The Advantages of Combining the Simultaneous–Sequential Paradigm with Systems Factorial Technology

15

Piers D.L. Howe, *Adam Ferguson*
University of Melbourne, Melbourne, VIC, Australia

Both the simultaneous–sequential paradigm (Eriksen & Spencer, 1969; Shiffrin & Gardner, 1972) and systems factorial technology (Townsend, 1992; Townsend & Nozawa, 1995) attempt to identify the type of processing underlying a particular mental act. However, these competing methods attempt to do this utilizing different behavior measures and by making different assumptions. In this sense, the two methods are complementary with the strengths of one being able to compensate for the weaknesses of the other and vice versa. We argue that this sometimes makes it advantageous to utilize both methods because then if they both arrive at the same conclusion this would amount to very strong evidence in favor of that conclusion.

Systems factorial technology typically measures the reaction time of an observer to two events that occur simultaneously. This allows direct inferences about the processing of the two events to be made (Townsend, 1992; Townsend & Nozawa, 1995). Conversely, the simultaneous–sequential paradigm typically divides a given task into two halves and compares the observers' performance (usually quantified by percentage correct on the task) when the two halves are performed sequentially as opposed to when they are performed simultaneously. In effect it measures the observer's processing capacity. From this the practitioners attempt to infer the underlying mental processing, specifically whether the processing is likely to be serial or parallel (Eriksen & Spencer, 1969; Shiffrin & Gardner, 1972). The advantage of the simultaneous–sequential paradigm is that it is usually relatively easy to adapt it to a given experimental paradigm. Its disadvantage is that it is more susceptible to parallel–serial mimicry since different mental architectures can in principle have the same capacity (Townsend, 1990). While a finding of unlimited capacity is indicative, though by no means proof, of parallel processing (Townsend, 1990), a finding of limited capacity is much harder to interpret as it is readily predicted by both parallel and serial models. Conversely, systems factorial technology suffers less from model mimicry but it can be harder to apply it in practice. For both methods it is easiest to describe them with reference to a particular task. In this chapter we will illustrate their use by using them to determine whether identity–location bindings are maintained and updated by a serial, parallel or coactive process.

The binding problem is central to visual perception. It is the problem of distinguishing the properties of an object from the other objects in a visual scene and associating together only those properties that correspond to the same object (Milner, 1974;

Systems Factorial Technology, DOI: 10.1016/B978-0-12-804315-8.00018-5

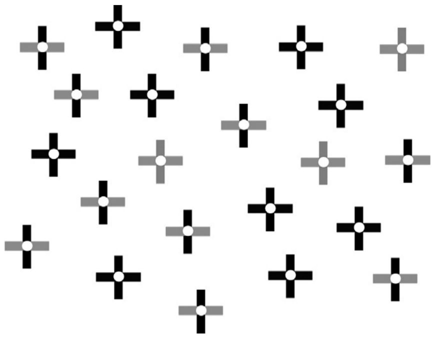

Figure 15.1 Before you attend to the display you are probably aware that it contains a large number of plus signs, some of them black or with black parts and others gray or with gray parts. However, if you were asked to search for a particular plus sign, say one with a gray vertical and a black horizontal bar, you would need to attend to the items, probably in a sequential fashion.

von der Malsburg, 1981; von der Malsburg & Schneider, 1986; Treisman, 1996). For example, in Fig. 15.1 you are probably aware of a number of plus signs and that some of these plus signs are black or have black parts and others are gray or have gray parts. However, before you attend to any of the objects, you are probably not aware whether any of the plus signs comprise a gray vertical bar and a black horizontal bar. If you search for this particular target, you will probably find yourself doing so by attending to each item in turn (Treisman & Gelade, 1980; Wolfe, Cave, & Franzel, 1989). In this display, only by attending to an item can you determine whether it is the target of your search. Attention is needed to bind item's features together so as to allow you to determine what it is (Treisman & Gelade, 1980). Without binding you can be aware of the individual features present in the scene (Treisman & Gelade, 1980; Wolfe et al., 1989) but not whether a particular combination of these features is present (Milner, 1974; von der Malsburg, 1981; von der Malsburg & Schneider, 1986; Treisman, 1996). For this reason, the binding problem strongly constrains visual awareness (Crick & Koch, 1990; Howe et al., 2009; Wolfe, 2012).

Although often not recognized, the binding problem takes a number of different forms (Treisman, 1996; Di Lollo, 2012). The one most commonly considered is the feature binding problem, which is the problem of determining which features belong to which object (Treisman, 1996). The above illustration is an example of the feature binding problem. Conversely, the identity–location binding problem is the problem of associating the correct object identity with the correct location. In other words, it refers to knowing what is where in the visual scene (Treisman, 1996). This binding problem is typically studied using objects whose identities are known but whose locations are continuously changing. In this chapter we will consider only this type of binding problem.

