Evaluating the Effects of Fatigue-Induction on Mice's Cognitive Behavior

Abstract

Characterized as one of the most commonly distressing symptoms of cancer treatment, cancer-related fatigue (CRF) is reported to be more severe and persistent than "normal" fatigue. CRF has been shown to manifest in higher intensities and longer durations among cancer patients, impair mood and functional abilities, and, most relevantly, correlate with impairments in cognitive functioning, leading to impairments in other areas such as recognition memory and explicit memory. Despite its prevalence, CRF continues to be underreported and untreated due to a lack of information about the biological mechanisms underlying this symptom and its correlated impairments in cognitive systems. In order to uncover these mechanisms, the study conducted by Wolff, et. al. (2020) observed pelvic irradiation's produced fatigue, and how it affects performance during various cognitive tasks, such as spontaneous altercations and reversal learning, as well as changes in whole-brain levels of mature and proBDNF. However, due to the possible behavioral confounds of the original study's assessment of cognitive mechanisms and the inconclusive BDNF results presented, we aim to design a new list of cognitive tasks to more effectively assess cognitive impairments manifested after fatigue-induced pelvic irradiation. We hope that this new battery of cognitive tasks can aid future research attempting to specify the underlying cognitive mechanisms responsible for the cognitive impairments seen in cancer-related fatigue.

Introduction

Characterized as one of the most commonly distressing symptoms of cancer treatment, cancer-related fatigue (CRF) is reported to be more severe and persistent in comparison to

“normal” fatigue. CRF has been shown to manifest in higher intensities and longer durations among cancer patients than healthy controls, as well as studied to impair mood and functional abilities (Bower, 2014). CRF has been studied to correlate with impairments in cognitive functioning, leading to impairments in other areas such as recognition memory and explicit memory (Feng, 2018). One of the most notable of the many expressed symptoms through CRF is cognitive impairments. Despite its prevalence, these have been largely ignored and CRF continues to be underreported and untreated due to a lack of information about the mechanisms underlying this symptom and its correlated impairments in cognitive systems (Bower, 2014).

Considering this lack of information on the biological mechanisms underlying CRF, the study conducted by Wolff, et. al. (2020) observed pelvic irradiation’s produced fatigue, and how it affects performance during various cognitive tasks, such as spontaneous alternations and reversal learning, as well as changes in whole-brain levels of mature and proBDNF.

To induce radioactive therapy effects, the study split mice subjects into “Irrad” (which experienced irradiation targeted to the pelvic region) and “Sham” (which were placed outside of the irradiator) groups. Mice were transferred into cages to measure their voluntary wheel running activity (VWRA), the number of minutes during which the wheel rotated. The mice were placed into one of the two spontaneous area tests to test radioactive effects on spatial memory: Open Field or Y-maze. The Y-maze recorded the number of “spontaneous alternations” (when the mice entered a new arm of the maze). There was a large, statistically significant decrease in VWRA in Irrad mice compared to Sham. The total distance in the Y-maze and spontaneous alternation rates were also lower in Irrad mice; however, neither were statistically significant.

To measure cognitive behavioral changes of the mice, a reversal learning behavioral paradigm was utilized to observe how quickly the mice could adapt their learned behavior

(cognitive flexibility). This task involved “training days” where mice had to poke their nose through the left hole of a three-holed cognition wall to access food, which was then switched to the right-most hole to assess reversal learning. Mice in the Irrad group were found to take more time making the equivalent number of pokes, suggesting that the mice experiencing fatigue also experienced slower learning.

The study hypothesized changes in BDNF expression were associated with fatigue-induced irradiation, though Western blot analyses showed insignificant effects of irradiation on proBDNF. The analysis of BDNF levels in whole-brains also results in ambiguity, as there is no clear distinction between certain areas of the brain that reflect BDNF changes and how these affected areas lead to changes in cognitive behaviors.

Although the results reported by Wolff, et al. (2020) reflected changes in behavior and task performance following irradiation, the mechanisms behind the seemingly impaired cognitive behaviors remain ambiguous. In addition, it is unclear whether performance differences among the Irrad and Sham mice groups were characteristic of the physical demands of the tasks and general fatigue, or if it could actually be attributed to cognitive impairments. To compensate for the lack of cognitive aspects of the rather physical activities and assess other cognitive functions that have been associated with fatigue, we designed a new group of cognitive tasks in this extension that will more effectively assess the possible cognitive impairments that are introduced after pelvic irradiation.

