

This electronic thesis or dissertation has been downloaded from the King's Research Portal at <https://kclpure.kcl.ac.uk/portal/>



Machine learning, robotics and neuroplasticity

Kakanos, Sotiris

Awarding institution:
King's College London

The copyright of this thesis rests with the author and no quotation from it or information derived from it may be published without proper acknowledgement.

END USER LICENCE AGREEMENT



Unless another licence is stated on the immediately following page this work is licensed

under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International

licence. <https://creativecommons.org/licenses/by-nc-nd/4.0/>

You are free to copy, distribute and transmit the work

Under the following conditions:

- Attribution: You must attribute the work in the manner specified by the author (but not in any way that suggests that they endorse you or your use of the work).
- Non Commercial: You may not use this work for commercial purposes.
- No Derivative Works - You may not alter, transform, or build upon this work.

Any of these conditions can be waived if you receive permission from the author. Your fair dealings and other rights are in no way affected by the above.

Take down policy

If you believe that this document breaches copyright please contact librarypure@kcl.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

Machine Learning, Robotics and Neuroplasticity

Sotiris Gabriel Kakanos

Thesis submitted to King's College London for the degree of:
Doctor of Philosophy in Neuroscience
January 2020

Department of Neurorestoration
Wolfson Centre for Age-Related Diseases
The Institute of Psychiatry, Psychology and Neuroscience
King's College London

Supervisors: Dr. Lawrence Moon, Dr. Dhireshan Gadiagellan

Acknowledgements

To Lawrence, my supervisor: thank you for the guidance, lessons, for the support, for getting as excited as me with new ideas but for also keeping me grounded. Having you as my PhD supervisor has been a blessing.

Thank you to Dhireshan, my friend-turned second supervisor, who not only inspired me to tinker with Raspberry Pis, but has also been an indefatigable help; regardless of the day of the week or hour of the day, and despite his busy schedule.

To Emily Burnside, who sent us the DeepLabCut paper when it was first published in BiorXiv and for being an inspiring scientist.

To Rekha Patel, the lab's MSc student and now Research Technician, who helped me with my animal behaviour and who did the manual scoring of the horizontal ladder videos (thank you also for being as calmly chilled as you are).

Thank you to my Mum, who supports me always. I strive to have the same work ethic as you. To my Dad who always encouraged my interest in science as a child, and to my sister, Mags, who I could and can always turn to when I want to chat.

To my surfing partner Merrick who made the Wolfson CARD even more fun when he was here. To Nick James for making the first 6 months in the behaviour rooms a lot more tolerable.

To Ariana, my dancing and brunch partner, for always being a joy to hang out with when PhD-life got a bit too much. Thank you to Lizzie for being such a wonderful person, and who checked in with me at random, unexpected times to hear how the project was going and how I was doing. I really appreciated that. To Rob, an amazing teacher of maths, who helped consolidate my ideas for LadderScorer into algorithms. To Maria Jakobsdottir for being a massive support when we first started together.

To Professor Liz Bradbury, who gave me helpful advice after my transfer viva, and for letting us use her lab's equipment throughout my PhD project.

Thank you to the members of 'Shiteshirts', Emily Bell-Clarke and Ed(ardashian) Fletcher, for the laughs, crosswords and teas and for intermittently checking in with me during thesis writing times.

Thank you of course to the ~lovely~ members of the Moon Lab; to Dr. Aline Spejo, always on hand to discuss and challenge new ideas, to Vanessa for all the nice chats and laughs, and for indulging me in my silliness. To Ilyaas who's become my engineering partner-in-crime and who's responsible for the cool MouseBot animation, and finally to Jared for consistently baking the nicest cakes. I can honestly say I haven't tried a better carrot cake than yours.

Thank you to the Hodgkin BSU, and especially to Claire Pearce and Gary Fulcher, who were always super helpful with animal care.

Finally, a big thank you to the Python, Raspberry Pi and Stack Overflow community who made troubleshooting very easy, to the software developers responsible for DeepLabCut, which has been so influential to my project, and to Sentdex, whose YouTube channel has both been inspiring but incredibly helpful for learning new data analysis techniques with Python. I recommend his channel unreservedly.

Table of Contents

1 Introduction.....	20
1.1.1 Stroke _____	20
1.1.2 The risk of ischaemic stroke increases with age _____	20
1.1.3 Prolonged ischaemia can result in irreversible cell death _____	22
1.1.4 Upper limbs are often impaired after ischaemic stroke _____	27
1.1.5 Promoting plasticity after stroke as a way of reversing the loss of motor function _____	28
1.2 Neurotrophin-3: a candidate drug therapy for the rehabilitation of motor function after traumatic brain or spinal cord injury _____	32
1.2.2 How is NT-3 promoting this sensorimotor recovery? _____	35
1.2.3 Tools to dissect NT-3's mechanism of action _____	37
1.3 Preclinical evaluation of rehabilitative drug candidates _____	39
1.3.1 Traditional assays _____	39
1.3.2 Automating preclinical assessments of motor function _____	46
2 Understanding how NT-3 restores the loss of motor function.....	60
2.1 Introduction _____	60
2.2 Methods _____	63
2.2.1 Genotyping and confirmation of mutation _____	63
2.2.2 Dosing with 1NMPP1 _____	64
2.2.3 Photothrombotic stroke _____	64
2.2.4 Histology _____	66
2.2.5 Behavioural assessment _____	66
2.2.6 Subcutaneous dosing of NT3 _____	67
2.2.7 Detection of NT3 levels in serum and DRGs _____	67
2.2.8 Automatic analysis of horizontal ladder videos _____	68
2.2.9 Statistical analysis _____	69
2.3 Results _____	69
2.3.1 Motor cortex and corticospinal tract ablated unilaterally following photothrombotic stroke _____	69
2.3.3 Confirmation of the F617A mutation in the the TrkCF617A mice _____	72
2.3.4 The effects of photothrombosis on motor performance is long lasting, and mice do not spontaneously recover _____	74
2.3.5 Mice do not spontaneously recover their grip strength or ability to reach and grasp for sucrose pellets using the Montoya staircase following photothrombosis _____	75
2.3.6 Subcutaneously delivered NT3 following injury modestly alters reaching performance _____	76
2.3.7 NT3 is elevated in serum but not in the DRGs of treated mice _____	77
2.3.8 Staircase performance was affected by subcutaneous NT3 but not the latency to cross the horizontal ladder _____	78
2.4 Discussion _____	83
2.4.1 Photothrombosis _____	83
2.4.2 Mice do not spontaneously recover their grip strength or ability to reach and grasp for sucrose pellets 6 weeks following photothrombosis. _____	83

2.4.3	Subcutaneous NT3 is detectable in blood serum, not cervical or thoracic level DRGs, but does seem to impact reaching performance	85
2.4.4	The maximum distance reached by mice treated with NT-3 is maintained after injury, while control-treated animals decline with time	86
3 MouseBot: Automating the single pellet reaching task for mice.....		91
3.1	Introduction	91
3.3	Methods	94
3.3.1	Microchipping mice with Radio Frequency Identity tags	94
3.3.2	Component Design	94
3.3.3	Fused Deposition Modelling (FDM) 3D Printing	94
3.3.4	Laser Cutting	94
3.3.5	Software	95
3.3.6	Calibration	98
3.3.7	Motor Control	98
3.3.8	Pellet dispensing	98
3.3.9	Onboard Camera	100
3.3.10	Encoding Video	103
3.3.11	Pellet Detection using Blob Detection	104
3.3.12	Pellet positioning	105
3.3.13	MouseBot identifies each mouse using RFID technology	106
3.3.14	Data Storage	108
3.3.15	Validation	109
3.3.16	The firmware implements timeouts in the case of failure events	109
3.3.17	Training neural network for Pose-Estimation	110
3.3.18	Training a Long Short Term Memory (LSTM) recurrent neural network for trial classification	110
3.4	Results	114
3.4.1	Overview of device	114
3.4.2	Pellet Handling	115
3.4.4	Using a neural network for pose-estimation and a recurrent neural network to classify the outcomes of trials	119
3.6	Discussion	121
3.6.2	Classification of trial outcomes	125
3.7	Supplementary Materials	126
3.7.1	Figure legends for each supplementary video	126
4 Using a deep convolutional neural network to automate the scoring of walking accuracy on a horizontal ladder by mice		133
4.1	Introduction	133
4.2	Methods	135
4.2.1	Training of artificial neural network to monitor mice crossing a horizontal ladder	135
4.2.2	Development of analysis software	136
4.2.3	The configuration file tells LadderScorer the file paths of the trained feature detector's model, so that the appropriate model is used to analyse videos.	137

4.2.4	Unilateral photothrombotic stroke _____	142
4.2.5	Correlation Analysis and Agreement between manual and automated methods _____	142
4.2.7	Labelling videos with LadderScorer _____	143
4.3	Results _____	143
4.3.1	Deep neural network predicts the pose of mice _____	143
4.3.2	LadderScorer segments a video with multiple ladder crossings by a mouse into individual runs and calculates its latency to cross _____	144
4.3.3	LadderScorer can be used to determine the number of slips made by the affected limbs of mice crossing the horizontal ladder _____	147
4.3.4	Comparison of automated with a manual method of assessment revealed good agreement between methods for paws closest to the camera, but not those furthest away _____	149
4.5	Discussion _____	153
4.5.1	Pose-estimation is a good candidate for feature reduction for the quantification of locomotion metrics of mice traversing a horizontal ladder _____	153
4.5.2	LadderScorer individually segments videos containing multiple ladder traversals _____	154
4.5.3	LadderScorer automatically predicts the location of the ladder's rungs _____	154
4.5.4	LadderScorer predicts more accurately the number of slips made by the limbs in closest view of the camera _____	155
4.5.5	Conclusions _____	156
4.6	Supplementary Materials _____	156
4.6.1	Figure legends for supplementary videos _____	156
5	<i>Conclusions and General Discussion</i>	158
6	<i>References.....</i>	166

List of Figures

Figure 1 Atherosclerosis can lead to ischaemic stroke.....	22
Figure 2 The key timeline of events following onset of ischaemia detailing both injury and repair mechanisms.....	26
Figure 3 A typical horizontal ladder.....	42
Figure 4 A single pellet reaching task for mice.....	44
Figure 5 A Raspberry Pi computer	48
Figure 6 The Automated Pellet Presentation (APP) robotic system created by Fenrich et al 2015.....	52
Figure 7 The Behaviour Box proposed by Wong et al in 2015 uses robotics and computer vision to automate the single pellet reaching	54
Figure 8 The RatBot as designed by Dhiresan Gadiagellan.....	57
Figure 9 Using the photothrombotic model to unilaterally ablate motor cortex.....	70
Figure 10 Histological analysis of the brain and spinal cords of mice 6 weeks following photothrombotic stroke	71
Figure 11 Sequence alignments of the PCR amplicons obtained from <i>Ntrk3</i> ^{m1Ddg} mice homozygous for the F617A mutation and their wildtype littermates against the known sequence.....	73
Figure 12 Mice do not spontaneously recover performance on pellet reaching and bilateral grip strength 6 weeks following injury.....	76
Figure 13 Subcutaneous NT3 detected in blood serum but not DRGs.....	78
Figure 14 Subcutaneous NT3 treatment affects the reach of mice but not their ability to grasp for sucrose pellets using their affected paw.....	80
Figure 15 Subcutaneous NT3 does not impact the latency for mice to cross the horizontal ladder	82
Figure 16 The MouseBot's python-based firmware which is fully customisable	97
Figure 17 A rendered schematic of the dispenser module.....	100
Figure 18 MouseBot: a 3d-printed and laser cut unit controlled using a Raspberry Pi.....	103
Figure 19 Identifying individual mice in a multi-housed cage with a retrofitted MouseBot.....	108
Figure 21 A schematic of the pipeline used to generate a trial outcome classifier.	113
Figure 22 Validating blob detection as a method of pellet detection	115
Figure 20 Stills taken from a reaching trial where a pellet is dislodged and missed	117
Figure 23 Pellet positioning is consistent	118
Figure 24 An LSTM neural network is used to automatically classify the outcomes of trials	120
Figure 25 A screenshot taken from a video, having been analysed by our trained pose-estimation model.....	136
Figure 26 The predicted coordinates of the back left paw overlaid on the first frame of a video	140
Figure 27 The horizontal ladder used in this study	141
Figure 28 Data from the feature detector can be used to segment videos into individual runs and generate latency data	146

Figure 29 LadderScorer predicts the position of the rungs and establishes a slip-threshold	148
Figure 30 LadderScorer detects a deficit in motor function after injury	149
Figure 31 Agreements between manual assessment of foot-slips and LadderScorer-generated trial outcomes	151
Figure 32 LadderScorer determines slips of the affected hindlimbs and forelimbs of mice traversing the horizontal ladder	152

List of Tables

Table 1. Calibration document contained within an online mongoDB server which each MouseBot calls upon before initiating.....	128
Table 2. The metadata contained along with each video when uploaded to a MongoDB server following each trial.	131
Table 3. The evaluation results of the trained neural network used to analyse the videos for pose-estimation with a p-cut off value of 0.9.....	144

List of Supplementary Videos Captions

Videos can be seen by following this link: <https://sotiriskakanos.com/thesis-materials/>

Supplementary Video 1 A walkthrough of the MouseBot and its main mechanical components	126
Supplementary Video 2 A pellet being dispensed from the hopper, onto the spoon and finally into the view of the frame	127
Supplementary Video 3 A movie showing the pellet being positioned from the perspective of the camera	127
Supplementary Video 4 An example video of a trial where the pellet is missed by the mouse	127
Supplementary Video 5 An example video of a trial where the mouse successfully reaches the pellet	128
Supplementary Video 6 Animal Pose-Estimation applied to reaching videos (part 1)	128
Supplementary Video 7 Animal Pose-Estimation applied to reaching videos (part 2)	128
Supplementary Video 8 Markerless tracking of a mouse traversing a horizontal ladder	156
Supplementary Video 9 Validation video generated by LadderScorerV1	157