In our experiments, we studied the identity–location binding problem using the multiple identity tracking (MIT) task (Oksama & Hyona, 2004; Horowitz et al., 2007). In the MIT task, the observers are shown a display containing a number of moving objects. In our case, these objects were differently colored disks. A typical trial starts

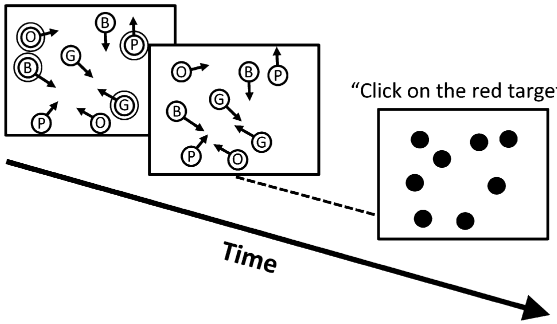

Figure 15.2 At the start of the trial the observer is shown a display containing a number of colored disks (blue, green, purple and orange denoted by the letters "B", "G", "P" and "O", respectively) a subset of which are briefly ringed to indicate that these are the targets to be tracked. The rings then disappear and the disks move around. At a random point in time, all the disks stop moving, turn black and the observer is asked to indicate the location of a particular target, for example, the red one.

with a subset of these disks being briefly ringed to indicate that these are the targets to be tracked during the trial. The rings then disappear and all the objects move around the screen in a pseudo-random fashion. At the end of the trial the objects all stop moving and become the same color. The observer is then asked to indicate the location of a particular target, for example, the red one (Fig. 15.2).

It is important to emphasize that the duration of the trial is random so the observer does not know in advance when the trial is likely to end. This forces the observer to continuously keep track of where each target is located. This ensures that the observer must continuously solve the identity–location binding problem for each target.

A potential concern with this task is that the observers may be tempted to make eye movements. Eye movements are commonly observed when observers freely view a tracking display (Fehd & Seiffert, 2008; Zelinsky & Neider, 2008; Fehd & Seiffert, 2010; Huff, Papenmeier, Jahn, & Hess, 2010). These eye movements necessarily introduce a serial component into the task. Thus, even if the mental processing underlying tracking was itself parallel, it might appear to be serial because of the serial component introduced by the eye movements. To avoid this potential confound, in all our experiments observers were required to maintain fixation on a fixation cross. We know that they were successful in doing this because we did not observe any seriality in their tracking.

Simultaneous–Sequential Paradigm

In our first two experiments we utilized the simultaneous–sequential paradigm. To do this, in each experiment we introduced two conditions. In both conditions, the disks did not move continuously. Instead each disk moved and paused for half the trial. Thus the motion discontinuities were exactly the same in the two conditions and in each condition each disk traveled exactly the same distance. The difference between the two conditions was whether the disks moved and paused synchronously. In the *simultaneous* condition, all the disks moved and paused together. Conversely, in the *sequential* condition, half the disks moved while the other half paused and then the

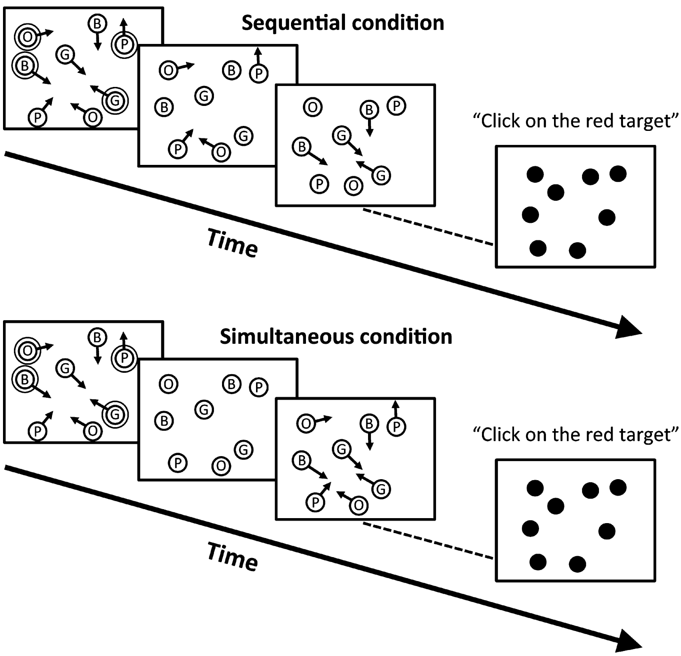

Figure 15.3 In both conditions, each disk moves and pauses for exactly half the trial. In the sequential condition, at any one time half the disks are moving. In the simultaneous condition, all the disks move and pause together.

first half paused while the second half moved. This sequence repeated throughout the trial so that at any one time only half the disks were moving (Fig. 15.3).