We utilize tasks that have been designed to specifically target cognitive functions, such as explicit memory and learning in mice, and are not heavily influenced by locomotor or physical activity and fatigue. To test learning and memory, we plan to apply the continuous delayed non-matching to sample with objects-odors (DNMS odor) task, which establishes the learned

association between odor and rewards. Parallel to the reversal learning task in the original study, this task will hopefully provide a more comprehensive assessment of abilities and visualize a clearer image of whether fatigue-induced radiation has an effect on cognitive ability in mice. Another version of reversal learning is the Paddling Pool Task (PPT), which utilizes an "escape hole" in a clockmaze filled with water to test spatial learning impairments caused by fatigue. PPT adds value to fatigue impacts on spatial and reference memory, adding additional context on spatial with distress and adding more comprehensive insight. Finally, to test the simultaneous processing of information and tasks, we will administer the attentional set shifting task, which will also measure the attention capabilities between Irrad and Sham groups.

Overall, due to the questionable validity of the original cognitive mechanisms implicated in Wolff's study, we aim to design a new list of cognitive tasks to more effectively assess cognitive impairments manifested after fatigue-induced pelvic irradiation.

Replication Methods

For our replication, we made use of the data published with the original study by Wolff et al. (2020). We then ran analyses of their original data in Python on Jupyter notebooks to see if we could replicate their statistical descriptives and figures. In going through this process, we encountered multiple obstacles with the data provided.

First, due to limited amounts of useful description and commentary, as well as a lack of consistent coding between the dataframes, we were only able to use the following datasets: correctEntries.csv for the spontaneous alternation replication, distance.csv and pellets.csv for the reinforcement learning, and vwra.csv for the VWRA measurements.

Secondly, the README description and guidelines were not useful and vague. The description did not provide any explanation for the datasets, the labeling, or the methods used to gather, organize, and clean the data. This meant that we had to make various assumptions in our analyses and had to deal with the inconsistencies of one dataset to another by ourselves. In some cases each mouse was labelled one way, in others they were not labelled the same. This lack of information complicated the process and made it much more difficult to investigate their reported results. Two very important bits of information that had to be discovered through trial and error were the analysis methods and data cleaning techniques used by the researchers. Often, the paper was very vague as to what specific mathematical methods were used to create the plots which made it extremely difficult to replicate and analyze them. For instance, it was not always clear the timeframe of testing that the data represented and how they calculated their reported "rates."

Finally, we noticed that distance.csv and pellet.csv (which contained one column per each mouse tracked) contained three columns with a significant amount of 'NaN' values. Upon further analysis, we decided to replace the NaN with zeros given that the filtering relied on grouping and adding these values and thus the zeros would have no effect on our results. It is important to note though that although we assumed that the NaN meant that the researchers did not collect information for those mice, it was never explained whether that was the case or what the NaN meant. If our assumptions are correct, then it would be important to explain why these mice (201808 reversal1-25:Irrad, 201808 reversal1-26:Sham, 201808 reversal1-28:Irrad) were not fully measured, and how this might have affected the conclusions.

Replication Results

As mentioned before, there were multiple limitations in terms of the data made available for replication and independent insights. Thus, we were forced to limit ourselves to only some

useful bits of data for further analysis. More specifically, we were only able to run analyses for the data utilized in the main figures of the original paper. Although we wanted to explore potential correlates with these tasks (such as mice age), we were unable to do so due to the lack of data for these variables that the original paper claimed to assess.

First, we focused on the Reversal Learning task. We plotted the percent change in distance from the pre-irradiation period to the post-irradiation period (Figure 1). We used the total distance traveled on the day before the irradiation began as the pre-irradiation baseline. Qualitatively our results match those reported in the paper, but the spread of both the Sham and the Irrad groups is reduced and the medians are different. This may be due to differences in the normalizing procedure. The authors state that values were *'normalized to the total on the day before irradiation'* but do not provide further details so we are not entirely sure whether we followed similar procedures. Additionally, we were not able to find their methods for this calculation in their original code due to the lack of commentary.

Next, we plotted the pellet consumption for the Reversal Learning task, averaged in each mouse over the first 6 days post-irradiation period (Figure 2). In this case, both our recreation and the paper box plots correspond exactly.