Abbreviations

Abbreviation	Definition
1NMPP1	4-Amino-1-tert-butyl-3-(1'-naphthylmethyl)pyrazolo[3,4-d]pyrimidine
BDNF	Brain-derived neurotrophic factor
chABC	Chondroitinase ABC
CNS	Central Nervous System
CPU	Central processing unit
CRE	Cre-recombinase
CV	Computer Vision
CST	Corticospinal tract
DLC	DeepLabCut
DRG	Dorsal Root Ganglia
FDM	Fused Deposition Modeling
FPS	Frames Per Second
IR	Infrared
LED	Light Emitting Diode
LSTM	Long Short Term Memory
MAC	Media Access Control
NT3	Neurotrophin-3
PBS	Phosphate Buffer Saline
PCR	Polymerase Chain Reaction
PLA	Polyactic Acid
PNS	Peripheral Nervous System
RFID	Radio Frequency Identity
ROI	Region of Interest
RPI	Raspberry Pi
Trk A, B, C	Tropomyosin Receptor Kinase A, B, C

Abstract and Aims

Stroke is a leading cause of disability worldwide. Three quarters of survivors leave the hospital with a limb impairment which makes their daily living activities difficult and reduces their quality of life. Currently, there are no drug therapies available for patients with stroke-induced limb impairments, other than intensive physical rehabilitation that require specialists and expensive rehabilitation facilities. New therapies are desperately required. Our lab and others have shown that delayed administration of neurotrophin-3 (NT-3), 24 hours following ischaemic stroke in rats, reverses the loss of motor function after 6 weeks. This is promising because NT-3 treatment has already been tested in humans in Phase I and II trials for other conditions and is safe and well tolerated. A long term goal of our research group is to understand how NT3 brings about recovery and to optimise its delivery with a view to treating humans after stroke. Further to this, preclinical assessment of therapies that might rehabilitate motor function after brain or spinal cord injury, are often done using manual behavioural test batteries, which are time consuming and labour intensive. We are leveraging robotics, computer vision, and machine learning to automate the scoring and assessment of these types of tasks. The specific aims of my PhD research are:

1. To set up in mice a photothrombotic model of large cortical stroke;
2. Assess the degree of spontaneous recovery after large cortical stroke;
3. Assess whether silencing TrkC kinase affects any spontaneous recovery;
4. Determine whether subcutaneous NT-3 improves recovery after large cortical stroke in mice;
5. Develop a device to train and assess reach-and-grasp function in mice before and after stroke, respectively;
6. Develop software to analyse mouse walking on a horizontal ladder.

Chapter 1 describes the work carried out to tackle (Aims 1 – 4)

Background There are currently no therapies available for the treatment of upper limb impairments caused by ischaemic stroke other than physical rehabilitation, and new therapies are desperately required. Our lab has shown using adult and elderly rats that peripherally administered NT-3 has the potential to restore the loss of sensorimotor function, possibly by promoting the plasticity of neurons spared from injury, such that they target neuronal pools previously signalled to by neurons lost from the injury. Understanding the mechanisms underlying NT-3 induced motor recovery will aid in the optimisation of the therapy used in the clinic. Namely, through which receptor does it mediate this rehabilitative effect, and where in the mammalian system does this occur?

Methods In the experiments described in this chapter, we have begun trying to answer these questions by firstly, expanding a colony of TrkC^{F617A} knock-in mice whose trkC receptors; thought to be the constituent receptor through which NT-3 promotes its motor recovery after stroke; can be pharmacologically and genetically modulated. We also sought to establish in mice, the photothrombotic model of stroke and characterise it both anatomically and behaviourally. We then treated these mice with 1NMPP1 (a drug that inhibits their trkC kinase activity) after inducing unilateral photothrombosis, and explored whether the trkC-NT3 signalling pathway plays a role in spontaneous recovery. We also treated a separate cohort of mice with subcutaneously delivered NT-3 following unilateral stroke of the motor cortex to extend the work carried out by the lab, which have shown that in rats, NT-3 restores the loss of motor function of the upper limbs.

Results Photothrombotic stroke induced consistently large ablations of the sensorimotor cortex and an 82 % loss of the corticospinal tract on the contralateral side. Our stroke model also induced sustained deficits in motor function as measured by grip strength, the latency to cross a horizontal ladder, and Montoya staircase. In

the 1NMPP1 study we show that, unexpectedly, mice, regardless of treatment group, do not show partial spontaneous recovery of their ability to reach and grasp for sucrose pellets (shown using the Montoya staircase), or regain their grip strength (shown using a bilateral grip-strength meter). As the vehicle control group did not partially recover, we could not determine whether silencing *trkC*'s kinase activity affected spontaneous recovery: however, it did not induce a further decline in motor performance compared to the control group either. In the subcutaneous NT-3 study, we show that mice, again, do not spontaneously recover their ability to reach and grasp for sucrose pellets using their affected limbs, in both treatment groups when tested using Montoya's staircase. Although, interestingly, some aspects of their reaching performance do differ: NT-3 mice were shown to reach for lower, more difficult-to-reach steps than the PBS treated group using their affected paw, while reaching distance with the affected paw of the PBS group seemed to decline over time. The latency to cross a horizontal ladder was not affected by NT-3 treatment. Also, we showed that NT-3 is elevated in serum after 28 days of infusion but in contrast to our lab's rat studies, we show that an elevation of NT-3 is not detectable in the cervical or thoracic DRGs .

Conclusions Photothrombosis is an effective model of stroke induced upper limb motor impairments because of its consistency in terms of lesion profile, and because of how well it is tolerated by mice (95% survival rate). A striking finding of these studies is the lack of spontaneous recovery of sensorimotor performance observed after photothrombotic stroke, which could prove advantageous, as it provides a large treatment window to detect the rehabilitative properties of candidate drug therapies: however, as a result we were not able to determine whether *trkC* signalling is involved in spontaneous recoveries. We were also unable to show that subcutaneously delivered NT-3 improved upper limb function in mice, and this could be because it is not detectable in the DRGs. It may be necessary to use a therapy which is more likely

to deliver NT-3 to DRGs, such as with a viral vector encoding NT-3 injected intramuscularly.

Chapter 2 details the work carried out to automate the training and assessment of the single-pellet reaching task (Aim 5)

Introduction The single pellet reaching task is an often-used behavioural test to study motor learning, neurodegeneration and to assess the efficacy of therapies that might reverse paralysis in survivors of stroke or spinal cord injury. The single pellet reaching task is typically carried out manually by research personnel. A mouse or rat is placed in a clear plastic box with a thin vertical opening in one of its walls. Two sugar pellet rewards are placed beyond the slot in fixed positions, and each is accessible only by the left or right paw, depending on the position. This task is an analogue for human reaching and as such is often used in animal research, as the ability for rats and mice to do this task is diminished when a stroke is induced that affects the motor cortex. However, carrying out this task is cumbersome and time consuming. Others have automated the single pellet reaching task but all of which are either desktop solutions or require that mice are individually housed. In this chapter, I propose a new solution that is able to determine the individual mouse executing the reach-and-grasp such that it can be used in-cage with mice which are socially housed. The solution proposed is called the MouseBot.

Methods The components of the MouseBot are laser-cut or 3D printed, and is Raspberry Pi-controlled. MouseBots use high-speed video capture, computer vision algorithms and machine learning to position sucrose pellets, while actively video recording reaching trials. MouseBots can discern which mouse is executing the trial in a multi-housed mouse cage. A custom 3D-printed hopper and dispenser mechanism feeds 20 mg pellets to a spoon which is then presented to a mouse in positions with varying difficulties to reach. The mice, which are microchipped with radio frequency tags, are recognised by the device by its onboard radio frequency identification tag reader. To automatically classify the outcomes of each trial, a convolutional neural network was trained to detect key features within the frames of

videos generated by the MouseBot. These key features were subsequently used to train a second artificial neural network (Long Short Term Memory) to classify the outcomes of trials as “hits or misses”.

Results The MouseBot, to date, has automatically run 1000s of trials. Automatically captured video footage and metadata are kept in cloud storage, accessible from anywhere via a simple web application. We trained a convolutional neural network and a Long Short Term Memory neural network to automatically classify the outcomes of trials whose accuracy is currently 91% and 80% for hits and misses, respectively.

Conclusion We show that with 3D printing, laser-cutting and open source software tools such as Python, OpenCV and Tensorflow, bespoke devices can be created with the purpose of automating preclinical tasks. Work is now required to further validate the device in a preclinical study before it is made open-source for other research facilities to manufacture their own.

Chapter 3 details the development of a software package, the aim of which is to automatically score mouse walking as they traverse a horizontal ladder.

Introduction The horizontal ladder is a behavioural task used in preclinical research to assess deficits in proprioception and fine motor function. Unlike many other tests, animals require minimal training from technicians; mice and rats can readily cross a horizontal ladder with little to no slips. Following spinal cord or brain injury, the number of slips between the rungs of the ladder increases, making it a good task to assess rehabilitative drug candidates. The rate limiting step of the analysis of this behaviour is the manual scoring of trials, which requires that recorded footage of mice or rats are watched back in slow motion so that errors made for each individual limb are quantified.

Methods Here we describe a method to automate the analysis of horizontal ladder videos using a Google Colab-based software package called LadderScorerV1. LadderScorerV1 is written in Python 3 and begins its analysis by extracting the pose of mice traversing the horizontal ladder using a convolutional neural network trained with the DeepLabCut toolbox. The features extracted from the video are the animal's limbs, the snout, the tail tip and the base of the tail. With this data, a video containing multiple traversals can be separated automatically into individual runs, which is achieved using k-means clustering; this output is also used to quantify the latency for mice to cross the ladder. The rungs are automatically located within a given video and are used to establish a "slip-threshold" which is relative to the average body length of the animal; a slip is counted when the slip-threshold be exceeded by a given limb. Finally, a labelling tool was created so that the scores and predictions of LadderScorerV1 can be overlaid on the video that was analysed so users can visually validate its outputs.

Results We demonstrate that it is possible to automatically determine the latency for mice to cross a horizontal ladder and the number of slips made by traversing animals.

LadderScorerV1 detected an increase in both the latency for mice to cross and the number of slips they made 7 days following unilateral photothrombotic stroke. We show the LadderScorer's ability to segment footage containing multiple traversals into individual runs, and quantify the latency for mice to cross with each traversal, and that it is accurate and comparable to a human scorer, respectively, with an r value of 0.99 as determined by Pearson's correlation analysis. LadderScorerV1 can also detect the number of slips made by the affected limbs of mice traversing the horizontal ladder comparably to a human scorer with Pearson's correlation coefficients of 0.94 and 0.86 for the hindlimbs and forelimbs respectively. LadderScorer's ability to quantify the slips made by the limbs furthest from the camera do not correlate well to the human assessor with Pearson's correlation coefficients of 0.51 and 0.43 for the hindlimbs and forelimbs respectively.

Conclusions LadderScorer is a Google Colab based software package which means that it can readily be used on any device to analyse and score footage of horizontal ladder footage, which would otherwise require expensive and dedicated graphics processing units. Work is still required to more accurately quantify the slips made by the limbs furthest from the view of the camera but we show that the software is capable of accurately quantifying the latency for mice to cross the ladder, as well as slips made by the limbs in closest view of the camera.

1 Introduction

1.1.1 Stroke

Stroke is caused by cerebrovascular aetiology and is defined by the World Health Organisation as an acute neurologic deficit lasting greater than 24 hours. Stroke can be subdivided into two categories: ischaemic stroke or haemorrhagic. The former being the result of vascular occlusion or stenosis (a narrowing of the arteries), and the latter, the result of vascular rupture, with subsequent intraparenchymal and/or subarachnoid haemorrhage. Ischaemic stroke accounts for 85% of stroke cases; 15% of strokes are hemorrhagic (Stroke Association, stroke.org.uk). Ischaemic stroke is a worldwide leading cause of mortality and morbidity. 17 million people worldwide experience a first time stroke each year, which equates to a stroke every 2 seconds (Feigin et al., 2014). Survivors of stroke will often have to leave work and rely on family, friends or carers to help with daily living activities and the annual economic burden this has is estimated to be in the region of £9 billion in the UK (Stroke Association, stroke.org.uk).

1.1.2 The risk of ischaemic stroke increases with age

The brain's maintenance relies on a constant supply of oxygen and glucose, and interruptions of blood flow can result in irreversible neuronal death. There are a number of ways that an ischaemic stroke can occur. For example, by atherosclerosis, which is the progressive narrowing (stenosis) and closure of arteries from the accumulation of fatty deposits over time, referred to as atheromas. Arteries become harder and narrower with age, which is a key risk factor of ischaemic stroke, and atherosclerosis can accelerate this. Lifestyle factors such as low activity or unhealthy diets are risk factors in the development of atherosclerosis as well as other medical conditions such as high blood pressure, high cholesterol levels and diabetes. Atheromas can occur anywhere in the circulatory system, including the carotid arteries, which provide blood flow to the brain. Atheromas, as well as narrowing the

passage of blood flow through the arteries, can break off and cause an upstream blockage (embolism), or become inflamed and lead to the formation of local blood clots (thrombosis) (Figure 1).

As well as atherosclerosis, other factors can cause ischaemic stroke, including small vessel disease, patent foramen ovale, atrial fibrillation and arterial dissection. The cause of stroke is important in the acute phases of its incidence so that the most optimal treatment is given. The cause is determined by the myriad brain scans and tests conducted upon admittance to a hyper-acute stroke unit, including magnetic resonance imaging (MRI), computed tomography (CT) scans and electrocardiograms (ECG). When a person is suspected of having an ischaemic stroke, it is critical that they are diagnosed quickly after onset of symptoms so that perfusion of the brain can be re-established should the type of stroke be ischaemic. This is either done surgically (thrombectomy) or with recombinant tissue plasminogen activator (rt-Pa), a clot-busting agent (thrombolysis). If reperfusion by thrombolysis is established within 4 and a half hours of stroke onset, neuronal damage can be mitigated. Mechanical thrombectomy up to 16 hours after ischemic stroke improved outcomes for patients with evidence of penumbra by perfusion imaging (Albers et al., 2018) DEFUSE trial. Public health service efforts to educate people on the signs of stroke have been successful in preventing stroke induced-deaths and more people are surviving than ever before; 8 out of 10 people who are eligible for rt-Pa treatment receive it, however, only 10% of people who arrive at the hospital with ischaemic stroke are eligible (Stroke Association, stroke.org.uk).