To understand the utility of this paradigm, it worth considering what predictions a parallel model with independent channels and a serial model with independent stages would make. Because the channels are independent, the parallel model by definition has unlimited capacity. Conversely, the capacity of the serial model is assumed to be limited. In the literature these are often referred to as a standard parallel model and a standard serial model (Townsend & Asby, 1983). A standard parallel model assumes that all targets are tracked continuously and independently, regardless of their state of motion. Thus it predicts that the maximum speed at which the targets can be tracked at should be the same in both the simultaneous and sequential conditions. Conversely, a standard serial model makes a very different prediction, as discussed below.

The model of multiple identity tracking (MOMIT) is currently the only model of MIT and is an example of a standard serial model in that it assumes that the targets are attended independently, one at a time in a sequential fashion (Oksama & Hyona, 2008). Every time a target is attended, the observer notes the target's location. When it is time to reattend a given target (i.e., after attending to all the other targets in turn), the model assumes that whichever object is closest to the last remembered position for that target is assumed by the observer to be that target. If this assumption turns out not to be true, for example, if the observer has taken too long to reattend a particular target or if the targets are moving too quickly, then a tracking error will be made. Crucially, the model assumes that the observer will attend preferentially to the targets whose locations are least certain, which in this case would be the moving targets. Assuming that there are four targets in total, it follows that in the simultaneous condition the effective tracking load, i.e., the number of disks to which the observer needs to attend at any given time, will sometimes be four disks because sometimes all four disks will be moving at the same time so will need to be attended. Conversely, in the sequential condition only two of the disks will be moving at any one time. Thus the effective tracking load will be approximately two disks in the sequential condition. According to the serial model, halving the effective tracking load should halve the time it takes the observer to reattend any given target and so halve the distance traveled by each target during the time it takes the observer to reattend it. This model predicts that the probability of losing a target is determined by the distance the target travels in the time it takes the observer to reattend it. It follows that, to hold the probability of making a tracking error constant, when the effective tracking load is halved the speed at which the targets move needs to be doubled so that the total distance traveled by each target in the time it takes for the target to be reattended is held constant. Consequently, the maximum speed at which the targets can be tracked is predicted to be approximately double in the sequential condition than in the simultaneous condition.

The serial and parallel models therefore make different predictions with regards the maximum speed at which the targets can be tracked in the two conditions. Thus, by measuring this maximum speed we can determine which model is more likely to be correct. However, this analysis rests on the assumption that if tracking occurs serially then the tracking mechanism would prioritize the moving targets. This assumption is reasonable for at least three reasons. First, it is known that motion onsets preferentially attract attention to moving targets, thereby presumably facilitating their tracking (Abrams & Christ, 2003). Second, there is both fMRI (Howe, Horowitz, Morocz, Wolfe, & Livingstone, 2009) and EEG (Drew, Horowitz, Wolfe, & Vogel, 2011) evidence that in tracking attention is indeed directed preferentially to the moving targets as opposed to the stationary ones. Finally, it has been shown that during tracking observers attend preferentially to the targets that are most in danger of being lost, which in the case of Experiments 1 and 2 would be the moving ones (Iordanescu, Grabowecky, & Suzuki, 2009). The above supporting evidence notwithstanding, we could come to a stronger conclusion if we did not have to make this assumption. For this reason we also investigated the same issue using systems factorial technology since we were able to do this in a way that avoided making the above assumption. Specifically, when we used systems factorial technology, all disks moved continuously, with the consequence that we did not have to make any assumptions regarding the difference in processing of moving versus stationary disks. Given that the two techniques make different assumptions, if they both arrive at the same conclusion then this would be strong evidence in favor of the veracity of that conclusion.

Systems Factorial Technology

As systems factorial technology has been discussed extensively elsewhere (see the tutorial in this book, Altieri, Fifić, Little, & Yang, 2017) in this chapter we will provide only a brief overview of how this method applies to our particular paradigm, referring the reader to the original sources for the mathematical proofs (Townsend & Nozawa, 1995; Townsend & Wenger, 2004). In our study, observers were presented with an MIT stimulus similar to the experiment outlined above except that all the disks moved continuously, during the tracking phase all the disks were the same gray color and there were only two targets. At a random point during the movement phase the two targets darkened by a small amount (S), by a large amount (L) or not at all (N). If they both darkened, they always did so simultaneously. For the situation where both targets did darken, there were three possible conditions. They could both darken by a large amount (the LL condition), one could darken by a large amount and the other by a small amount (the mixed condition) or both might darken only by a small amount (the SS condition). Note that because both targets were free to move, we do not distinguish between the SL and the LS condition. Instead, we refer to both conditions jointly as the mixed condition. The observer was requested to push a response button as quickly as possible if either target darkened. The greater the darkening, the faster it is likely to be detected, so the shorter the reaction time is likely to be. By measuring the distributions of the reaction times it is possible to determine the type of processing, for example, whether the targets are processed in series, in parallel or coactively (Townsend & Nozawa, 1995). The most straightforward way to do this is to measure the mean reaction time for the three conditions and use this to calculate the mean interaction contrast (*MIC*) given by the following formula (Townsend & Wenger, 2004):