Lastly, we plotted the correlation between distance and pellet (Figure 3). We noticed that this plot was somewhat misleading in the original paper. Rather than correlating the pellet consumption with the distance traveled in the same task, the distance information came from the VWRA. Due to the lack of commentary and inconsistent coding across dataframes, it was unclear how they were able to correlate this data from two different tasks. Furthermore, as far as we could tell, there was no VWRA information for the mice that engaged in the reversal learning task which leaves open the question of whether they correlated data between two different

groups of mice or simply did not make the data accessible to perform this correlation. Regardless, it is a bit misleading to correlate metrics from two separate tasks, given that there could be other factors influencing these results (i.e., different environmental or timing factors between the tasks) and it is not clear whether the data from the two separate tasks encompasses different timescales (where there would most likely be a further effect of fatigue in results encompassing 3 versus 6 days post-irradiation, given the short effect of the irradiation indicated in the original paper). Given these practical and theoretical constraints to their original correlation of distance and pellet, we correlated the distance traveled *within* the reversal learning task with the pellet information. While the original paper reported positive correlations for both groups (using VWRA instead), our results (Figure 3) show a slight negative correlation for the sham group, where the amount of pellets consumed seemed to decrease with distance traveled.

For the Arena Tests, we gathered and analyzed the data for VWRA and for the Y-Maze to plot the changes in time activation and spontaneous alternation. After plotting, we noticed that there was somewhat of a discrepancy between our plot and the results obtained by the researchers (although the relationships between sham and irradiated remained relatively the same). The spread for VWRA was larger and started at a lower time (Figure 4), but we were not able to figure out the source of this difference due to the lack of proper commentary provided in the code and data. We also plotted a boxplot of spontaneous alternation and we can observe that though the Irrad group is not that far off, but it still had a significantly lower median than the Sham group in the experiment (Figure 5). Lastly, we correlated VWRA with spontaneous alternation (Figure 6). Although the original paper reported a relatively similar positive correlation between these two metrics for both the irradiated and sham groups, we found a much stronger positive correlation for the sham group. It is still unclear whether this discrepancy is due

to different analyses of this data between the original paper and our replication (since they did not provide sufficient commentary for us to know if we analyzed the raw data in the same way) or a larger problem with the analyses conducted in the original paper.