Figure 1 **Atherosclerosis can lead to ischaemic stroke.**

The progressive build up of fatty deposits in key arteries that supply blood to the brain lead to the development of atheromas. Atheromas narrow the passage of the arteries and reduces blood flow. An ischaemic stroke can occur when an atheroma breaks off and blocks a small vessel upstream, or when inflammation occurs at the atheroma itself, causing a blood clot. (source of image: bupa health, bupa.co.uk)

1.1.3 Prolonged ischaemia can result in irreversible cell death

Re-establishing blood flow as quickly as possible following the onset of ischaemia is critical in preventing neuronal damage and limiting disability; reperfusion with rt-tPa or mechanical thrombectomy is associated with an improved clinical outcome. Ischaemic strokes are focal, and at the central core of ischaemia (where blood flow is 10 to 25% of normal), cell death occurs within minutes (Bejot et al., 2007). The ischaemic-cascade theory, first proposed in 1987 (Nesto and Kowalchuk, 1987) and refined more recently (Maznyczka et al., 2015), has been reviewed extensively

elsewhere (Xing et al., 2012, Candelario-Jalil, 2009). However, briefly, the ischaemic-cascade is the events that follow the disruption of blood flow that ultimately lead to irreversible cell death and can be summarised by these key events:

1. Adenosine triphosphate (ATP) synthesis is disrupted, leading to the loss of ion homeostasis in neurons (Candelario-Jalil, 2009).
2. Glutamate, the primary excitatory neurotransmitter is released in abundance from cells whose ion homeostasis is compromised, and binds to its postsynaptic receptor targets, NMDA, AMPA, iGluR and mGluRs on surrounding cells resulting in excitotoxicity. Glutamate's reuptake by glia and astrocytes is also compromised (Lipton, 1999).
3. Cortical spreading depression is induced, which increases the energy demand of surrounding penumbral tissue (Dreier, 2011).
4. The uncontrolled cellular influx of calcium leads to the production of free radicals, which damage the cells further (Lo et al., 2003); the intracellular influx of ions creates a detrimental osmotic gradient, causing uptake of water and the formation of edemas.

The events that follow ischaemia do not unfold in a way suggested by its name "cascade", i.e. one event leading to the other, and may also occur when re-perfusion is established (after, for example, thrombolysis) and have likely already taken place by the time the patient is admitted to the hospital. The tissue surrounding the central core of ischaemia, the penumbra, is hypoperfused but not infarcted, and contains salvageable tissue that can remain viable for several hours. The penumbra can be therefore be rescued with the re-establishment of blood flow, however, if re-perfusion is not re-established this tissue will undergo cell death.

Researchers have fought to protect the penumbra with novel neuroprotective therapies. However, to date, a clinically protective neuroprotective agent does not exist. Pre-clinical trials have shown that therapies that interrupt the progression of the

ischaemic-cascade are promising, but the window to deliver them is very narrow (Figure 2), which may be part of the reason they have failed in clinical trials. For example, AMPA or NMDA receptor inhibitors reduced the infarct volume following the onset of ischaemia in animals (Bruno et al., 2001), but failed in clinical trials (O'Collins et al., 2006, Ikonomidou and Turski, 2002).

1.1.3.1 Neuroinflammation in the post-stroke brain as a target for therapeutic intervention

Ischaemic stroke is also associated with a marked neuroinflammatory response, with the immediate reactive oxygen species (ROS) produced leading to cytokine and chemokine release from ischaemic cells (Wang et al., 2007). Astrocytes and microglia become activated and remain so for up to at least a month after the onset of ischaemia (Ahmad et al., 2014, Li et al., 2014). The activation of both cell-types is associated with both the progression of and protection against further neural damage as a result of ischaemia, as described below.

Astrocytes proliferate and adapt rapidly following the onset of stroke, and secrete pro-inflammatory cytokines, chemokines and metalloproteins that contribute to the damage and opening of the blood brain barrier (Dong and Benveniste, 2001). Astrocytes and fibroblasts also contribute a protective barrier, or scar, around the damaged tissue. While the supposed purpose is to protect from further damage, with it surrounding the infarcted area, the heterocellular scar may later impede functional recovery, if new connections cannot be formed through this scar (Hao et al., 2016). The activated astrocytes can also have positive roles in the ischaemic brain, taking up excess extracellular glutamate and mitigating further excitotoxicity, as well as releasing neurotrophic factors (Amantea et al., 2015), so their presence and activation in the post-stroke brain is complex.

Similarly, activation of microglia in the post-stroke brain is associated with both positive and negative effects on the pathological progression of ischaemia. Microglia

are the resident immune cells within the brain and serve to support a stable physical and chemical micro-environment in the CNS. Microglia, when activated after ischaemia, broadly exist in two distinct states: M1 and M2. The M1 state is largely seen as one which is pro-inflammatory and is the most prevalent microglial phenotypes following the initial stages of ischaemia. M1 microglia contribute to ischaemic damage by releasing pro-inflammatory mediators such as TNF and IL-1 (del Zoppo et al., 2007). Its switching to the M2 state, the anti-inflammatory form, reduces this inflammatory response and contributes to tissue regeneration by releasing IL-10 and TGF β (Picascia et al., 2015).

Damage to the blood brain barrier leads to the infiltration of immune cells from the peripheral circulation, and these too, contribute to the progression of neuronal damage, particularly in the peri-infarcted region (Price et al., 2004). The opening of the blood-brain-barrier leads to the infiltration of granulocytes and monocytes, which in the acute stages following stroke are predominantly pro-inflammatory (Gliem et al., 2012). Contributing further to damage are neutrophils, whose number in the brain is correlated with the size of the ischaemic volume (Ahmad et al., 2014), both $\alpha\beta$ T cells and $\gamma\delta$ T cells, which are observed in the neural tissue following ischaemia (Gelderblom et al., 2009), and natural killer cells whose presence in the post-ischaemic brain, particularly after re-perfusion, is associated with neuronal death likely owing to the perforins and the interferon- γ they release (Gelderblom et al., 2012). Conversely, several studies have shown that regulatory T cells, whose source is also from the peripheral circulation, limits ischaemic damage (Liesz et al., 2009, Gill and Veltkamp, 2016).

In summary, the neuroinflammation/immune response that follows ischaemic stroke is complex with aspects of its events being both initially damaging, but later protective and it is possible that for this reason many potential drug candidates have not been approved for stroke patients beyond clinical trials. Therapeutic interventions

that disrupt this inflammatory cascade have shown promise in experimental models of stroke (Han and Yenari, 2003), and although none are yet to be approved, many clinical trials (summarised and reviewed by Drieu et al) of drugs which target post-stroke inflammation have taken place (Drieu et al., 2018).

While there are still opportunities to reduce the infarct size and mitigate the extent of damage cause by ischaemia with immediate therapeutic intervention, whether that be through anti-inflammatories or agents which scavenge free-radicals, an alternative strategy with a larger therapeutic window is promoting recovery after the ischaemic stroke has occurred.

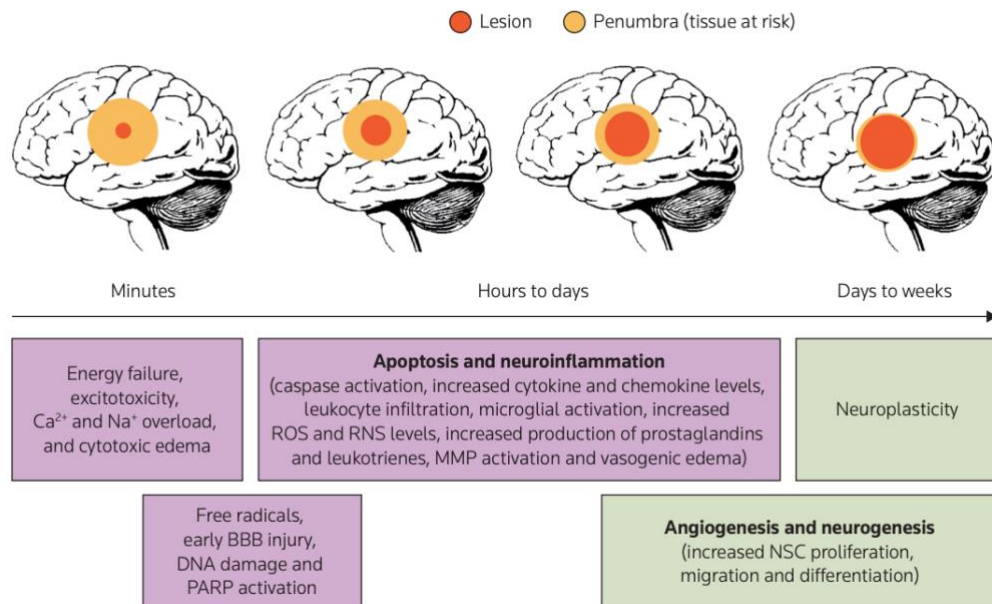


Figure 2 **The key timeline of events following onset of ischaemia detailing both injury and repair mechanisms.**

After the immediate onset of ischaemia, cell death occurs within minutes at the central core of blood occlusion (depicted by the red circle). The tissue that surrounds this infarcted zone is the penumbra, which is also at risk of death by way of apoptosis and neuroinflammation unless reperfusion of blood is re-established promptly. In the days that follow the ischaemic stroke, the brain-initiated mechanisms to repair itself from the ischaemic damage; namely angiogenesis and neurogenesis with increased

neuronal stem cell proliferation, migration and differentiation. The initiation of neuroplasticity also occurs, with the intact, surviving peri-infarcted regions re-wiring in ways that are associated with a spontaneous recovery. Image source: (Candelario-Jalil, 2009).

1.1.4 Upper limbs are often impaired after ischaemic stroke

Approximately 80% of people who suffer an ischaemic stroke will leave hospital with a motor impairment on one side of their body (Langhorne et al., 2009), with two thirds of those patients still perceiving a loss of their arm function as a problem 4 years after (Broeks et al., 1999); just 38% of patients who displayed upper limb paralysis early after stroke regained some dexterity following 6 months (Kwakkel et al., 2003). Upper limb impairments may manifest as limb weakness, spasticity or numbness. Currently, there are no NICE, EMA or FDA approved drug therapies for the purpose of reversing lost motor function due to stroke. The options available are intensive rehabilitation or electrical stimulation which both require specialists and costly facilities. The resources to deliver intensive early rehabilitation following stroke are few and far between, and for patients with chronic stroke, virtually non-existent (Ward, 2017). Furthermore, while the initially observed severity of upper limb impairment most robustly predicts the level of recovery one can expect to make, especially those whose initial impairment is mild to moderate [termed the “proportional recovery rule” (Prabhakaran et al., 2008)], it is still not entirely predictive; approximately half of patients whose observed initial impairment of injury is severe will not recover their upper limb function (Ward, 2017), despite rehabilitation. Understanding the factors that limit the recovery of upper limb function following stroke in some people will aid in the discovery of new therapies, which are desperately required. Care after stroke in the acute and hyperacute periods has greatly improved over the past 20 years, but our attention must now turn to treatments that actively promote recovery.

1.1.5 Promoting plasticity after stroke as a way of reversing the loss of motor function

Humans have the capacity to spontaneously recover motor function after stroke, which can sometimes persist years after the initial injury (Cramer, 2008). This has been replicated in pre-clinical models of stroke, although the time course is a lot quicker. The initial impairment brought about by stroke can recover to some degree, but often not to the performance level prior to injury. The level of spontaneous recovery one can expect from animal models of stroke also depends on the severity of the injury; smaller focal ischaemic injuries in rats can be recovered from fully, and truly, without compensation (Moon et al., 2009), while larger strokes often lead to motor tasks being completed by compensatory behaviours. For example, rats will compensate for their loss of upper limb function and reduced range of motion executing a reaching task by adjusting their posture (body angle, or shoulder and trunk angle adjustments) (Whishaw, 2000).

Underlying these changes is neuroplasticity: the adaptation and re-wiring of neurons spared from the injury, that form new functionally meaningful circuits. Parallels are often drawn between the developing and the post-stroke brain (Cramer and Chopp, 2000); for example, there are similarities between recovery of motor function in humans after stroke and the acquisition of new motor skills in infants. There are key differences, though, in that the capacity for plasticity to happen is diminished in the elderly: animal studies looking at the substrates important for neuronal growth, synaptogenesis and proliferation decline with age (Hattiangady et al., 2005). However, after stroke, there seems to be a critical period of heightened expression of these aforementioned substrates (proteins and genes) which allow for plasticity to occur once again, but only for a short period of time; rats who undergo intensive rehabilitation 30 days after stroke have little recovery compared to those who initiate rehabilitation at 5 or 14 days after, which recover significantly (Biernaskie et al., 2004). There are also both positive and negative factors which exist to promote and

sequester plasticity in the post-stroke brain, respectively, and have been the targets of therapeutic intervention such that post-stroke plasticity can be made more functionally meaningful. For example, chondroitin sulphate proteoglycan and Nogo-A, which serve to repel sprouting axons and inhibit outgrowth have been targeted by chondroitinase ABC (Soleman et al., 2012) and anti-Nogo-A (Wahl et al., 2014) in animal models of stroke, respectively, with positive functional outcomes. Promoting recovery of plasticity is a promising area of research gaining a lot of traction. We are exploring the use of a positive factor, the mammalian growth factor neurotrophin-3, to promote a recovery in lost motor function post-stroke, which is discussed in more detail later in this chapter.

1.1.5.1 Modelling Ischaemic stroke preclinically with photothrombosis

As with most other preclinical models of disease, capturing the full breadth of human stroke in an animal model is difficult. Our lab's ultimate goal is to restore the movement of patients who have suffered a stroke. The preclinical models of focal cerebral ischaemia used in our lab historically have included the Middle Cerebral Artery Occlusion (MCAO) model, induced by electrocoagulation and tandem carotid occlusions (Duricki et al., 2018b) and the endothelin-1 model of ischaemic stroke (Duricki et al., 2016, Soleman et al., 2012). Briefly, the former, MCAO, is achieved by exposing the carotid arteries of an anaesthetised animal, permanently occluding one side with a thread, while temporarily occluding the other for a user specified duration of ischaemia (typically between 45 minutes to an hour). Then, the middle cerebral artery is occluded by electrocoagulation, having exposed it by craniotomy.

The MCAO model with tandem carotid occlusions is a technically challenging surgery. Mortality rates can be high and the lesion sizes and their location are prone to be variable.

