40(2), 233-238.
- James, W. (1890). *Principals of psychology* (2 Vols.). New York: Holt.
- Kim, H., Stolte, M., & Humphreys, G. W. (2019). The relations between temporal and social perceptual biases: Evidence from perceptual matching. *Attention, Perception, & Psychophysics*, 81(3), 599-606.
- Kurzban, R., & Athena Aktipis, C. (2007). Modularity and the social mind: Are psychologists too self-ish? *Personality and Social Psychology Review*, 11(2), 131-149.
- Leshikar, E. D., Dulas, M. R., & Duarte, A. (2015). Self-referencing enhances recollection in both young and older adults. *Aging, Neuropsychology, and Cognition*, 22(4), 388-412.
- Linville, P. W. (1987). Self-complexity as a cognitive buffer against stress-related illness and depression. *Journal of Personality and Social Psychology*, 52(4), 663-667.
- Little, B. R. (1987). Personal projects and fuzzy selves: Aspects of self-identity in adolescence. In T. Honess & K. Yardley (Eds.), *Self and identity: Perspectives across the life span* (pp. 230-245). New York: Routledge & Kegan Paul.
- Marinsek, N. L., Gazzaniga, M. S., & Miller, M. B. (2016). Split-brain, split-mind. In *The Neurology of Conciousness* (pp. 271-279): Elsevier.
- Markus, H., & Nurius, P. (1986). Possible selves. *American Psychologist*, 41(9), 954-969.
- McConnell, A. R. (2011). The multiple self-aspects framework: Self-concept representation and its implications. *Personality and Social Psychology Review*, 15(1), 3-27.
- McConnell, A. R., Rydell, R. J., & Brown, C. M. (2009). On the experience of self-relevant feedback: How self-concept organization influences affective responses and self-evaluations. *Journal of Experimental Social Psychology*, 45(4), 695-707.

- Miller, J. (1982). Divided attention: Evidence for coactivation with redundant signals. *Cognitive psychology*, 14(2), 247-279.
- Mordkoff, J. T., & Yantis, S. (1991). An interactive race model of divided attention. *Journal of Experimental Psychology: Human Perception and Performance*, 17(2), 520.
- Nancekivell, S. E., Friedman, O., & Gelman, S. A. (2018). Ownership matters: People possess a naïve theory of ownership. *Trends in cognitive sciences*, 23(2), 102-113.
- Newman, G. E., De Freitas, J., & Knobe, J. (2015). Beliefs about the true self explain asymmetries based on moral judgment. *Cognitive Science*, 39(1), 96-125.
- Newman, G. E., & Knobe, J. (2019). The essence of essentialism. *Mind & Language*, 34(5).
doi:10.1111/mila.12226
- Ogilvie, D. M. (1987). The undesired self: A neglected variable in personality research. *Journal of personality and social psychology*, 52(2), 379-385.
- Oyserman, D., & James, L. (2011). Possible identities. In *Handbook of identity theory and research* (pp. 117-145). New York: Springer.
- Parfit, D. (1984). *Reasons and persons*. Oxford: Oxford University Press.
- Rafaeli-Mor, E., Gotlib, I. H., & Revelle, W. (1999). The meaning and measurement of self-complexity. *Personality and Individual Differences*, 27(2), 341-356.
- Reid, P. T., & Comas-Diaz, L. (1990). Gender and ethnicity: Perspectives on dual status. *Sex Roles*, 22(7-8), 397-408.
- Roberts, B. W., & Donahue, E. M. (1994). One personality, multiple selves: Integrating personality and social roles. *Journal of Personality*, 62(2), 199-218.
- Rogers, T. B., Kuiper, N. A., & Kirker, W. S. (1977). Self-reference and the encoding of personal information. *Journal of personality and social psychology*, 35(9), 677.

- Sanz-Vergel, A. I., Demerouti, E., Bakker, A. B., & Moreno-Jiménez, B. (2011). Daily detachment from work and home: The moderating effect of role salience. *Human Relations*, 64(6), 775-799.
- Sperry, R. W., Gazzaniga, M. S., & Bogen, J. E. (1969). Interhemispheric relationships: the neocortical commissures; syndromes of hemisphere disconnection. In P. Vinken & G. Bruyn (Eds.), *Handbook of Clinical Neurology* (pp. 273-290). Amsterdam: North-Holland Publishing.
- Strong, R. W., & Alvarez, G. (2019). Hemifield-specific control of spatial attention and working memory: Evidence from hemifield crossover costs.
- Sui, J., & Gu, X. (2017). Self as object: Emerging trends in self research. *Trends in Neurosciences*, 40(11), 643-653.
- Sui, J., He, X., & Humphreys, G. W. (2012). Perceptual effects of social salience: evidence from self-prioritization effects on perceptual matching. *Journal of Experimental Psychology: Human Perception and Performance*, 38(5), 1105.
- Sui, J., & Humphreys, G. W. (2013). Self-referential processing is distinct from semantic elaboration: Evidence from long-term memory effects in a patient with amnesia and semantic impairments. *Neuropsychologia*, 51(13), 2663-2673.
- Sui, J., & Humphreys, G. W. (2015). The interaction between self-bias and reward: Evidence for common and distinct processes. *Quarterly Journal of Experimental Psychology*, 68(10), 1952-1964.
- Sui, J., Rotshtein, P., & Humphreys, G. W. (2013). Coupling social attention to the self forms a network for personal significance. *Proceedings of the National Academy of Sciences*, 110(19), 7607-7612.
- Sui, J., Sun, Y., Peng, K., & Humphreys, G. W. (2014). The automatic and the expected self: Separating self-and familiarity biases effects by manipulating stimulus probability. *Attention, Perception, & Psychophysics*, 76(4), 1176-1184.

Sui, J., Yankouskaya, A., & Humphreys, G. W. (2015). Super-capacity me! Super-capacity and violations of race independence for self-but not for reward-associated stimuli. *Journal of Experimental Psychology: Human Perception and Performance*, 41(2), 441.

Yankouskaya, A., Booth, D. A., & Humphreys, G. (2012). Interactions between facial emotion and identity in face processing: Evidence based on redundancy gains. *Attention, Perception, & Psychophysics*, 74(8), 1692-1711.