Figure 1

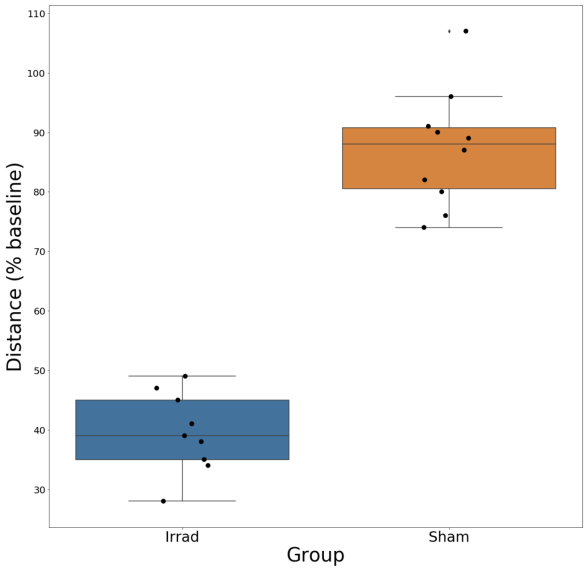

Figure 2

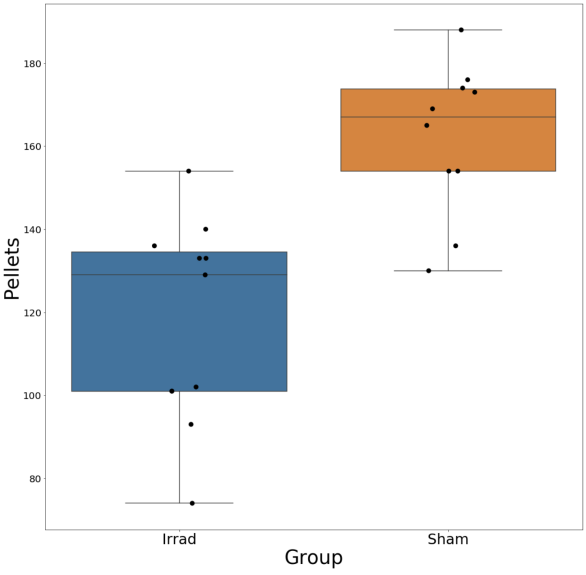

Figure 3

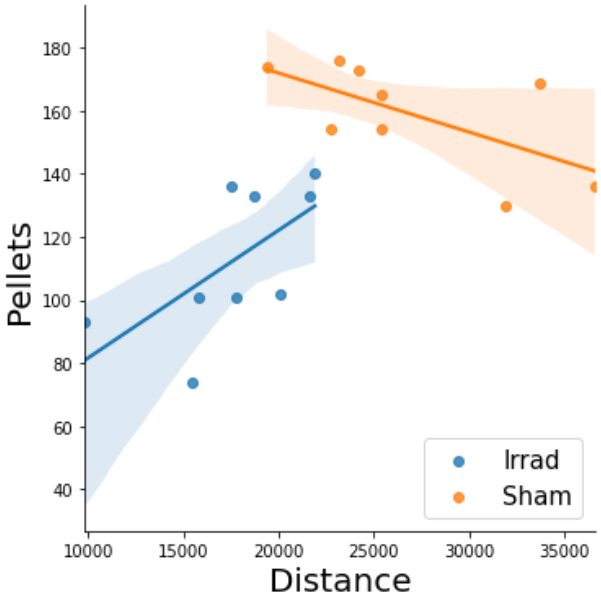

Figure 4

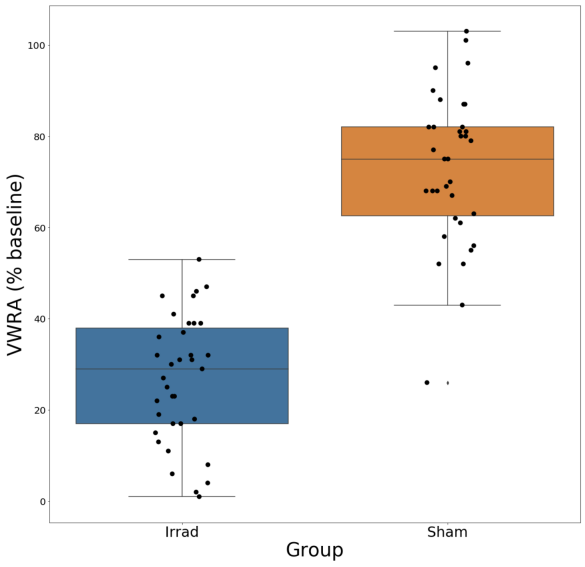

Figure 5

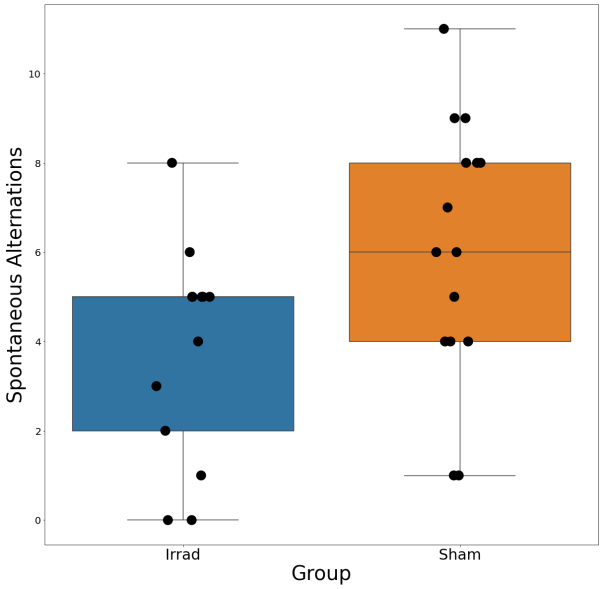

Figure 6

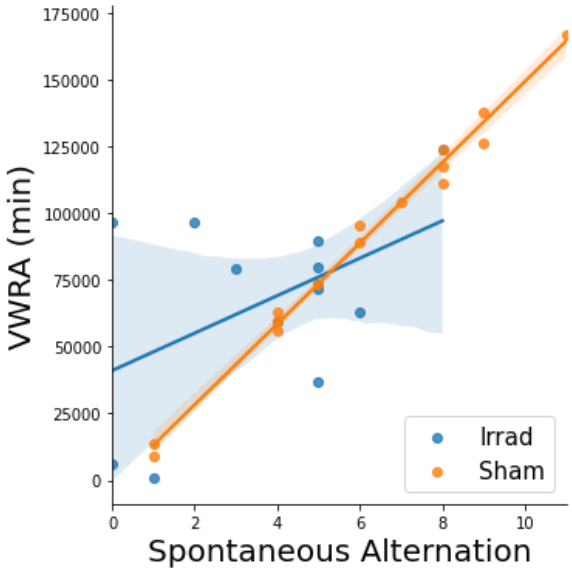

Extension Methods

Delayed Non-Matching to Sample with Objects (Odor)