The endothelin-1 model of focal stroke also works by occlusion of blood flow at the MCA. While the carotid arteries need not be occluded, a craniotomy is still

required so that endothelin-1 (ET-1), which is a long-acting and potent venous and arterial vasoconstrictor, can be applied to the MCA. The craniotomy and application of the ET-1 is done by stereotaxic injection, however, due to the variability in MCA size and location, infarct sizes can vary between experimental animals.

The most commonly used model of focal cerebral ischaemic stroke in mice and rats is induced using the intraluminal suture-filament model (Longa et al., 1989). This is achieved by creating a physical barrier in the middle cerebral artery using a thread, or suture, passed through the common carotid artery where it is extended intracranially through the internal carotid artery where it blocks blood flow. The filament can then be removed after some time, allowing for re-perfusion once more, i.e. the 'transient model' or left in place and tied at the stump, achieving permanent occlusion. The benefit of this model is the option to re-establish blood-flow, quickly, which allows researchers to study the molecular events that occur shortly after blood-flow is re-established in the clinic. However, this model causes infarcts with large variability, again owing to intra-species differences in vascular architecture, and deficits in dexterity that are not long lasting.

A good model of stroke-induced motor impairments would be one that causes ischaemia of the motor cortex reproducibly, with consistent infarct sizes and brain locations across animals, causing large, long-term, measurable motor impairments. This model will also ideally be well tolerated by mice, with high survival rates following injury induction, and induce long lasting deficits in motor performance such that the rehabilitative effects of a potential drug candidate can be detected against control treated animals.

The photothrombotic model of stroke fits these criteria. Photothrombotic model of stroke was first characterized by Watson et al in 1985 for use in rats (Watson et al., 1985). It has since been widely used in other experimental animals such as mice (Labat-gest and Tomasi, 2013). Photothrombosis is achieved by systemically

administering a photoactive dye, Rose Bengal, and illuminating the brain through the intact skull, above the cortical region to be lesioned. Illumination of the dye leads to electron transfer from dye to circulating oxygen, yielding high energy singlet oxygen. This leads to peroxidase damage to the vessel endothelium and vasoconstriction, which, in turn, leads to increased platelet adhesion, aggregation and subsequently, reproducible distal cortical infarcts (Watson et al., 1985). This method of inducing stroke has many advantages compared to middle cerebral artery occlusion (MCAO) for modelling motor impairments of stroke. Surgery to induce photothrombotic stroke is less technically challenging and takes around 45 minutes to complete. Data from our lab (unpublished) has shown the survival rate of lesioned animals is high (95%) and the position of cortical infarcts is directed by the surgeon because the infarction occurs where the illumination is directed, not the pattern of vasculature, which varies between experimental animals even of the same strain (Fox et al., 1993).

Preclinical models of disease do not completely capture the full extent of human pathology, and thus it is important to select the model that is most applicable to the hypothesis under question. Given that our lab's work aims to restore motor function after stroke, it is most important that we consistently lesion the motor cortex, that the infarct size remains consistent, that these lesions cause long lasting functional deficits, that they ablate the corticospinal tract, and that the surgery-induced mortality of animals remains as low as possible, especially as elderly mice may be used. For this reason, photothrombosis is the most appropriate model of inducing stroke to test hypotheses/therapies concerning sensorimotor recovery over time frames greater than a few weeks.

1.2 Neurotrophin-3: a candidate drug therapy for the rehabilitation of motor function after traumatic brain or spinal cord injury

Neurotrophin-3 (NT-3) is a nerve growth factor whose synthesis takes place in muscle spindles during development, and is essential for the survival of large diameter sensory neurons (Ernfors et al., 1994, Wright et al., 1997) and for their correct patterning in the CNS. NT-3, as well as other neurotrophins, are processed in the endoplasmic reticulum as preproteins (Longo and Massa, 2013). After signal peptide cleavage, the preprotein form of NT-3 becomes a pro-neurotrophin and is further proteolytically processed by proprotein convertase subtilisin kexin (PCSK) enzymes to yield the mature form of NT-3, which is secreted into the extracellular space. TrkC is the principal receptor through which NT-3 binds, and with the highest affinity. NT-3, bound to the full-length tropomyosin receptor kinase 3 (TrkC), leads to the autophosphorylation of tyrosine residues in its intracellular domain.

This autophosphorylation forms adaptor binding sites that allows TrkC's intracellular domain to couple to and activate either mitogen-activated protein kinases (MAPKs), phosphoinositide 3-kinase (PI3K), or phospholipase C-gamma-1 (PLCgamma1) pathways. The binding of NT-3 to full-length TrkC, expressed on neurons, ultimately results in neurite outgrowth or survival. NT-3-null mice die shortly after birth and lack large diameter sensory neurons; motor neurons mostly remain unaffected (Ernfors et al., 1994). In TrkC knock-out mice, deficiencies in sensory neurons are lesser than in NT-3 null mice (Tessarollo et al., 1997). The differences in the outcomes of knock-outs of both NT-3 and TrkC suggest that NT-3 can signal via alternative pathways. Indeed, NT-3 has been shown to interact with a related receptor, TrkB, resulting in similar effects to its binding of TrkC. Survival of large diameter, proprioceptive neurons can be maintained by NT-3 in the absence of TrkC, if they express TrkB (Fariñas et al., 1998).

After development, the roles of NT-3 and TrkC are largely unknown with NT-3 levels in serum greatly reduced in adults and elderly (Murase et al., 1994). The expression of TrkC receptors persist into adulthood in the DRGs (McMahon et al., 1994), and elsewhere (Altar et al., 1994). While little is known about the roles they play after development, mRNA levels of both TrkC and NT-3 have been shown to increase following voluntary exercise in muscle and the spinal cord (Ying et al., 2003), and after rehabilitation following spinal cord injury (Hutchinson et al., 2004, Côté et al., 2011). NT-3 may act as a muscle derived signal which prompts the reorganisation or adaptation of proprioceptive circuits for efficient motor function and control after neurological injury. In addition to this, it has been shown that mRNA levels of NT-3 in the ischaemic cortex are significantly elevated 72 hours following occlusion compared to sham operated rats. Further, NT-3 was also found to be significantly elevated in other brain structures at this same time point including the contralateral cortex, and the subcortical areas of both hemispheres (Dmitrieva et al., 2016). As it has been shown that NT-3 plays a significant role in the formation of new synapses in the hippocampus (Shimazu et al., 2006a), its upregulation following ischaemia in the aforementioned brain structures may be part of a wider repair mechanism that regulates the formation of new connections made to restore those lost due to ischaemic injury. Taken together, it is possible that NT-3 has endogenous functionally restorative properties owing to its influence on neuroplasticity, and administering it exogenously may act to promote a functional recovery further.

1.2.1.1 Intramuscular injections of AAV-neurotrophin-3 promote the plasticity of intact spinal circuits and a sensorimotor recovery in post-stroke rats

NT-3 has long shown potential to be a good rehabilitative therapy following unilateral spinal cord injury. Early work carried out by Professor David Shine's group and others showed that NT-3, when expressed in motor neurons at the lumbar level

ipsilateral to a unilateral pyramidotomy, induced sprouting across the spinal midline of axons originating from the motor cortex on the ipsilesional side (Zhou et al., 2003).

Duricki et al later demonstrated that NT-3 restored the loss of upper limb motor function of rats who had undergone unilateral cortical strokes, induced by endothelin-1. In this study, NT-3 was introduced via an adeno associated virus (AAV) with a gene encoding the human variant of pre-pro-neurotrophin-3 (AAV-hNT3) into the biceps and triceps muscles of one of the affected forelimbs. This delayed introduction of NT-3, 24 hours after unilateral stroke, induced sprouting of the less-affected CST and a sensorimotor recovery in the rats treated (Duricki et al., 2016).

The mature, active form of NT-3 is conserved between both species of rats and humans and ELISA assays are unable to distinguish between the two. Quantitative real time PCR can be used to quantify levels of their respective mRNA, as the transcripts vary in the gene sequences that encode the mature NT-3 precursors. Human NT3 mRNA was not detectable in the spinal cords of the animals treated by Duricki et al, nor were they present in the DRGs, despite the functional recovery meaning this tissue had not been infected by the AAV-hNT3. On the other hand, the mature NT-3 protein itself was elevated in the muscles in which the AAVs were injected, and also in the ipsilateral DRGs and could either have been transported there from the circulating blood, or directly, by uptake of NT-3 by afferent fibres terminating in the muscle (Kathe et al., 2016). This data suggests that NT-3 expression in the spinal cord is not necessary to induce functionally meaningful plasticity in the spinal cord (in the way demonstrated by Zhou et al) after intramuscular delivery of NT-3, and that NT-3 is acting on a receptor target outside of the CNS.

1.2.1.2 Intramuscular infusion of NT-3 protein also restores lost sensorimotor function after ischaemic stroke

Accordingly, Duricki et al later demonstrated that infusion of human prepro-NT3, via an osmotic mini-pump whose outlet catheter terminated in the affected arm triceps muscle, also improved sensorimotor function after induction of cortical stroke

(Duricki et al., 2019). Anterograde tracers administered to the non-lesioned side of the motor cortex showed that axons from the less affected side of the corticospinal tract crossed the midline and sprouted to the injured hemicord in NT-3 treated animals.

Subcutaneous delivery of NT-3 is a clinically translatable method of administration following ischaemic stroke. In phase I and II clinical trials for disorders other than stroke, repeated high doses of recombinant NT-3 was shown to be safe and well tolerated (Sahenk, 2007, Parkman et al., 2003, Chaudhry et al., 2000). Furthermore, it has been shown that peripherally administered neurotrophins accumulate in sensory ganglia where they can cause gene transcription, axon growth of primary afferents, and synapse strengthening within locomotor circuits (Lee et al., 2012, Chen et al., 2002) and induce the same effects of muscle-derived NT-3.

In summary, NT-3 administered peripherally, using a clinically feasible method of delivery, can still bring about functionally meaningful anatomical changes in the CNS. Parallels can be drawn from the NT-3's involvement in the development of proprioceptive circuits. It has been shown that proprioceptive neurons can be rescued in NT-3 null mice during development when NT-3 is overexpressed in the muscle alone, and that this target-derived source of NT-3 is perhaps more preeminent in supporting the development of proprioceptive neurons than intraganglionic NT-3 (Wright et al., 1997).

1.2.2 How is NT-3 promoting this sensorimotor recovery?

Most of the previous studies from our laboratory mentioned show that NT-3 measurably induced sprouting of the CST spared from injury. Reorganisation of supraspinal axons spared from injury is a high-level mechanism through which function is restored. It has been shown that the sprouting of supraspinal axons is an integral part of functional recovery following other interventions for stroke. By using a

Designer Receptor Exclusively Activated by Designer Drug (DREADD) strategy selectively targeting the corticospinal axons sprouting across the midline, it has been shown that inhibition of these axons reversibly impairs the animal's reach-and-grasp to a pre-recovery performance level after treatment with antibodies against Nogo-A (Wahl et al., 2014).

Untreated control animals still spontaneously recover following injury, however, this recovery is hastened by NT-3 and increases the size of recovery. Kathe et al show that it is normalisation of proprioception circuits which are perturbed after injury that may drive sensorimotor recoveries (Kathe et al., 2016). Proprioception is an awareness of our limbs whereabouts in space, and relies on NT-3:TrkC signalling during development in order for its underlying spinal cord circuits to correctly form. This circuitry is required for correct motor control, coordination and balance and is orchestrated by signals carried from muscle spindles to the spinal cord by group Ia and II afferent fibres, whose cell bodies reside in the dorsal root ganglia (Sherrington, 1913). This circuitry conveys information about muscle stretch to the spinal cord. When action potentials arrive in the spinal cord from muscles they either modulate muscle contraction directly through their monosynaptic connections to motor neurons in the ventral horn of the spinal cord, or by inhibitory interneurons in the dorsal or intermediate horn, which results in the relaxing of the antagonistic muscle from which the signal was derived. Intact spinal proprioceptive circuits play an important role in the spontaneous recovery of animals who have undergone spinal cord injury. Ablation of proprioceptive signalling below, but not above, the spinal cord level of injury prevents the onset of spontaneous recovery (Takeoka and Arber, 2019). Proprioception is altered following ablation of corticospinal inputs and it has been hypothesised that it is the normalising of hyper-excited proprioceptive circuits that drives NT3's rehabilitative effect (Kathe et al., 2016).

The studies from our lab outlined in the earlier section (Duricki et al., 2016, Duricki et al., 2018b, Duricki et al., 2019) provide evidence which indicates that

intramuscularly or subcutaneously delivered NT3 heightens and accelerates a sensorimotor recovery after ischaemic stroke by acting in the PNS alone, but it cannot be ruled out that NT-3 also acts, at least in part, or fully, in the CNS. NT-3 can cross the CNS-blood barrier (Poduslo and Curran, 1996). The corticospinal tract is rich in TrkC receptors, more so than any other descending tract (Keefe et al., 2017), and local injections of NT-3 into the CST has been shown to increase its sprouting (Schnell et al., 1994). The CST is the major benefactor among all of the descending tracts in terms of the sprouting NT-3 induces (Grill et al., 1997). Therefore, it is possible that NT-3 is acting in the CNS to bring about the sensorimotor recoveries we have seen.

Another unknown is the receptor through which NT-3 binds. While NT-3 binds with highest affinity to trkC, it does still, although with weaker affinity, interact with trkA, trkB and p75^{NTR}; co-activation of trkC and one or more of these receptors may still be important to NT-3 induced sensorimotor recovery. Understanding which receptor will aid in the development of drugs which have higher serum half-lives than NT-3; a drug may need to bind to and activate one more of the trk receptors. Using the methods described in the next section, and whose experiments are described in the results Chapter 1, the following two working hypotheses are tested:

Hypothesis 1. NT-3 mediates the reversal of lost motor function after ischaemic stroke by binding to and activating TrkC receptors.

Hypothesis 2. NT-3 restores the loss of motor function after ischaemic stroke by acting on receptors in the dorsal root ganglia.