$$MIC = RT_{\mathrm{LL}}(t) + RT_{\mathrm{SS}}(t) - 2 \times RT_{mixed}(t) \tag{15.1}$$

and analogously for the *SIC*:

$$SIC = S_{\mathrm{LL}}(t) + S_{\mathrm{SS}}(t) - 2 \times S_{mixed}(t) \tag{15.2}$$

where the survivor function is the probability that on a trial where the observer will respond, he has not yet done so by time t.

In our experiment, the observers were instructed to respond when either target darkened. Thus, if the observers notice that one of the targets has darkened they do not need to check whether the other target has darkened before responding. In other words, the observers can perform the task by applying a self-terminating stopping rule (Townsend & Nozawa, 1995). However, despite this, the observers may in practice always check both targets, even though this is not always necessary, in which case the observers would be said to be applying an exhaustive stopping rule (Townsend & Nozawa, 1995). By calculating the *SIC*, we can determine which stopping rule the observers employ.

It is also useful to quantify the capacity of the system to process multiple items simultaneously. In the present MIT task, this can be done by measuring the change in processing speed when the observers are asked to process two stimuli as opposed to just one stimulus. Detailed capacity predictions are presented in the tutorial chapter.

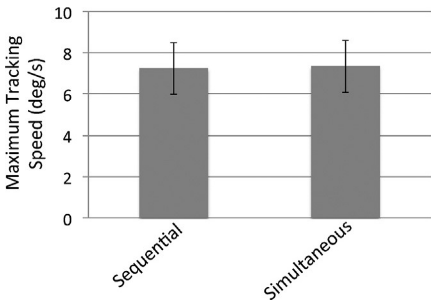

Figure 15.4 The results from the first experiment. There is no difference in the maximum tracking speed in the sequential and simultaneous conditions. Error bars represent within-subject standard error of the means (Cousineau, 2005; Morey, 2008).

Experiments

Below we will briefly summarize three experiments that show how the simultaneous–sequential paradigm can be profitably combined with systems factorial technology to mitigate the potentially limiting assumptions of either method. For further details on these experiments, the reader is referred to the original article (Howe & Ferguson, 2015).

Experiment 1: Simultaneous–Sequential Paradigm with Targets Whose Unique Identities Were Continuously Visible

This experiment utilized the paradigm outline above and illustrated in Figs. 15.2 and 15.3. Fifteen observers were asked to keep track of four targets among eight disks while maintaining fixation on a fixation cross. To prevent the observers from verbally reciting the color names of the targets, the observers were required to repeatedly say the word "animal" during the tracking phase of each trial. At the end of each trial two disks were highlighted one at a time, and the observers were asked whether each disk in turn was a target. In simultaneous and sequential conditions, during the movement phase, each disk moved only half the time. In the simultaneous condition all the disks moved and paused simultaneously. Conversely, in the sequential condition, only half the disks moved at any one time. The QUEST staircase routine (Watson & Pelli, 1983) was used to find the maximum tracking speed at which the observers could track the disks at 75% tracking accuracy. In other words, the speed at which on 75% of the trials the observer was able to correctly answer whether the two disks sequentially highlighted at the end of the trial were targets or distractors. Observers were not provided feedback on their performance.

As explained previously, the standard serial model predicts that the maximum tracking speed in the sequential condition should be double that of the simultaneous condition. Conversely, the standard parallel model predicts that the maximum tracking speed should be the same in the two conditions. The results are shown in Fig. 15.4. Consistent with the prediction of the parallel model, a within-subjects t test found no

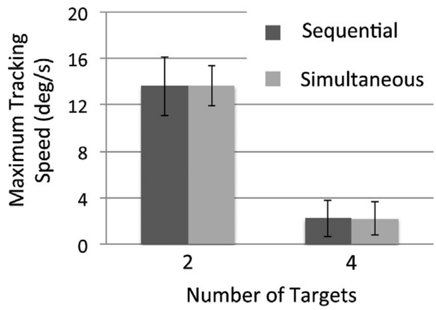

Figure 15.5 The results from the second experiment. As before, there is no difference in the maximum tracking speed in the sequential and simultaneous conditions. Error bars represent within-subject standard error of the means (Cousineau, 2005; Morey, 2008).

difference between the maximum tracking speeds of the two conditions, $t(14) = 0.71$, $p = 0.49$, $Cohen's\ d = 0.22$.