All mice will be individually trained in a square chamber that has an odor port, a hole from which the mice can smell a presented odor, and a water port, which is placed above the odor port and can produce a reward (a food pellet). The task includes 18 odors: oregano, curry, dried onions, cinnamon, fennel, dried garlic, marjoram, nutmeg, celery seeds, cloves, turmeric, savory, cumin, tarragon, rosemary, dill seeds, chervil, and sage.

In trial 0, the mouse will be prompted to make a nose poke into the odor port and illicit one of the 18 available odors to be presented. Initially, a random odor will be presented through the port. If the odor is different than the one previously presented and there was a nose-poke to the water port, then a reward food pellet will be presented. If the odor is a repeat of the previous odor in trial 0 and there was a nose-poke, then there will be no reward and the nose-poke will not be reinforced. There will be 30 trials for each rat. The number of correct and incorrect responses will be recorded. A correct response means a "go" (a nose poke) on a different odor or a "no-go" (no nose poke) on the same odor as the one before. An incorrect response would mean a nose poke on the same odor or no nose poke on a different odor.

Parallel to the original study's procedures, mice will be split into a Sham and Irrad group. The Sham group will act as the control, while the Irrad Group will receive the irradiation described in the original experiment. Importantly, half of the mice from each group will be trained after the irradiation to test learning deficits; the other half will be trained before irradiation to test memory deficits. The Sham and Irrad groups will be split in half, leaving 2 Sham and 2 Irrad groups. To test memory, one pair of Sham and Irrad mice groups will be

trained to complete this task before irradiation. To test learning, the other pair of Sham and Irrad mice groups will be trained for this task after irradiation.

The Paddling Pool Task (Spatial Memory)

The Padding Pool Task (PPT) utilizes a circular arena dubbed the "clockmaze" filled with cold water to stimulate a sense of urgency in the mice and prompt them to escape the clockmaze via an escape hole. The clockmaze consists of corridors surrounding the clear arena in which the mouse can escape through as well as view the surroundings.

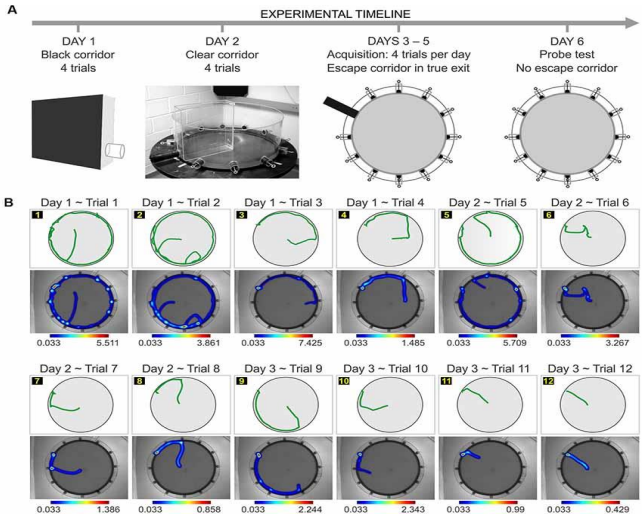

Prior to testing, each mouse undergoes pre-training on two consecutive days to familiarize them with the procedural aspects of the task (such as navigating in the cold water, going into the open hole and the adjacent pipe as a means of escape, and staying inside the pipe when they are being transported to the drying station). The first day consists of the mice being placed in the furthest point from the exit in a black corridor filled with cold water to learn the principle of escaping into the tube without spatial cues of the outside environment to indicate the correct escape corridor. The second day, the mice are presented with a clear corridor so that they *can* see the distal cues in the testing arena.

For the actual trials, eleven escape holes in the clockmaze are sealed while one is the correct exit. The Irrad and Sham groups are then incorporated into the clockmaze before and after Irrad receives pelvic irradiation. The path traversed, time to find the true exit, and number of visits to decoys are recorded.