1.2.3 Tools to dissect NT-3's mechanism of action

Genetic mouse lines whose neurotrophic receptors trkA, trkB and trkC, could be manipulated both chemically and genetically have been created, each containing a point mutation within the exon coding for an intracellular ATP binding pocket (Chen et al., 2005). ATP binding of trkC is necessary for its kinase activity. The F617A mutation introduced in these mice in an ATP binding domain of Trk receptors allows

them to be reversibly and selectively inhibited by 1NMPP1 treatment with 1NMPP1 (a CNS-penetrant drug). The F617A mutation is innocuous in the absence of 1NMPP1, and mutant mice develop into healthy adults with no distinct phenotype when compared to their heterozygous or wildtype littermates. However, when pregnant homozygous mutant mice are treated with 1NMPP1, the offspring lack the proprioceptive circuits in the DRGs like previous knock-out studies of trkC (Chen et al., 2005). This method of inhibition allows us to test **hypothesis 1** and determine whether NT-3 mediates its rehabilitative effects through trkC. If trkC inhibition by 1NMPP1 prevents a functional recovery, it would provide evidence that it is trkC through which this effect is mediated.

These mice also contain loxP sites on exon 13 of the ntrk3 gene, flanking the F617A mutation, meaning trkC can be knocked-down in cells expressing cre-recombinase (Leonard et al., 2012). To address the question of where exactly NT-3 is binding to exert its rehabilitative effects and to test **hypothesis 2**, one can use this cre-loxP strategy. Large diameter, mechanosensory neurons express trkC in the most abundant quantity on the cell body and along the axons. Advillin, a protein involved in microtubule assembly is expressed in all sensory neuron subtypes (as well as sympathetic neurons and a small number of CNS neurons) (Hunter et al., 2018), and using its promoter to drive expression of tamoxifen-sensitive cre-recombinase (Cre^{ERT2}) means it is possible to selectively knock-down trkC in the sensory neurons of mice (Lau et al., 2011). Once the ischaemic stroke model is established and the behavioural assays made, we can use this strategy to test **hypothesis 2**.

1.3 Preclinical evaluation of rehabilitative drug candidates

1.3.1 Traditional assays

The assessment of therapies for the improvement of functional outcomes of stroke animals is typically achieved in three phases;

1. **Training:** Pre-training of animals on a complex skilled reaching task.
2. **Injury:** Induction of injury and providing post-operative care.
3. **Assessment:** Weekly assessments of performance on the pre-trained behaviour in which the treatment is assessed compared with vehicle control.

An *in vivo* experiment, designed as above, from start to finish, including analyses of data, can take between 4 - 6 months. Higher-throughput methodologies are needed.

An ideal behaviour test paradigm would be one that does not require training; a behaviour that mice or rats are innately able to perform; but one that is reliant on the motor cortex and the corticospinal tract, such that ablation of either of these systems leads to a persistent drop in performance, with a large window between pre and post-injury performance such that the therapy efficacy can be detected with relatively few animals per group. The animal will also need to be motivated to continue performing this behaviour, after injury, and for there to be little to no chance of a compensatory mechanism the animal could learn in order to achieve the goal of the behavioural task, or that compensation can be discriminated from true recovery of preinjury strategy.

Phases one and three take the most time and human labour and limits the number of studies that can be done at any one time, and the number of animals one can use in these studies, reducing the statistical power. The first phase can be made high throughput by use of behavioural assays in which mice and rats both spontaneously partake in without training; it is likely that there will be less variability in the performance of these tasks because of the omission of the training component. These might include, for example, the grip strength test [used after ischaemic stroke

in mice (Kilic et al., 2008)], the voluntary running wheel (Manzanares et al., 2018), more recently, a rope pulling test (Blackwell et al., 2018) and a dynamic assay which challenges learned motor function after injury (Lopes et al., 2017). While these tasks interrogate different aspects of motor performance and cannot and should not be used interchangeably, they are highlighted for their advantage over many other tasks in that they do not require as much time for animals to become competent.

While these behavioural paradigms circumvent the need for training, they still need to be carried out on a weekly basis through the assessment phase which is arguably the most difficult phase to scale and make higher throughput, as it depends on the efficacy of the therapy. It is the most rate limiting. Therapies that encourage plasticity in the spinal cord typically require a minimum of 4 to 5 weeks before this manifests as a measurable improvement in motor ability (Wahl et al., 2014). Animals are commonly assessed for up to 12 weeks post-injury. It is at this point animals are sacrificed, and their post-mortem tissues analysed for anatomical changes such as increased plasticity of spared tracts.

Apart from the sheer volume of time required to train and assess animals, there are also many other caveats in their manual training and assessment: (1) Animals are normally assessed during their light cycle, when they are least active (Ellens et al., 2016); (2) human handling of animals can affect their behaviour (Hurst and West, 2010); (3) animals are typically left untrained or unassessed on weekends, which can impact on their inter-day performance; and (4) the work required means that animals are not exposed to as much training one might require for experimental paradigms. Intensive rehabilitation of reaching and grasping, which has increasingly been shown to be an important adjunct to drug therapies both preclinically and clinically, requires a lot of work to achieve manually. A typical reaching and grasping trial will last 20 minutes where the rodent is presented with 20 pellets to retrieve; intensive rehabilitation paradigms can have as many as 100 trials in a session (Wahl et al., 2014). The problem is not the behavioural tasks themselves, which are often

excellent analogues of human movements, it is doing them at scale while maintaining data quality.

In this introductory section I will detail two behavioural tasks which, using robotics, machine learning, and computer vision, I have begun attempting to automate. I will detail the tasks, the issues faced when scaling these assays, and the challenges one must overcome in order for it to become automated. They are the horizontal ladder task and the single pellet reaching tasks.

1.3.1.1 Horizontal Ladder

The horizontal ladder task is used widely in preclinical trials for the assessment of proprioception, fine motor control and walking accuracy (Metz and Whishaw, 2009). Proprioception is the awareness of the position and movement of limbs in space and is critical for correct execution of the horizontal ladder task. An aversive and a homecage-like environment are bridged by a horizontal ladder with irregularly spaced rungs (Figure 3). Mice are motivated to cross the ladder as quickly as possible, and after two or three sessions of familiarisation, can cross in a timely manner, making very little to no foot slips. After injury to the cortex by ischaemia, the number of slips made by the paws contralateral-to-injury increases. The latency to cross the ladder also increases. With time [4-6 weeks (Duricki et al., 2018a)] rats recover to some degree and the number of slips made decrease. With rehabilitative treatments, this recovery can be accelerated. Rats are typically video recorded as they traverse across the horizontal ladder and the number of slips made by each paw are quantified offline, and in slow motion.




Figure 3 A typical horizontal ladder

Horizontal ladders bridge two environments; a neutral cage and a home cage. Mice are motivated to cross from the neutral to home cage and do so by crossing the horizontal ladder which may have regularly or irregularly spaced rungs. (image source: MazeEngineers.com)

An animal slipping on the horizontal ladder is a very quick event. It is necessary to watch captured video footage in slow motion to accurately determine the number of slips made by each paw. Also, while efforts have been made to standardise the quantifications of errors made by animals crossing the horizontal ladder (Metz and Whishaw, 2009), it is subjective. Manual quantification is also very time consuming; we estimate that scoring foot slips made from each paw by a mouse in a three-minute video containing four ladder traversals takes approximately 20 minutes. A typical study may have upwards of ten animals per treatment group, and behavioural outcomes are measured at weekly timepoints, sometimes for up to 12 weeks. This means that 130 hours are required to score all the videos in a preclinical trial consisting of 20 animals with 13 timepoints, including the baseline assessment. The scoring of this task would benefit from being automated. As this task is typically

recorded by video capture, it lends itself to being automated by the use of computer vision algorithms.

Single Pellet Reaching Task

An example of a skilled reaching assessment is the single-pellet reaching task (Whishaw, 1996). A pellet is placed on a shelf beyond a small aperture which the animal must reach through to grasp its reward. For the correct execution of this task, animals must be trained. Technicians place a mouse or rat into a clear Plexiglass box typically with a thin vertical slot (Figure 4). A plinth, or shelf, is appended to its outside with two concave wells offset laterally on either side of the shelf, where sucrose pellets can be placed. sucrose pellets placed in the divots can only be retrieved by the contralateral paw, which is advantageous as it means each individual paw can be assessed independently. This task has evolved from early versions: rather than a shelf, they had a tray containing multiple pellets. The use of a shelf, and singular pellet, enables researchers to more easily quantify “failed” trials, in which a pellet placed in the concave well is missed or dropped.



Figure 4 **A single pellet reaching task for mice**

The sugar pellet reward is separated by a Plexiglass with a thin slot from which that mouse can reach through (image source: Farr and Whishaw 2002).

The successful grasp of a sugar pellet is dependent on the correct coordination of multiple muscle groups, modulated by the firing of unique and sequential clusters of corticospinal neurons (Wang et al., 2017).

The movement can be subdivided into three stages:

1. **Extension**
2. **Grasping**
3. **Retraction**

The rat first senses the presence of a pellet by olfaction (Whishaw, 1996) (not visually, such as with humans - a blindfolded rat can still sense the presence of a sugar pellet) it places its snout through the slot and registers that a pellet is available. The head is then lifted slightly, allowing room for the paw to reach through the slot. It then extends its arm through the aperture, supinates the forearm, and closes its digits around the sugar pellet. It must then retract its arm while grasping the pellet, rotating the wrist once more to finally bring the sugar pellet to its mouth. This movement is incredibly transient, and lasts around a third of a second. In the early training phase, the ability for rats to reach for sucrose pellets is low. Many reach attempts are made (where a reach attempt is a repetition of the extension and grasping stages) and pellets are often missed or dropped. As animals become more familiar with the task the number of successful reach-and-grasps increases, and the number of reach-attempts required before a successful grasp is reduced. With training, rats will progressively gain reaching and grasping proficiency, but will not reach 100% success rates, which is indicative of the complexity of the task. A lesion that destroys corticospinal neurons lead to deficits in this reaching performance, but the motivation to continue attempting to reach for pellets persists after injury.

This single pellet reaching task is an excellent analogue for reaching actions in humans (Klein et al., 2012), and ablation of one side of the primary motor cortex leads to deficits in the rat's performance on the contralesional paw (Duricki et al., 2016). This task is used in preclinical trials to evaluate the efficacy of potentially rehabilitative drug therapies. A drug therapy is assessed on its ability to accelerate and increase the extent of recovery of lost motor function required by animals to execute the single pellet reaching task.

The single pellet reaching task for mice has not been as widely used as for rats (Farr and Whishaw, 2002, Wang et al., 2017). This might be due to the fact that mice do not readily reach for pellets and require many more training sessions for familiarisation. Food restriction prior to training is required to motivate mice to execute

the task, but their performance as determined by their number of successful reaches often does not surpass 50% (Farr and Whishaw, 2002), far lower than rats, which can reach up to 80% (Wahl et al., 2014). It is possible that to achieve the same success rates in mice, that more exposure to the training task is required.

Quantifying some aspects of the single-pellet reaching task is more straightforward than the horizontal ladder. A successful trial is typically said to be when the animal grasps the sugar pellet and brings it to its mouth, irrespective of the number of reach attempts made prior to the successful retrieval. This is usually done in real-time, and successes and drops are tallied by hand. Any additional metrics a researcher wishes to gather from animals executing this task are gathered from video footage, such as the number of reach attempts, and the trajectories of the paw in space. These additional metrics are often not pursued because of the time required to score the videos.

Behavioural assessments of motor performance are arguably the most important outcomes of these types of preclinical trial but to gather data is very time consuming. For these reasons, alternative solutions are sought. One such solution is automation; using robotics to deliver sucrose pellets, and bespoke software to analyse both new and historic video footage. With automation, one can make assessments of the single pellet reaching task more objective, less variable and would allow for more trials to take place.

1.3.2 Automating preclinical assessments of motor function

There are many commercially available software packages and devices available for purchase to automate many preclinical tasks. For example, EthoVision is a software package that can be used to automate the assessment of many behaviours through its live or offline tracking of animals. The DigiGate is a device that can be used to assess the gait of animals walking on a treadmill before and after injury. The ErasmusLadder is an alternative to the classical horizontal ladder task,

and is a hardware device with an array of pressure sensitive rungs and optical sensors which can quantify the errors made by an animal traversing across it. Due to the limited market, these solutions are often very expensive, and therefore inaccessible to many research groups.

1.3.2.1 Single Board Computers

Behavioural tests can be automated inexpensively using rapid prototyping methods. Recent advancements in the DIY maker space has made the creation of hardware more accessible to researchers in academic institutions. Low cost single board computers (SBC) such as the Raspberry Pi (Figure 5) or Arduino have enabled researchers to go from concept to prototype in a cost-effective manner. Raspberry Pis are a completely self-contained computer, whose current flagship model, the Raspberry Pi 4, costing as little as \$35 US dollars at the time of writing (including a Linux operating system). Their small footprint and low power requirements mean that they can be battery powered or placed in-cage and left unattended, and the myriad of communication protocols integrated within the board (for example: WiFi, Ethernet, BlueTooth, Serial and I₂C) means that data transfer can be achieved in a number of different ways.




Figure 5 A Raspberry Pi computer

A labelled diagram of a Raspberry Pi view from the top down. Raspberry Pi's are equipped with various connectivity ports and a header with tens of general purpose input and output pins that allows users to connect different sensors and actuators (image source: <https://www.seeedstudio.com/Raspberry-Pi-4-Computer-Model-B-4GB-p-4077.html>).

Raspberry Pis have already been used for this purpose by Tim Murphy's group in Alberta to automatically gather weight information from mice which are housed together (Noorshams et al., 2017), to gather data from mice who voluntarily dock their heads onto a device that can image their brains through a skull window (Murphy et al., 2016), and to automatically detect slips of mice traversing across a tapered beam (Ardesch et al., 2017). Raspberry Pi's provide researchers with the ability to query sensors, modulate actuators, capture images at high frame rates, and transfer data wirelessly and securely. For this reason, they are an excellent tool to automate preclinical assessments of motor function. In chapter 2 of my thesis I describe the use

of Raspberry Pi SCBs, controlling a camera, stepper motors and a limit switch to automate the single pellet reaching task.

1.3.2.2 3D Printing

In the last decade, the use of 3D printers has greatly risen. Once a very expensive prototyping method, it is now a commercially viable manufacturing process with fusion deposition modelling (FDM) printer technology being the most common and accessible. These printers work by melting, extruding and then re-cooling plastic, onto a bed to which it readily adheres, following a toolpath generated from the designed model. It is an additive manufacturing process: layer by layer, a model, designed using CAD software, can be made with high tensile strength, chemical resistance and high melting temperatures with a very quick turn-around time: parts can be printed within a few hours, enabling day-to-day iterations in prototype development. Spools of Polylactic Acid (PLA) filament are inexpensive, making this a cost-effective and fast manufacturing process for academics. 3D printing, together with Raspberry Pis, cheap sensors, and custom firmware provide the tools required to rapidly prototype bespoke devices.