Experiment 2: Simultaneous–Sequential Paradigm with Targets Whose Unique Identities Are not Continuously Visible

In most previous studies of MIT, the identities of the objects were continuously visible during the tracking phase (Horowitz et al., 2007; Oksama & Hyona, 2008; Pinto, Howe, Cohen, & Horowitz, 2010; Cohen, Pinto, Howe, & Horowitz 2011). The advantage of this procedure is that it mimics real life where the identities of objects are also typically visible during tracking. However, a potential disadvantage is that having the identities of the object continuously visible makes it possible for observers to sometimes recover from tracking errors. Thus, the performance of the observers on the tracking task reflects not just their tracking proficiency but also their ability to recover from errors. Because we wished to have a pure measure of the ability to track objects and their distinct identities, in Experiment 2, objects were initially presented in different colors but became identical during the tracking phase (Pylyshyn, 2004; Hudson, Howe, & Little, 2012), thereby making target recovery impossible. This ensured that the performance of observers on this task reflected only their tracking ability. It is important to emphasize that this experiment still required observers to continuously refresh the identity–location bindings of the targets since if they failed to do that they would not know where the targets were located at the end of the trial. Thus, their tracking performance was a true measure of their ability to solve the identity–location binding problem.

After running the previous experiment it occurred to us that observers may employ different strategies when tracking different numbers of targets. For example, they might have the capacity to track two targets in a parallel fashion but are forced to track four targets in a serial fashion. To investigate this possibility, in Experiment 2 we sometimes had observers track two targets and sometimes four targets. Other than these differences, the experiment was conducted the same as previously.

The results are shown in Fig. 15.5. A two-way ANOVA found a significant main effect for the number of targets, with the maximum tracking speed being significantly greater in the two target condition than in the four target condition, $F(1, 14) = 75.1$, $p < 0.001$, *partial* $\eta^2 = 0.86$. However, no main effect was found for type of motion with the maximum tracking speed not being significantly different between the sequential and simultaneous conditions, $F(1, 14) = 0.01$, $p = 0.93$. No significant interaction was observed, $F(1, 14) = 0.02$, $p = 0.90$. This experiment therefore supports our previous finding that performance is equivalent in the simultaneous and sequential conditions. This result is consistent with the standard parallel model but not with the standard serial model.

Experiment 3: Using Systems Factorial Technology

While the previous two experiments provide strong evidence against serial accounts of tracking, they both rely on the assumption that if tracking is serial then the tracking mechanism would prioritize the tracking of the moving targets thereby reducing the effective tracking load in the sequential condition. While this assumption is clearly plausible, we could be more confident in our conclusions if we could reach them without making this assumption. That is why in Experiment 3 we used systems factorial technology as this method does not make this assumption.

Systems factorial technology requires one to measure the reaction time to a discrete event, such as the appearance of two stimuli (Townsend & Nozawa, 1995). However, tracking is a temporally extended process, so it is not obvious how one can directly measure tracking performance in terms of a reaction time. Instead, in Experiment 3, we assumed that tracking is mediated by attention and used systems factorial technology to measure whether the two targets were attended simultaneously or sequentially by measuring the response time to the darkening of the two targets. This in turn allowed us to infer whether tracking occurs in a serial or in a parallel fashion.

This assumption that tracking is mediated by attention is widespread (Scholl, 2001, 2009; Cavanagh & Alvarez, 2005) and is supported by experimental evidence. For example, performing an attentionally demanding task concurrently with tracking is known to reduce tracking performance (Kunar, Carter, Cohen, & Horowitz, 2008). Furthermore, probes that are briefly flashed on targets are detected more readily than probes that are briefly flashed on distractors, indicating that targets are preferentially attended (Pylyshyn, 2006; Flombaum, Scholl, & Pylyshyn, 2008; Pylyshyn, Haroutioun, King, & Reilly, 2008). Despite this evidence, the issue is still debated and at least one theory of tracking claims that tracking is preattentive (Pylyshyn & Storm, 1988).

Previously it has been suggested that tracking may occur independently in the left and right visual hemifields (Alvarez & Cavanagh, 2005). If so, serial tracking might occur only when there was more than one target in a given visual hemifield. Consequently, for Experiment 3 there were two sets of four disks, with one set of disks confined to the left hemifield and the other to the right hemifield. Both targets were confined to same hemifield thereby ensuring that seriality could be observed regardless of any hemifield independence.