The Attentional Set Shifting Task (Attention)

The attentional set shifting task is a multi-stage test that evaluates how quickly mice are able to associate and differentiate variables to food rewards to measure attention and cognitive flexibility. The actual experimentation period is split into two days, during which the first day focuses on testing simple discrimination (SD), compound discrimination (CD), reversal 1 (R1), and intra-dimensional shifts (IDS1). The second day continues with intra-dimensional shifts (IDS2, IDS3), reversal (R2), before concluding with extra-dimensional shift (EDS). "Cue" refers to the general stimuli, which refers to either digging medium or odor; "stimuli" refers to the specific types of cues, such as paper for digging medium or nutmeg for odor. Simple discrimination (SD) allows the mice to pair the relevant cue (such as digging medium) with the food reward. Compound discrimination (CD) allows the mice to respond to the relevant cue (such as digging medium) while ignoring the irrelevant cue (such as odor), while

Stage	Relevant Stimuli	Bowl 1 (Food Reward)	Bowl 2 (No Food Reward)
Simple Discrimination (SD)	Digging medium	Felt	Paper
Compound Discrimination (CD)	Digging medium (Trial 1)	Felt/Nutmeg	Raper/Rosemary
	(Trial 2)	Felt/Rosemary	Paper/Nutmeg
Reversal (R1)	Digging medium (Trial 1)	Paper/Rosemary	Felt/Nutmeg
	(Trial 2)	Paper/Nutmeg	Felt/Rosemary
Intradimensional Shift (IDS 1)	Digging medium (Trial 1)	Pompoms/Cinnamon	Sequins/Clove
	(Trial 2)	Pompoms/Clove	Sequins/Cinnamon
IDS 2	Digging medium (Trial 1)	Pipe cleaners/Red thyme	Googly eyes/Ginger
	(Trial 2)	Pipe cleaner/Ginger	Googly eyes/Red thyme
IDS 3	Digging medium (Trial 1)	Ribbon/Vanilla	Metallic strips/Lemon
	(Trial 2)	Ribbon/Lemon	Metallic strips/Vanilla
R2	Digging medium (Trial 1)	Metallic strips/Vanilla	Ribbon/Lemon
	(Trial 2)	Metallic strips/Lemon	Ribbon/Vanilla
Extradimensional shift (EDS)	Odor	Citronella/Raffia	Anise/Foam

still reinforcing the paired association of cue (such as digging medium) and food reward. Reversal (R1, R2) refers to swapping the previous trials' negative and positive stimuli, such as making the previously negative stimuli of paper positive and the previously positive stimuli of felt negative, still maintaining the association of relevant cue of digging medium to food reward. Intra-dimensional shift (IDS1, IDS2, IDS3) refers to swapping stimuli into new stimuli that have not been introduced, such as replacing the previous stimuli of paper to pipe cleaners, still maintaining the association of relevant cue of digging medium to food reward. Lastly, extra-dimensional shift (EDS) refers to completely swapping relevant and irrelevant cues, such as implementing the new paired association of the previously irrelevant cue of odor to food reward.

Mice are individually placed in a testing chamber, consisting of a waiting area and testing area separated by a removable screen that functions as a start gate. The testing area is also separated by a screen where two ceramic ramekins are placed. In each stage, a mouse is placed in the waiting area with the start gate closed. Pots are placed on either side of the testing chamber filled with relevant cues and food reward. The trial begins when the start gate is removed and a timer is started; the mice are allowed three minutes per trial to make a choice (denoted by poking their nose far enough or displacing the digging medium far enough to make a space). If no choice is made within the three minutes, the trial is simply marked as "No Choice (Incorrect)." Once a choice is made for a pot, the unchosen pot is removed. If a mouse has chosen the correct pot (the pot containing the food reward), the mouse is allowed to finish the food reward before being placed back into the waiting area. If the mouse

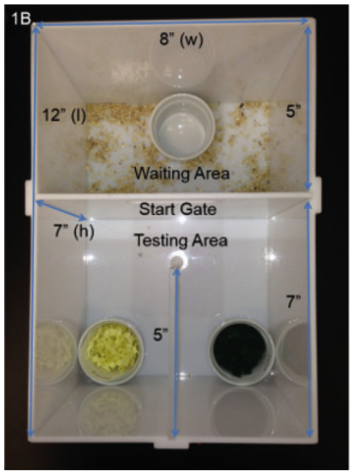

instead has chosen the incorrect pot (no food reward), the mouse is allowed to explore the chamber to indicate no food reward was presented. The measure for this task is the total number of trials the mice take per stage in order to meet a criterion value, which in this particular paper was consecutively passing eight correct trials within fifty trials to advance to the next stage. If they are unable to do so, they fail the stage and are unable to move on. Six consecutive "No Choice" choices will also result in a failure to participate and move onto the next stage. In both cases, mice that fail are ineligible to continue with the experiment, but their data in previous passed stages (if applicable) is still recorded.