1.3.2.3 Existing automated solutions

Many research groups have begun automating preclinical behavioural tests for the assessment of therapies that might restore lost motor function. Fenrich et al sought to automate aspects of the single pellet grasping task with a view to alleviate the work required for high intensity rehabilitation paradigms (Fenrich et al., 2016, Torres-Espín et al., 2018b). High-intensity rehabilitation is being shown both pre-clinically, and clinically to be an effective therapy to restore motor function, in isolation and in combination with other therapies, however, it is very difficult to achieve at scale. A rehabilitation cohort, for example, might require that rats undergo 100+ reaching and grasping trials on a daily basis, which would require many hours of human labour (Wahl et al., 2014). The automated pellet presentation (APP) device is a robotic

solution which removes the procedural weakness of manual training, will allow for overnight reaching, and can deliver a number of trials which exceeds that of a human.

The APP is accessible from an additional chamber attached to the rat's home cage, which has ad libitum access to food and water (

Figure 6). Pellets are dispensed on an arm and are positioned relative to a vertical slot very similarly to the manual task proposed by Whishaw. A mechanical barrier prevents the rat from prematurely reaching for a not-yet positioned pellet. The advantage of this device is that it could also vary the difficulty of the task depending on the rat's performance level. The limitation of this device is that they cannot be retrofitted to standard rat housing and that animals need to be individually housed, meaning this solution may not scale efficiently. Also, while this device did have an integrated camera and infrared LEDs to illuminate the reaching area, trials still required manual assessment by watching video footage back. It was very evident that this automated solution was superior to manual testing in that rats achieved more daily trials than they would have manually, most trials were carried out during the dark cycle, animals with access to the APP performed better on the task than manually trained rats, and that there was less inter-day variability in reaching performance by rats trained with the APP than traditional testing methods (Fenrich et al., 2016). This suggests that automated 24-hour delivery of pellets is superior to traditional methods.




Figure 6 **The Automated Pellet Presentation (APP) robotic system**

The APP is a device that is appended to a modified home cage and is accessible to singly housed rats by a tunnel. Rats in the task enclosure can be sensed from sensors positioned below, and are used to modulate when a pellet is presented to the rats. Pellets are accessible through a slit, visible from the front view. A pellet is positioned relative to the slot by an arm from the side and can be adjusted based on the preferred paw of the animal.

While the APP did have an integrated camera system, the video footage of trials was manually validated. There is potential scope to automate the performance of animals with the captured video frames using computer vision and machine learning algorithms. This is arguably a superior approach for the automation of trial scoring (as opposed to using physical sensors), as captured video, analysed automatically, can still be validated manually should it be required. Also, new computer vision algorithms can be retrospectively applied to stored video footage that might quantify behaviours in new and innovative ways.

Indeed, there have been methods and devices published such that some level of trial scoring automation is achieved with computer vision. In the device proposed by Ellens et al, the authors describe one of the limitations in using IR sensors to trigger video capture when a paw passes through the reaching slot: the rat often peers through the slot using its snout (Ellens et al., 2016). Infrared beams are indiscriminate of the feature occluding the beam, and therefore is not optimal for this purpose. Instead, the authors painted the preferred paws of the animals in green, and used computer vision algorithms (pixel-intensities) to monitor the presence of a paw beyond the slot. This solution was successful in triggering video acquisition. However, again, this would not scale up easily given the need to paint the preferred paws of animals and that this would have to be repeated on a weekly basis as the ink wears away. The advantage of the device is that it uses automatic pellet delivery, an array of mirrors and triggered video capture at very high speed (>300 frames per second) and

resolution to record trials. The quality of the video footage means that large amounts of data could be gathered from it. This solution's disadvantage is that animals are individually assessed in a non-homecage chamber, not 24/7, that the success of trials still need to be scored manually and that the preferred paws of the animal need to be painted with nail varnish. Also, the difficulty of the task could not be modulated.

Another device which aimed to automate the single pellet reaching task with computer vision was described by Wong et al in 2015 (Figure 7). A commercially available pellet dispenser was used, connected to two chutes, driving pellets to positions either left or right of the reaching window. With this device, unlike those described previously, pellet presence, and trial success was detected automatically in real time with an accuracy of 97.5% and 92%, respectively. Pellet placement was detected using infrared beams positioned in a top-down orientation. Beams were occluded, which indicated the pellet was present, when a pellet was successfully dispensed to its concave cut-out. This was validated using the integrated camera. Pellet drops were detected using computer vision, in real time. This was achieved with Image subtraction applied to a region of interest on the ramp, where a pellet would roll down into a reservoir, should it be displaced and not eaten (Wong et al., 2015). The weakness of this device is that animals need to be individually housed, and the devices are a desktop and therefore cannot be used 24 hours a day, 7 days a week.




Figure 7 The Behaviour Box proposed by Wong et al in 2015 uses robotics and computer vision to automate the single pellet reaching

(A) Front view of the Behaviour Box. Pellets are detected in their divots by top down infrared LED beams. Pellets are placed there by chutes which are attached to the mechanical barrier. A side-view camera is positioned to monitor the presence of a pellet. (B) The ramp, outlined in green, has a high contact against the falling pellets which enables the presence to be detected with image subtraction methods.

Individual housing is advantageous as it ensures that each animal participates in the task equally, however, its limitation is in its scaling: training and assessing as many animals as possible. Also, singly housing animals for long periods of time is not ideal for their welfare.

A former PhD student in our laboratory, Dr. Dhireshan Gadiagellan, used a Raspberry Pi as the basis to begin automating the single pellet reaching task for rats. The goal was to create a device that would require as little supervision as possible, that could run unattended in the home cages of animals, and could differentiate between animals using the device such that animals need not be singly housed.

Gadiagellan's development lead to the inception of the "RatBot" (Figure 8), a Raspberry Pi powered device programmed in C Sharp, that using a set of actuators and sensors, could train and assess rats on the single pellet reaching task. That

RatBot's hopper has a capacity of one thousand 45 mg sucrose pellets. Within the hopper is a stepper-motor controlled augur that drives pellets, in single-file, up and over the lip of the hopper, into a chute directing the pellets to a "spoon". The spoon is equipped with a proximity sensor which detects the presence of a pellet. Once a pellet has dispensed, the spoon can then be positioned relative to the slot in various positions depending on the performance level and handedness of the animal executing the task.

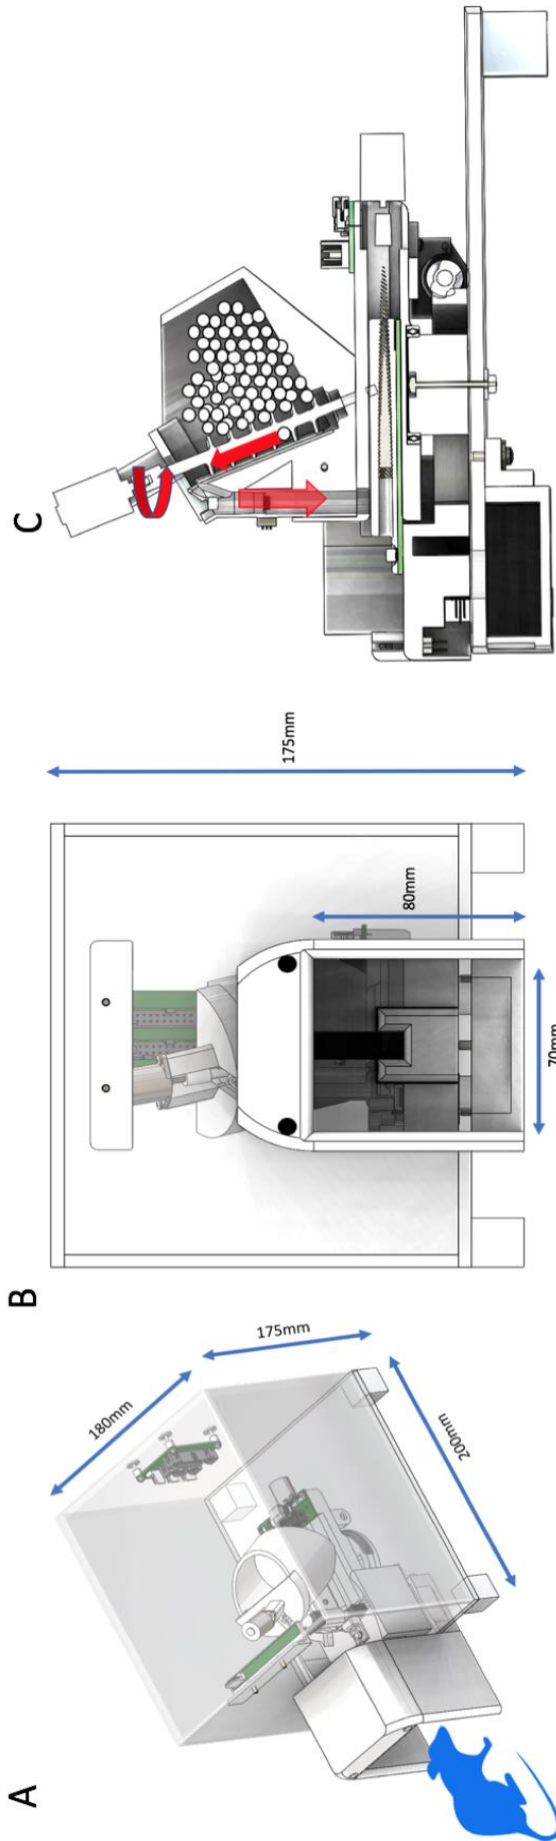


Figure 8 The RatBot as designed by Dhireshan Gadiagellan.

(a) An isometric view of the ratbot, which is an in-cage device that can operate with multiple animals. The rat is recognised by the RatBot by the radiofrequency identification (RFID) reader situated within the tunnel positioned around the reaching window. (b) The tunnel restricts access to the slot to one animal, ensuring that only the correct implanted RFID tag is read when a rat executes the task. (c) A cross-sectional view of the rat bot's dispenser mechanism which uses an augur to drive pellets up and over a hopper into a chute, leading to a spoon, that presents the pellets to the animal.

It can move both away from the slot and laterally, and even beyond the slot into the cage, for the training phases of preclinical studies so that rats can reach the pellet initially with their tongue; the task is made progressively harder by retracting the spoon further from the slot. Rats, each implanted with glass radio frequency identity (RFID) capsules, approaching the slot are identified by an RFID reader located above the slot. The number of reach attempts are quantified by the number of breaks of an infrared beam across the reaching slot. A drop of the pellet is recorded when a beam break occurs in either of the two drop-zones located below the spoon accompanied by the absence of a pellet on the spoon as determined by the proximity sensor. Conversely, absence of a pellet from the spoon without a break in either drop-zone beams indicate the successful retrieval of a pellet.

Gadiagellan showed that an array of RatBots were able to train a cohort of animals to reach for single-pellets in difficult-to-reach positions, achieving daily trial numbers that would require many hours of human labour to achieve (100s rather than 20). The RatBot also detected a drop in performance when the animals underwent photothrombotic surgery. The RatBot proved to be a viable solution for the automation of the single pellet reaching task in rats. The RatBot solved the single housing issue with RFID technology.

Currently, an automated single-pellet-reaching device does not exist for mice.

A device for mice would ideally be capable of the following:

1. The ability to position pellets relative to the slot with varying levels of reaching difficulty
2. The ability to record and store high quality video footage capture with a high frames per second
3. Can be retrofitted to standard mouse housing
4. Does not require that animals are singly housed
5. Can run for a long period of time and unattended
6. Transmits data wirelessly such that it is accessible from the cloud

1.3.2.4 Animal Pose-Estimation using Deep Convolutional Neural Networks

Many animal behaviours are video recorded so outcome measures can be gathered “offline” in ways that cannot be done in real time, such as those previously described; kinematic analysis, number of reach attempts made by animal executing the single pellet reaching task; the number of slips made by animals traversing a horizontal ladder. Manual video analysis is subject to human bias and requires a lot of time to analyse. Many software packages exist to help researchers automate the analysis of their behavioural assays. One notable software package is EthoVision (Spink et al., 2001), which is licensed software that allows users to gather metrics on behavioural assays such as open field locomotion, or the water maze test. Object detection and tracking algorithms track mice and rats as they locomote around their arenas, and can output metrics such as the distance covered in a given time, or the amount of time spent by animals at the edges of an open-field environment (as a measure of anxiety).

Until recently it was very difficult to non-invasively automate the kinematic analysis of animals carrying out reaching and grasping behaviours, and the same is true for the horizontal ladder. Pose-estimation is a field of computer vision whereby the posture of humans, and now other animals, can be estimated. More specifically, pose-estimation can be used to locate various features of a human in video footage.

This has been used widely in different consumer applications and implementation by researchers has also recently become easier. Pose-estimation works by applying a trained model, often a deep convolutional neural network, to individual frames of a video and having it estimate, or predict, the whereabouts of a feature, such as a hand or digit in a frame. The model is trained based on manually annotated frames. An estimation of the whereabouts of body parts provides researchers with a non-invasive, inexpensive way to automate the analysis of video footage.

Recently, a toolbox called “DeepLabCut” was published which is revolutionising behavioural neuroscience (Mathis et al., 2018). It allows users to train deep convolutional neural networks to predict the whereabouts of a body part in a frame of video footage (see mousemotorlab.org). The team behind DeepLabCut took the neural network architecture (ResNet50) of a model which was trained to classify images from the ImageNet dataset (Deng et al., 2009) and using a further training dataset with the annotated joints of humans, applied transfer learning to the ResNet50 model with these additionally annotated frames (Insafutdinov et al., 2016). This accurately predicted the pose of humans. DeepLabCut’s innovation is the application of transfer learning to create animal pose-estimation models and the release of an open-source toolbox that allow others to do the same. This means that users need only annotate as little as 200 frames of an animal doing its behaviour for a model to accurately predict its pose. Once a network is trained, videos separate to those used in the training dataset can be fed through the network such that coordinates of body parts within a frame are predicted. Rows of predictions are generated for each frame in a video, containing the x and y coordinates alongside a “likelihood” value. The likelihood value lies between 0 and 1 where one is most ideal; it indicates that the predicted coordinate is very likely to demarcate the feature of interest. These predictions provide the basis for the automation of behavioural assays and their scoring. In chapters 3 and 4 I describe how I am using DeepLabCut to automate the

classification of the trial outcomes from the MouseBot, and the scoring of the horizontal ladder.