At the start of each trial, each disk within a hemifield was colored differently and the two targets were indicated by being briefly ringed. The disks all then became the same color and moved around their hemifield. During this movement phase the

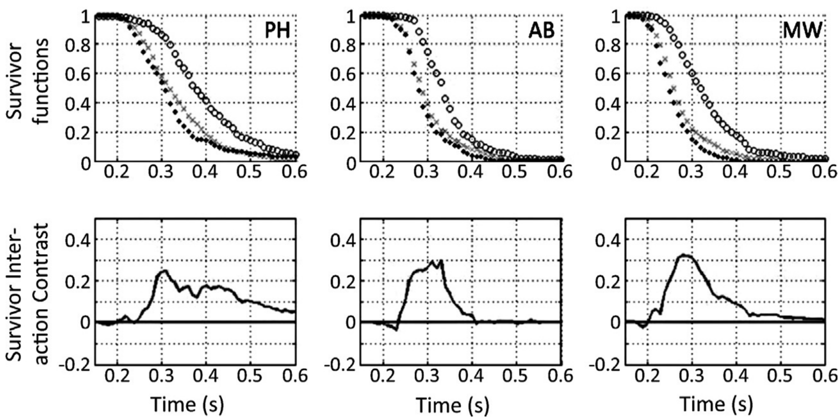

Figure 15.6 Results of the third experiment for each of the three observers. The top row represents the survivor functions (circles represent the SS condition, the gray crosses the mixed condition and the black diamonds the LL condition). The bottom row represents the survivor interaction contrasts (*SIC*s).

disks would suddenly and unpredictably darken. Each disk would darken by either a small amount (S), a large amount (L) or not at all (N). If both disks darkened they would do so at the same time. The observer was instructed to push a button on a response box a quickly as possible whenever he noticed either target darkening. At the end of the trial all the disks would stop moving and the observers were asked to indicate the location of each target. On average observers were able to correctly indicate the individual locations of both targets on 95% of the trials. In calculating reactions times, we discarded any trial were observers were not able to do this. This ensured that we analyzed only those trials where observers had been able to maintain accurate identity–location bindings.

As it is not possible to directly test for selective influence, a key assumption of systems factorial technology (Townsend & Nozawa, 1995), we instead tested for the implied ordering of the survivor functions (Houpt, Blaha, McIntire, Havig, & Townsend, 2013). The survivor functions for the three observers are shown in the top row of Fig. 15.6. By performing a series of Kolmogorov–Smirnov tests we were able to show the expected ordering.

For all observers the *MIC* was found to be significantly greater than zero (adjusted rank transform test, average $MIC = 0.046$), allowing us to reject serial processing but not to distinguish between parallel and coactive processing. To make this distinction we need to consider the *SIC* of each observer, shown in the bottom row of Fig. 15.6. The Houpt–Townsend KS–SIC test showed that for each observer the *SIC* was at some point significantly greater than zero but never significantly less than zero. This allowed us to rule out coactive processing and conclude that processing must be parallel. Furthermore, by looking at the shape of the *SIC*'s we could conclude that each observer employed a self-terminating stopping rule (please see the tutorial chapter in this volume for further details). For each observer, the capacity coefficient was found to be significantly less than one under the Houpt–Townsend UCIP test leading us to conclude that our data is compatible with both a limited capacity independent parallel model (Townsend & Asby, 1983) and a inhibitory, interactive parallel model (Eidels, Houpt, Altieri, Pei, & Townsend, 2011).

Discussion

Both the simultaneous–sequential paradigm and systems factorial technology have arrived at the same conclusion, that in MIT the targets can be tracked in parallel implying that their identity–locations bindings must also be capable of being updated in parallel. The main difference between these two methods is the assumptions that they make.

The simultaneous–sequential paradigm assumes that any serial model would preferentially track the moving targets. While, this assumption is certainly plausible, we could reach a stronger conclusion if we did not have to make it. It was for this reason that we decided to also use systems factorial technology.

Systems factorial technology does not make this assumption; however, to use it we need to make a number of different assumptions, two of which are particularly important. One is the assumption of selective influence (Townsend & Nozawa, 1995). This assumption cannot be proven but does imply a particular ordering of the survivor functions. As this ordering of survivor functions was observed for all our observers in Experiment 3, we can be fairly confident that this assumption holds for that experiment.

The other assumption that we need to make to is that tracking is mediated by attention. Systems factorial technology requires one to measure the reaction time to a discrete event (Townsend & Nozawa, 1995) so it cannot be used to directly study tracking because tracking is a temporally extended process. Instead, we assumed that tracking is mediated by attention and used systems factorial technology to study the deployment of attention in tracking. We found that during tracking the targets were attended simultaneously and from this we concluded that if tracking is mediated by attention it follows that tracking must occur in parallel.

In summary, both our techniques arrived at the same conclusion but both required us to make significant assumptions. However, the assumptions made by the two techniques were different. We argue that while neither technique on its own can be considered to be conclusive, combined they provide strong converging evidence that tracking occurs in parallel. This is therefore a good example of when it is beneficial to utilize both methods so as to come to a stronger conclusion than would be possible if either method was used on its own.