Extension Expected Results & Significance

Delayed Non-Matching to Sample with Objects (Odor)

By measuring the accuracy of continuous delayed non-matching to sample with objects-odors (DNMS odor) tasks of a sample group of mice, we can establish the learned association between odor and rewards. Assuming the mice reinforce the reward-producing nose pokes, the mice will form a relationship between odor and action, as well as learn to differentiate novel and old odors.

To measure both memory and learning, this task is applied at different times in their training. If the mice were taught before irradiation, then any drop in performance among the Irrad group would suggest a decrease in motivation to participate, due to worse memory, as the mice had already been trained for this task. Whereas if the Irrad group experienced poorer performance when trained after irradiation, then the differences could be attributed to willingness or ability to learn a new task.

Comparing the accuracy results between the Sham (control) and the Irrad (experimental) group of mice will allow us to directly see the effect of fatigue-induced radiation on the two

groups' performances on the task. If we see that the Sham group maintains a better accuracy than that of the Irrad group, then we can infer that the fatigue-induced radiation did have an effect on cognitive ability in mice, especially in learning and memory.

The Paddling Pool Task (Spatial Memory)

The One Exit Test is a reference memory test to see if mice can remember the same exit using spatial memory in all twelve trials. The escape hole is changed in the One Exit Test before and after training periods. Escape latencies (time taken to find the exit) will be utilized as a spatial learning measure. The mouse's reaction to a novel exit is examined by calculating the latency difference between a trial with a new exit compared to the previous trial. We hypothesize that the Sham as well as the Irrad group will show a decreased escape latency, with the latter decreasing latency at a slower rate due to the impairments on spatial learning caused by fatigue. Increasing difference values would suggest that the mouse was confused by the switch and had to readjust their approach. Mice that take longer than sixty seconds to find the true exit are classified as having used a futile escape strategy. This is significant because we can test to see if there are increasing differences in escape latencies between the Sham and Irrad group and attribute statistically significant differences to the fatigue induced by pelvic irradiation. We expect more errors—measured by visits to the blocked off exits—to be made by the Irrad group when finding exits, as well as taking a longer time to reach the exits; this discounts impairments on the dimensions of path traversed. In analogy with humans, learning within the clockmaze is facilitated by a behavioral change away from exploratory sampling of blocked exits and towards directed movement to the true exit. The PPT adds value to the assessment of fatigue impacts on spatial and reference memory because the "5-30 minute arena tests [(Y Maze and Open Field Tests)] were not an effective way to measure fatigue induced by irradiation" (Wolff, 2020). The

PPT would add additional information on spatial memory in the context of distress (due to the cold water) that the original tests do not provide, therefore providing more holistic insight into whether fatigue's sphere of influence extends to when subjects are in distress or only with voluntary activity.

The Attentional Set Shifting Task (Attention)

In order to measure the effect of fatigue on irradiated mice, one important cognitive metric to consider is attention, specifically through tasks that require simultaneous processing of complex information and conflict-monitoring tasks. The attentional set shifting task is ideal as it allows for quantifiable measurements of the attention capabilities between groups of mice. In our specific case, we would be able to compare the performances between our Sham and Irrad group by comparing the average number of trials mice take during each stage in both groups. Although complicated in its experimental design, the principle behind it is quite straightforward: as mice form an attentional set on the relevant cue, their performance improves and is depicted by a reduction in trials in order to meet the criterion for success. In regards to our proposed follow-up of looking at the relationship of cancer-related fatigue and its effects on cognition, by simply comparing the number of trials both groups take at each stage, if all other conditions are kept constant, we will be able to quantify whether or not fatigue plays a role in reducing attentional capacities.