2 Understanding how NT-3 restores the loss of motor function

2.1 Introduction

Stroke is a leading cause of disability world-wide. Up to three quarters of stroke survivors will leave hospital with a motor impairment (Lawrence et al., 2001), which makes their daily living activities very difficult and reduces their quality of life. Currently there are no therapies available to reverse the loss of motor function, other than physical rehabilitation with or without electrical stimulation which are intensive, require specialists and dedicated facilities. New therapies are desperately required.

Our lab and others have shown that administration of exogenous neurotrophin-3 (NT-3) 24 hours following induced ischaemic stroke, both in adult and elderly rats, accelerates the recovery of their upper limbs (Duricki et al., 2016, Duricki et al., 2019). Further, subcutaneous infusion of NT3 for one month starting 24h after large cortical strokes in rats results in progressive recovery of sensorimotor function (Duricki et al., 2018b). NT-3 is a member of the neurotrophin family and plays key roles in neural differentiation, migration, and activity-dependent synaptic plasticity during development (Skaper, 2012, Tessarollo, 1998). We hypothesise that exogenous NT-3 promotes the recovery of lost motor function by activating its highest affinity receptor, trkC (Lamballe et al., 1991, Tessarollo et al., 1993). While NT-3 can cross the blood brain barrier (Poduslo and Curran, 1996), trkC receptors are expressed in both the CNS and PNS (McMahon et al., 1994, Mu et al., 1993, Muragaki et al., 1995, Takeo et al., 2003), and our data suggest that transport of NT-3 from the periphery to the CNS may not be necessary to reverse the loss of motor function after stroke; a viral vector encoding human NT-3 injected into the affected forelimbs of rats

improved upper limb function after stroke despite no detectable elevation of NT-3 in the CNS (Duricki et al., 2016). Conversely, earlier work from Professor David Shine's lab showed that when a viral vector encoding NT-3 is introduced into the sciatic nerve of rats whose CSTs had been unilaterally severed at the level of the medulla, contralateral to the injury, it encouraged plasticity of the intact CST fibres across the spinal midline (Zhou et al., 2003). If lost motor function can be restored by NT-3 solely activating upon targets in the periphery, then therapies that mimic NT-3 need not be designed to cross the blood brain barrier, mitigating one of the key challenges in drug design. TrkC receptors are expressed in and upon the large diameter neurons of the dorsal root ganglion (Mu et al., 1993) and these primary sensory neurons are important for mechanosensation and proprioception as evidenced by knockout studies (Klein et al., 1994, Tessarollo et al., 1994). It is possible that NT-3 promotes recovery after stroke through its interactions of trkC receptors on the axons or cell bodies of large diameter neurons of the DRGs and subcutaneously infused NT-3 could accumulate in the DRG from the bloodstream as the DRG has fenestrated capillaries through which molecules as big as NT-3 can cross (Jimenez-Andrade et al., 2008).

We are yet to know through which receptor it is that NT-3 interacts with to bring about this rehabilitative change. NT-3 also binds to and activates trkB receptors (Fariñas et al., 1998), which may be an important interaction in NT-3 induced motor recovery. Studying NT3-trkC signalling has historically been difficult because a selective inhibitor of trkC has not yet been synthesised, however, genetically modified mice were created whose high affinity neurotrophin receptors can be genetically and pharmacologically modulated (Chen et al., 2005) and have been used to study the roles of neurotrophin signalling in various models of health and disease (Gill et al., 2016, Mantilla et al., 2014). The TrkC^{F617A} knockin mice used in this study have a phenylalanine to alanine mutation within the ATP binding pocket of kinase subdomain V (amino acid position 617) of their trkC receptors, rendering one of their kinase

domains sensitive to inhibition by a small (CNS-penetrant) chemical compound, 1NMPP1. This inhibition prevents trkC from phosphorylating itself and its downstream targets (which occurs after NT3 binding). Furthermore, the gene sequence containing this mutation was created with flanking loxP sites, meaning trkC can be knocked down both spatially and temporally when cre-recombinase is introduced. With trkC^{F617A} knockin mice it is possible to explore whether NT-3 reverses a loss of motor function following stroke through its interactions with trkC using a temporal and conditional inhibition strategy.

Accordingly, we first characterised photothrombosis in mice behaviourally and anatomically, explored whether 1NMPP1 inhibited a spontaneous recovery in homozygous trkC^{F617A} mice, and treated mice with delayed subcutaneously delivered exogenous NT-3 with a view to build upon work carried out previously using rats in the lab. As spontaneous recoveries are observed in the clinic (Langhorne et al., 2009) and in rodent models of spinal cord injury and stroke (Tennant, 2014), and NT-3 accelerates recovery following lesioning (Bradbury et al., 1999, Duricki et al., 2016, Kathe et al., 2016), we hypothesized that endogenous NT3-trkC signalling may also play a role in spontaneous recovery. Spontaneous recoveries are lesser in the elderly (Franceschini et al., 2018), NT3 levels in serum also reduce with age (Murase et al., 1994) and intact proprioceptive circuits are required for the spontaneous recovery of motor function after spinal cord injury (Takeoka and Arber, 2019). NT-3 plays a crucial role in orchestrating the correct wiring of proprioceptive circuits during development (Tessarollo et al., 1994), and so it is possible that NT-3 influences these same circuits following injury and that this influence promotes a recovery. To test this hypothesis, we trained mice to reach for sucrose pellets using the Montoya staircase, and measured their grip strength, before inducing a photothrombotic stroke in the hemisphere contralateral to their preferred paw. We then treated trkC mutant mice with 1NMPP1 for 6 weeks following injury and monitored how they performed on the pre-trained tasks. We also hypothesised that NT-3, delivered subcutaneously 24

hours following photothrombosis, would restore the loss of motor function of mice trained to cross a horizontal ladder and reach for sucrose pellets in a Montoya staircase.

2.2 Methods

2.2.1 Genotyping and confirmation of mutation

Ntrk3^{tm1Ddg} (*TrkC* fl/fl, for homozygous floxed mice containing F617A mutations) mice were obtained from Jackson Labs and genotyped as per their instructions. Briefly, this was done using forward and reverse primers 5' GCACTTCTGCTAAACCTTTGGT 3', and 5' ATCGTCCATCCCCATGTCT 3', respectively. Ear clippings were taken from mice typically at postnatal day 21. DNA was extracted by placing clippings in 75 µl 25mM NaOH and 0.2 mM of EDTA buffer, heating at 98 degrees Celsius for 1 hour, before adding an equal volume of 40 mM Tris HCL (pH 5.5).

Confirmation of the amino acid substitution in the *TrkC*^{F617A} knockin mice was achieved by PCR amplification of a 536 bp sequence containing the F617A mutation. Primers were designed using primer-BLAST and were 5' GTCTCACGGGTCTTCCACTT 3' and 5' AGAAGACTCCCACGGCTCA 3' for the forward and reverse primers respectively. To confirm the presence of the F617A mutation, amplicons were purified using a Qiagen PCR clean-up kit before sequencing by Source Biosciences (Cambridge), the forward primer was also used as the sequencing primer. The PCR cycling parameters used to generate the final, sequenced amplicons, were: 94 °C hold for 5 minutes, followed by 94°C, 55°C, 72°C cycle temperatures for 30, 45, and 45 seconds respectively. After 30 cycles, a 5 minute 72 °C elongation was set, before cooling to 4°C. All PCR procedures were carried out with a PCR Master Mix obtained from Promega (catalogue number M7502). All gel electrophoresis procedures for genotyping were done using a 1% agarose gel and ethidium bromide (1:20,000) with TAE buffer.

2.2.2 Dosing with 1NMPP1

The preparation of drug and behavioural assessment was carried out by the same investigator, but blinded to each drug group by a second investigator. Mice were randomly allocated to treatment group by stratified sampling to ensure both groups contained equal numbers of male and females. 1NMPP1 was given to mice via their drinking water in a 0.3% DMSO solution. Treatment initiation occurred three days before training and assessment following the wash-out period. 1NMPP1 was purchased from Millipore, (catalogue number 529581-1MG, lot number, D00169919). 1NMPP1 was dissolved in 100% DMSO to achieve a stock concentration of 100 mM and stored in -20 degrees C until use.

1NMPP1 is not stable in aqueous solution for long periods (Millipore datasheet), thus, fresh solution was made every 3 days. To make 500 ml of 0.3% 1NMPP1 solution for mouse consumption, a 2 ml 2.5% solution of Tween-20 was first made, 160 ul of either 1NMPP1 or control (DMSO) was then added to this volume, before making this up to 500 ml of distilled water. Mice were singly housed so that consumption of their drinking water could be monitored.

2.2.3 Photothrombotic stroke

For the 1NMPP1 study, a total of 20 TrkC fl/fl male and female mice weighing 25 - 30 g aged between 8-12 weeks old were maintained on a 12-hour light and dark cycle with *ad libitum* access to food and water. Mice were given additional bedding, and acclimatized to DietGel for postoperative recovery prior to surgery. For the NT-3 study, 27 trkC fl/fl mice, male and female were used. All procedures were in accordance with the Animals (Scientific Procedures) Act of 1986. All protocols involving animals received prior approval by the King's College London Animal Welfare Ethical Review Board and were authorized by the UK Home Office Project (license number 70/7865, held by Dr. Lawrence Moon).

On the day of surgery, mice were weighed, anesthetized with 5% isoflurane, and maintained at 2% in an oxygen/air mixture through a nose cone attached to the stereotaxic frame. Fur was shaved above the skull and skin was swabbed using antimicrobial solution (4% chlorhexidine gluconate). Viscotears (Boots, UK) was applied to the eyes of the mice to prevent drying. A midline incision was made along the scalp from above the nose to between the ears, and skin was retracted using weighted clips. The periosteum was removed with the edge of a scalpel. Photothrombosis was confined to a 15 mm² rectangular window (5 x 3 mm). The window was positioned such that both forelimb and hindlimb representation was exposed to illumination by the cold light source used. An aluminium stencil was placed to expose a rectangular region of skull located from 0 to 3 mm laterally (hemisphere contralateral to preferred forelimb), and from 2.5 mm anterior to -2.5 mm posterior to Bregma to confine irradiation to the motor cortex. Strokes were unilateral and the side lesioned was contralateral to the preferred paw as determined from the staircase task. An optic fibre with a 6 mm aperture attached to a KL 2500 LED light source (Schott) set to produce a light intensity of 20 klux as determined by a light meter attached to the fibre optic at a set distance (120 mm) using a 3D printed adapter. The optic fibre was placed using a stereotaxic arm in light contact with the aluminium stencil and, once in place, with the light turned off, 30 mg/kg of 95% Rose Bengal (Sigma-Aldrich, catalogue number: 330000, diluted in saline) was injected intraperitoneally. Rose Bengal was allowed to circulate for 5 minutes, then photoactivated by turning on the light for 10 minutes. Saline was applied to the skull following the irradiation period and the incision was sutured. Animals were then placed in an incubator heated to 33 degrees Celsius until conscious.

2.2.4 Histology

2.2.4.1 Quantification of CST ablation

The fixed spinal cords were cut transversely using a cryostat with a thickness of 20 microns. Sections were immediately mounted to slides. Before undergoing immunolabeling with rabbit anti-PKC gamma (1:200/5µg/ml, Abcam, catalogue number: ab71558), sections were placed in an oven for 45 minutes and heated to 50 °C to maintain section adhesion during incubation steps. Sections were blocked with 10% donkey serum for 1 hour, incubated with rabbit-anti-PKC gamma (Abcam, ab71558) overnight at room temperature, washed, incubated with the secondary antibody, donkey anti-rabbit IgG Alexa Fluor 488 (Abcam, catalogue number: ab150073) (1:1000), for 1 hour and 30 minutes, washed again and then coverslipped with Vectashield containing DAPI. All washing steps were done with PBS and 0.2% triton X. Images were taken using an Apotome fluorescence microscope, and the integrated density of PKC gamma immunolabeling in each dorsal column was measured using ImageJ.

2.2.5 Behavioural assessment

2.2.5.1 Grip strength

Triplicate measurements of mouse grip strength were taken using a bilateral grip strength meter (BioSeb, model: Bio-GS3) and averaged before photothrombosis. These values were used as the baseline score of performance. Following photothrombosis, grip strength measurements were taken on a weekly basis.

2.2.5.2 Montoya Staircase Test

Mice were trained to reach and grasp on the Montoya staircase. Mice were habituated to the sucrose pellets (20 mg, DateSand, #F05301) and staircase devices one week before initiation of training by leaving them overnight in their home cages. Before training and assessment days, mice were deprived of food for 20 hours to increase their motivation to grasp and then allowed to feed *ad libitum* after testing.

Their weights were monitored daily. Staircases were baited with two sucrose pellets per step on both sides. Training and testing sessions lasted for 30 minutes. Upon termination of the 30 minutes trial period, the number of pellets from each side of the staircase were counted and recorded. The maximum reach was also determined, which was the lowest step containing less than two pellets. This is an indication of gross reaching ability, and not fine motor ability or dexterity. Mice were considered trained when they were able to grasp 8 pellets from either side of the staircase and any mice who did not reach criterion were excluded from analysis. Baseline performance was recorded and their preferred paw was determined.

2.2.6 Subcutaneous dosing of NT3

24 hours following induction of photothrombotic stroke, an osmotic mini-pump (Alzet Mini Osmotic Pump, model no. 1004, product code: 9922) which infused either 587.5 ug/kg/24h of recombinant human NT3 (267-N3, Peprtech) or vehicle control (0.1% BSA in 20mM sodium acetate 140mM sodium chloride pH 5.0) for 28 days. Mice were anaesthetised and an incision was made such that osmotic mini pumps could be placed between the shoulder blades, with the outlet of the device facing away from the incision.