References

Abrams, R. A., & Christ, S. E. (2003). Motion onset captures attention. *Psychological Science: A Journal of the American Psychological Society*, *14*, 427–432.

Altieri, N., Fifić, M., Little, D. R., & Yang, C.-T. (2017). Historical foundations and a tutorial introduction to systems factorial technology. In D. R. Little, N. Altieri, M. Fifić, & C.-T. Yang (Eds.), *Systems factorial technology: A theory driven methodology for the identification of perceptual and cognitive mechanisms*. Elsevier (in this book).

Alvarez, G. A., & Cavanagh, P. (2005). Independent resources for attentional tracking in the left and right visual hemifields. *Psychological Science*, *16*(8), 637–643.

Cavanagh, P., & Alvarez, G. A. (2005). Tracking multiple targets with multifocal attention. *Trends in Cognitive Sciences*, *9*(7), 349–354.

Cohen, M. A., Pinto, Y., Howe, P. D., & Horowitz, T. S. (2011). The what–where tradeoff in multiple identity tracking. *Attention, Perception, & Psychophysics*, *73*, 1422–1434.

Cousineau, D. (2005). Confidence intervals in within-subject designs: A simpler solution to Loftus and Masson's method. *Tutorial in Quantitative Methods for Psychology*, *1*(1), 42–45.

Crick, F., & Koch, C. (1990). Towards a neurobiological theory of consciousness. *Seminars in the Neurosciences*, *2*, 263–275.

Di Lollo, V. (2012). The feature-binding problem is an ill-posed problem. *Trends in Cognitive Sciences*, *16*(6), 317–321.

Drew, T., Horowitz, T. S., Wolfe, J. M., & Vogel, E. K. (2011). Delineating the neural signatures of tracking spatial position and working memory during attentive tracking. *Journal of Neuroscience*, *31*(2), 659–668.

Eidels, A., Houpt, J. W., Altieri, N., Pei, L., & Townsend, J. T. (2011). Nice guys finish fast and bad guys finish last: Facilitatory vs. inhibitory interaction in parallel systems. *Journal of Mathematical Psychology*, *55*(2), 176–190.

Eriksen, C. W., & Spencer, T. (1969). Rate of information processing in visual perception: Some results and methodological considerations. *Journal of Experimental Psychology*, *79*(2), 1–16.

Fehd, H. M., & Seiffert, A. E. (2008). Eye movements during multiple object tracking: Where do participants look? *Cognition*, *108*(1), 201–209.

Fehd, H. M., & Seiffert, A. E. (2010). Looking at the center of the targets helps multiple object tracking. *Journal of Vision*, *10*(4), 19.1–19.13.

Flombaum, J. I., Scholl, B. J., & Pylyshyn, Z. W. (2008). Attentional resources in visual tracking through occlusion: The high-beams effect. *Cognition*, *107*(3), 904–931.

Horowitz, T. S., Klieger, S. B., Fencsik, D. E., Yang, K. K., Alvarez, G. A., & Wolfe, J. M. (2007). Tracking unique objects. *Perception & Psychophysics*, *69*(2), 172–184.

Houpt, J. W., Blaha, L. M., McIntire, J. P., Havig, P. R., & Townsend, J. T. (2013). Systems factorial technology with R. *Behavior Research Methods*, *46*(2), 307–330.

Howe, P., Evans, K., Pedersini, R., Cohen, M., Horowitz, T., & Wolfe, J. (2009). Attention: Selectice attention and consciousness. In W. P. Banks (Ed.), *Encyclopedia for consciousness, Vol. 1* (pp. 61–75). Amsterdam: Elsevier.

Howe, P. D., Horowitz, T. S., Morocz, I. A., Wolfe, J., & Livingstone, M. S. (2009). Using fMRI to distinguish components of the multiple object tracking task. *Journal of Vision*, *9*, 10.1–10.11.

Howe, P. D. L., & Ferguson, A. (2015). The identity–location binding problem. *Cognitive Science*, *39*, 1622–1645.

Hudson, C., Howe, P. D., & Little, D. R. (2012). Hemifield effects in multiple identity tracking. *PLoS ONE*, *7*(8), e43796.

Huff, M., Papenmeier, F., Jahn, G., & Hess, R. F. (2010). Eye movements across viewpoint changes in mulitple object tracking. *Visual Cognition*, *18*(9), 1368–1391.

Iordanescu, L., Grabowecky, M., & Suzuki, S. (2009). Demand-based dynamic distribution of attention and monitoring of velocities during multiple-object tracking. *Journal of Vision*, *9*(4), 1–12.

Kunar, M. A., Carter, R., Cohen, M., & Horowitz, T. S. (2008). Telephone conversation impairs sustained visual attention via a central bottleneck. *Psychonomic Bulletin & Review*, *15*(6), 1135–1140.