Discussion

Our primary replication method was generating the figures presented in the paper, specifically recreating the reversal learning, pellet consumption, and arena test visualizations. Although we were able to successfully replicate the boxplots and draw the same significant relationships among fatigue, distance traveled, pellet consumption, and VWRA, a general lack of

data and data specificity created many difficulties. To begin, there was a general lack of data, such as change in mice weight, that prevented us from doing further analyses. Furthermore, the data that was presented was organized in a confusing manner, such as inconsistent file naming and lack of specificity in the code.

Another issue we faced with the original paper was its limited insight into the cognitive mechanisms and effects of CRF. For instance, while the original study's tasks suggested significant correlations between behavior and fatigue, these same tasks also incorporated new confounding variables that hindered our analysis into the exact cognitive mechanism of CRF. To begin, both the Voluntary Wheel running and spontaneous arena tasks, although meant to measure voluntary behavior after irradiation treatment, also induced physical fatigue as locomotive tasks, which may have played a role in augmenting the mice's already fatigued state. In fact, studies show that a possible cause of the behavioral changes seen in mice after targeted pelvic irradiation stems from systemic inflammatory response (Dantzer, 2007). Radiation elevates levels of pro-inflammatory cytokines (Lalonde, 2002) that can alter behavior, which may resemble fatigue-like behavior as well as cognitive impairments (Schaue, 2012). The amount of movement—exercise—is linked with spatial learning, memory, and the hippocampus. In addition, the reversal learning tasks utilized in the original study failed to concisely distinguish the learning from other related behaviors like attention, response time, and motivation. In our extension, we explore these measures and propose that incorporating these tasks into the original set can help understand how behaviors or symptoms related to fatigue can be distinguished or at points where they overlap and precisely pinpoint the cognitive mechanisms impaired after irradiation treatment.

Most studies have shown that cancer-related cognitive impairment frequently and negatively impairs attention, memory, and information processing (Pendergrass, 2018). Thus, our extension will also cover more specific areas of cognition, such as association learning, spatial memory, attention and vigilance in order to develop a more exact understanding of the cognitive mechanism behind CRF. The Delayed Non-Matching to Sample with Objects (Odor) is an associative learning task that will further clarify the significantly different results between Sham and Irrad groups in the area of accuracy. If the Sham group ends with a higher accuracy percentage than that of the Irrad group (after the latter receives the irradiation), this would strengthen the original experiment's results that fatigue-induced radiation had a negative effect on learning. The Paddling Pool Task targets spatial memory and would provide additional information on spatial memory in the context of induced distress that the Y Maze and Open Field Tests did not provide (Wolff, 2020). The additional factor of stimulating the mice with cold water and triggering them to feel uncomfortable provides a widened insight into whether fatigue is also correlated to spatial memory and learning in situations of duress as opposed to just in voluntary situations. The attentional set shifting task delves more into attention, specifically by quantifying the attentional capabilities between our Sham and Irrad groups of mice. In both groups, the mice develop an attentional set, in which they must not only learn to associate variables with a food reward, but also be able to distinguish between specific variables. The number of trials each mouse takes to pass a specific stage allows for comparison between the Sham and Irrad group in quantifying how irradiation treatment may affect their attentional skills.

Understanding these additional factors in our extension proposal could be a vital step towards shedding more light onto the biological mechanisms of CRF. This discovery would be

crucial to modifying treatments that induce CRF and allow for better management of CRF-induced cognitive deficits.

References

Bower J. E. (2014). Cancer-related fatigue--mechanisms, risk factors, and treatments. *Nature reviews. Clinical oncology*, *11*(10), 597–609. https://doi.org/10.1038/nrclinonc.2014.127

Dantzer R, Kelley KW. Twenty years of research on cytokine-induced sickness behavior. Brain Behav Immun. 2007; 21(2):153–60. https://doi.org/10.1016/j.bbi.2006.09.006 PMID: 17088043

Lalonde R. The neurobiological basis of spontaneous alternation. Neurosci Biobehav Rev. 2002; 26 (1):91–104. https://doi.org/10.1016/s0149-7634(01)00041-0 PMID: 11835987

Pendergrass, J. C., Targum, S. D., & Harrison, J. E. (2018). Cognitive Impairment Associated with Cancer: A Brief Review. *Innovations in clinical neuroscience*, *15*(1-2), 36–44.

Schaue D, Kachikwu EL, McBride WH. Cytokines in radiobiological responses: a review. Radiat Res. 2012; 178(6):505–23. https://doi.org/10.1667/RR3031.1 PMID: 23106210