Milner, P. M. (1974). A model for visual shape recognition. *Psychological Review*, *81*(6), 521–535.

Morey, R. D. (2008). Confidence intervals from normalized data: A correction to Cousineau (2005). *Tutorial in Quantitative Methods for Psychology*, *4*(2), 61–64.

Oksama, L., & Hyona, J. (2004). Is multiple object tracking carried out automatically by an early vision mechanisms independent of higher-order cognition? An individual difference approach. *Visual Cognition*, *11*, 631–671.

Oksama, L., & Hyona, J. (2008). Dynamic binding of identity and location information: A serial model of multiple identity tracking. *Cognitive Psychology*, *56*(4), 237–283.

Pinto, Y., Howe, P. D., Cohen, M. A., & Horowitz, T. S. (2010). The more often you see an object, the easier it becomes to track it. *Journal of Vision*, *10*(10), 4.

Pylyshyn, Z. W. (2004). Some puzzling findings in multiple object tracking (MOT): I. Tracking without keeping track of object identities. *Visual Cognition*, *11*(7), 801–822.

Pylyshyn, Z. W. (2006). Some puzzling findings in multiple object tracking (MOT): II. Inhibition of moving nontargets. *Visual Cognition*, *14*(2), 175–198.

Pylyshyn, Z. W., Haroutioun, H. H., King, C. E., & Reilly, J. E. (2008). Selective nontarget inhibition in multiple object tracking. *Visual Cognition*, *16*(8), 1011–1021.

Pylyshyn, Z. W., & Storm, R. W. (1988). Tracking multiple independent targets: Evidence for a parallel tracking mechanism. *Spatial Vision*, *3*(3), 179–197.

Scholl, B. J. (2001). Objects and attention: The state of the art. *Cognition*, *80*(1–2), 1–46.

Scholl, B. J. (2009). What have we learned about attention from multiple object tracking (and vice versa)? In D. Dedrick, & L. Trick (Eds.), *Computation, cognition, and Pylyshyn* (pp. 49–78). Cambridge, MA: MIT Press.

Shiffrin, R. M., & Gardner, F. T. (1972). Visual processing capacity and attentional control. *Journal of Experimental Psychology*, *93*(1), 72–83.

Townsend, J. T. (1990). Serial vs. parallel processing: Sometimes they look like Tweedledum and Tweedledee but they can (and should) be distinguished. *Psychological Science*, *1*(1), 46–54.

Townsend, J.T. (1992). On the proper scales for reaction time. In S. L. H. Geissler, & J. T. Townsend (Eds.), *Cognition, information processing and psychophysics: Basic issues, Vol. 3* (pp. 105–120). Hillsdale, NJ: Erlbaum.

Townsend, J. T., & Asby, F. G. (1983). Methods of modeling capacity in simple processing systems. In J. C. F. Restle (Ed.), *Cognitive theory, Vol. 3* (pp. 200–239). Hillsdale, NJ: Erlbaum.

Townsend, J. T., & Nozawa, G. (1995). On the spatio-temporal properties of elementary perception: An investigation of parallel, serial, and coactive theories. *Journal of Mathematical Psychology*, *39*, 321–359.

Townsend, J. T., & Wenger, M. J. (2004). A theory of interactive parallel processing: New capacity measures and predictions for a response time inequality series. *Psychological Review*, *111*(4), 1003–1035.

Treisman, A. (1996). The binding problem. *Current Opinion in Neurobiology*, *6*(2), 171–178.

Treisman, A. M., & Gelade, G. (1980). A feature-integration theory of attention. *Cognitive Psychology*, *12*(1), 97–136.

von der Malsburg, C. (1981). *The correlation theory of brain function*. MPI biophysical chemistry, Internal Report 81-2. (Reprinted in J. L. Van Hemmen, E. Domany, & K. Schulten (Eds.), *Models of neural networks II (1994)*, Berlin: Springer.)

von der Malsburg, C., & Schneider, W. (1986). A neural cocktail-party processor. *Biological Cybernetics*, *54*(1), 29–40.

Watson, A. B., & Pelli, D. G. (1983). QUEST: A Bayesian adaptive psychometric method. *Perception & Psychophysics*, *33*(113–120), 113–120.

Wolfe, J. M. (2012). The binding problem lives on: Comment on Di Lollo. *Trends in Cognitive Sciences*, *16*(6), 307–308; author reply pp. 308–309.

Wolfe, J. M., Cave, K. R., & Franzel, S. L. (1989). Guided search: An alternative to the feature integration model for visual search. *Journal of Experimental Psychology: Human Perception and Performance*, *15*(3), 419–433.

Zelinsky, G. J., & Neider, M. B. (2008). An eye movement analysis of multiple object tracking in a realistic environment. *Visual Cognition*, *16*(5), 553–566.