2.2.7 Detection of NT3 levels in serum and DRGs

In a separate cohort of naive mice, osmotic mini-pumps were inserted as detailed above, with the purpose of measuring the bioavailability of NT3 in blood serum and the DRGs. Four weeks after insertion of the pumps, mice were euthanized with an overdose of pentobarbital. Approximately 500 ul of blood was taken from each mouse by incising the heart's right atrium with scissors and allowing blood to drain into the cavity, before collection with a syringe. The mouse was then perfused with PBS and DRGs were dissected from the cervical and thoracic cord levels 2-6, and 1-4 respectively. These were flash frozen using liquid nitrogen.

2.2.7.1 Serum

Blood samples were kept overnight at 4 °C and then centrifuged at 8000 x g at room temperature for 15 minutes. The serum was then collected and stored at -80 °C. Measurement of total protein in serum (BCA) was done according to the manufacturer's instructions (BCA protein assay kit, Novagen catalogue no. 71285-2). The ELISA assay was also performed according to the manufacturer's instructions, except for the incubation step with TMB One-Step Substrate which was done for 25 minutes (instead of 30). The kit used was the NT-3 Human ELISA kit, Abcam, AB100615. Samples were diluted 1 in 2 using diluent A, according to the manufacturer's instructions.

2.2.7.2 DRGs

Protein extraction from cervical and thoracic DRGs from right and left sides were pooled, snap frozen after dissection and kept at -80°C until protein was extracted. Frozen DRGs were homogenized for 1 minute in tubes containing 300 µl of extraction buffer and protease inhibitor using a GentleMACS dissociator (Miltenyi Biotec GmbH). The tubes were centrifuged (3000 x g, 10 min, 4°C), and the supernatant was transferred to eppendorf tubes and further centrifuged (8000 x g, 5 min, RT). The supernatant was collected and stored at -80 °C. The tissue extraction buffer was from Invitrogen, FNN0071. (BCA protein assay kit, Novagen catalogue no. 71285-3). The ELISA assay was also performed according to the manufacturer's instructions (sample dilution 1:3 in PBS, except for the incubation step with TMB One-Step Substrate which was done for 25 minutes instead of 30). The kit used was the NT-3 Human ELISA kit, Abcam, AB100615.

2.2.8 Automatic analysis of horizontal ladder videos

Videos of mice crossing the horizontal ladder were analysed using LadderScorer V1 to quantify their latency to cross (this is described in detail in chapter

4). Videos were organised in directories according to the time point from which they were taken. Each video contained three or four ladder traversals.

2.2.9 Statistical analysis

All longitudinal behavioural data were analysed using SPSS's MIXED procedure and residual maximum likelihood (REML) estimation to accommodate for missing data points (due to mice dying before pre-planned experimental endpoints) (Duricki, Soleman, and Moon 2016). The covariance structure with the best fit, reflecting the most accurate model that represents the variation of values at each time point, was determined to be "compound symmetry" which had the lowest Akaike Information Criterion (AIC) value of all models tested (AR(1), and unstructured). For all outcome measures tested, the fixed factors included in the final model were the treatment group (NT-3, PBS, 1NMPP1 etc), time and the group and time interactions.

2.3 Results

2.3.1 Motor cortex and corticospinal tract ablated unilaterally following photothrombotic stroke

The goal of our lab is to restore the loss of motor function experienced among survivors of ischaemic stroke. Therefore, we opted to use photothrombotic stroke as our injury model due to its ability to reproducibly induce infarcts in terms of both its size and location (and therefore consistently lesion the motor cortex) and because survival rates after this surgery are high (Figure 9). 20 mice underwent photothrombotic stroke of the motor cortex following training on the Montoya staircase task.




Figure 9 Using the photothrombotic model to unilaterally ablate motor cortex.

- (a) An outline of the steps involved to induce photothrombotic stroke. (1) a midline incision is made to expose the skull, (2) Rose Bengal is injected and allowed to circulate for 5 minutes, (3) a fibre optic is positioned at the desired location and the brain is illuminated (through the intact skull) for 10 minutes, (4) the skin is then sutured and the mouse is placed in a heated incubator for recovery.
- (b) Left panel: a mouse in a stereotactic frame having just undergone photothrombotic stroke on the right hemisphere. There is a clear discoloration of the skull. Right panel: a perfused brain 6 weeks following induction of photothrombosis. Red rectangles denote the position of the aluminium stencil measuring 5 mm x 3 mm. The red scale bar is 3 mm.

Photothrombosis was tolerated well by mice, with a survival rate of 95% (compared to 65% in MCAo surgery) (Chen et al., 2005, Shanbhag et al., 2016). Photothrombosis has been shown to produce consistent lesion volumes across experimental animals: infarcts peak in size 2 days after surgery (Li et al., 2014). Analysis of the spinal cord sections (cervical level 3) 6 weeks post injury with PKC gamma staining, which labels the corticospinal tract in the dorsal column, revealed an 82% reduction in PKC gamma intensity on the side contralateral to the injury ($P < 0.05$; two-tailed t-test) (Figure 10).




Figure 10 **Histological analysis of the brain and spinal cords of mice 6 weeks following photothrombotic stroke**

(a) A coronal section of a cresyl violet stained mouse brain 6 weeks following photothrombosis the corpus callosum can be seen collapsing into the lesion centre and the ventricle beneath the lesion is enlarged, encroaching into the space where the cortex once was.


(b) A transverse cervical perfusion-fixed section of a spinal cord 6 weeks following photothrombotic stroke. Sections were immunostained with PKC gamma which is a marker of the corticospinal tract within the dorsal column. Quantification of the integrated density of PKC gamma staining showed an 82% reduction of intensity

of the contralateral-to-injury CST ($n = 11$, $p < 0.05$, two tailed t-test).

2.3.3 Confirmation of the F617A mutation in the the TrkCF617A mice

Homozygous $\text{trkC}^{\text{F617A}}$ knock-in mice can be differentiated from their wildtype littermates by exploiting the FRT and loxP sequences inserted along with the point mutation, which reside in the adjacent introns, and are approximately 100 bp in length. The mice used to begin the colony were heterozygous for the Trk F617A point mutation, which we paired for breeding, and approximately 25 % of their litters were homozygous for the mutant allele. Genotyping was done by extracting genomic DNA from ear clippings obtained from the litters and using primers designed to flank the location of the double 100 bp sequences, creating amplicons that could be visualised with gel electrophoresis. DNA amplicons with the F617A mutation have the additional 200 bp, and are 600 bp in length. Meanwhile, amplicons generated from wildtype mice are 400 bp. Figure 11a shows an example of an agarose gel imaged following gel electrophoresis of 10 mice (labelled arbitrarily from 5 - 14) alongside an amplicon from a mouse known to be homozygous for the F617A mutation (far right-hand side of panel). In this particular gel, 6 mice were wildtype, while 4 were heterozygous. To confirm the presence of the F617A mutation, we designed primers that flanked its position using Primer-BLAST. Sequencing of the amplicon obtained from the wildtype mice confirmed the presence of phenylalanine against the known mouse trkC sequence (Figure 11b), while sequencing of the amplicon obtained from $\text{TrkC}^{\text{F617A}}$ knock-in mice confirmed the presence of this F617A mutation (Figure 11b).

a



b




Figure 11 **Sequence alignments of the PCR amplicons obtained from *Ntrk3t^{m1Ddg}* mice homozygous for the F617A mutation and their wildtype littermates against the known sequence**

- (a) An example of an agarose gel image containing PCR amplicons obtained from the genomic DNA extracted from the ear clips of a *Ntrk3^{m1Ddg}* litter (their IDs were 5 - 14, alongside the positive control on the far right, which was known to be *TrkC^{F617A} +/+*). The ladder used was a 100 bp ladder. The mutant allele is contained within bands of approximately 600 bp in length. Additional amplicons were created from the mice contained within the red rectangles and Sanger sequenced to confirm the presence of the F617A mutation.
- (b) PCR amplification and subsequent sequencing of DNA obtained from a *TrkC^{F617A} +/+* mouse and its wildtype littermate (top sequence), compared to the amino acid sequence obtained from the known *TrkC* gene sequence (lower sequence): F617 was present in the wildtype, as expected (as highlighted by the blue rectangle). In the PCR amplicon obtained from *Ntrk3t^{m1Ddg}* mice homozygous for the 600 bp-length band, the F617A mutation was confirmed (highlighted by the red rectangle).

2.3.4 The effects of photothrombosis on motor performance is long lasting, and mice do not spontaneously recover

One strategy to show that *trkC* might be the receptor through which NT-3 mediates a functional recovery after stroke is to treat homozygous *trkC*^{F617A} knockin mice with NT-3 in combination with 1NMPP1 and monitoring whether 1NMPP1 prevents an NT-3 induced functional recovery (as seen previously in rats). Additionally, mice and rats spontaneously recover after stroke and spinal cord injury (Bachmann et al., 2014, Duricki et al., 2016) and intact proprioception may be important for this recovery (Takeoka and Arber, 2019).

As NT-3 is made by muscle spindles and other cells in the CNS and PNS (Shimazu et al., 2006b) and is critical for the development of circuits that mediate proprioception during development (Tessarollo et al., 1994), it might also be possible that endogenous NT-3 signalling may play a role in the spontaneous recoveries seen in rodent studies. While rats are known to spontaneously recover after large cortical ischaemic stroke, there is less data to suggest mice do the same. It has been demonstrated that mice do partially recover motor function of the affected forepaw after unilateral large photothrombotic cortical stroke 4 weeks after injury when tested using climbing assays, rotarod, rope climbing and grid walking (Bachmann et al., 2014), but whether they recover reaching and grasping function is unknown.

In an initial pilot study, we first wanted to study whether the silencing of *trkC*'s kinase activity in adult mice prevented a spontaneous recovery, in the absence of exogenous NT3, and therefore whether endogenous NT3-*trkC* signalling was an integral part of the spontaneous recovery seen in rats following brain or spinal cord injury. A preference test between water bottles of vehicle (which contains DMSO) and 1NMPP1 found mice did not find one or the other more palatable [data not shown]. Accordingly, *TrkC* mutant mice were given 1NMPP1 in their drinking water (containing 0.3% DMSO) throughout the assessment phase.

2.3.5 Mice do not spontaneously recover their grip strength or ability to reach and grasp for sucrose pellets using the Montoya staircase following photothrombosis

Mice (n=10 per group) were lesioned unilaterally on the brain hemisphere contralateral to the animal's preferred limb, as determined by staircase testing. A 3-day post-injury (DPI) measurement was taken on the staircase and bilateral grip strength, and their performance was assessed weekly for 6 weeks following stroke. There was a deficit in staircase and grip strength performance following photothrombosis, which dropped to around 50% of pre-injury performance. This deficit was sustained for the 6 weeks of assessment with no spontaneous recovery of either groups in both outcome measures tested (Figure 12). Analysis of the longitudinal data obtained for the staircase test, after controlling for baseline differences in performance, revealed that there was no effect of time ($F_{6,42} = 0.602$, $P = 0.727$), no effect of treatment, ($F_{1,7} = 0.296$, $P = 0.603$), and no interaction between time and group ($F_{6,42} = 0.820$, $P = 0.561$) (Figure 12). Analysis of the data for bilateral grip strength after controlling for baseline differences in performance revealed that there was no effect of time ($F_{6,90} = 1.65$, $P = 0.143$), no effect of treatment, ($F_{1,15} = 0.296$, $P = 0.595$), and no interaction between time and group ($F_{6,90} = 0.904$, $P = 0.496$). To summarise, unexpectedly, we did not see a spontaneous recovery in either treatment groups, and this meant we were unable to determine whether inhibition of TrkC kinase signalling using 1NMPP1 affected spontaneous recovery of grasping in mice after stroke. Also, 1NMPP1, did not reduce grip strength or the ability for mice to reach and grasp for sucrose pellets either throughout the assessment phase.

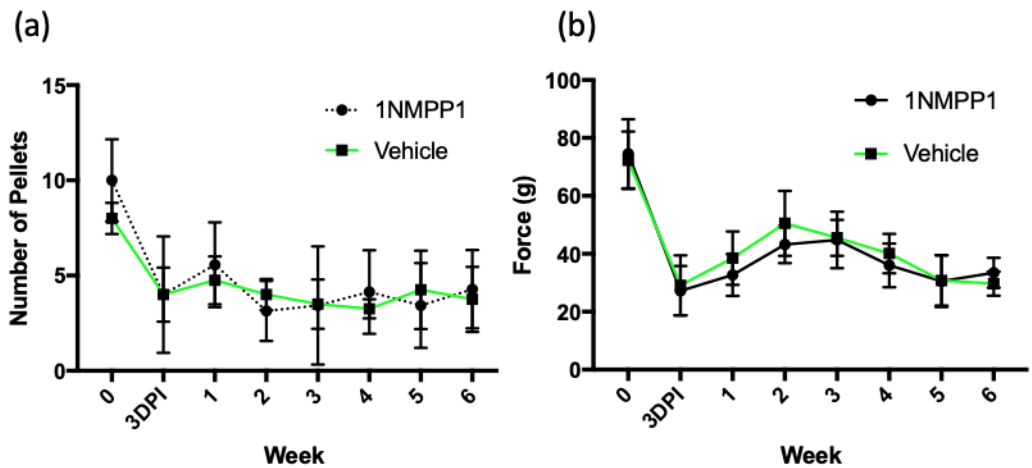


Figure 12 Mice do not spontaneously recover performance on pellet reaching and bilateral grip strength 6 weeks following injury

(a) Following photothrombotic stroke, the ability for mice (n=10 per group) to grasp for sucrose pellets using their preferred paw dropped by approximately 50% and this deficit was sustained throughout the 6 week testing period.

(b) Grip strength was also assessed weekly following photothrombosis. Bilateral grip strength was again reduced by approximately 50% following injury, and this was sustained over the testing period with no recovery or divergence of groups.

2.3.6 Subcutaneously delivered NT3 following injury modestly alters reaching performance

Given the control group did not spontaneously recover on either behaviour task (Montoya staircase and grip strength), which are known to happen in rats, we next wanted to confirm whether subcutaneously administered NT3 is rehabilitative in a separate, longer study using mice (n=10 per group). It is possible that spontaneous recoveries require more time to manifest in mice (> 6 weeks) using the Montoya staircase. Following training on the Montoya staircase, mice underwent unilateral photothrombotic stroke confined to the sensorimotor cortex on the side contralateral to the mouse's preferred paw. 24 hours following photothrombosis, mice had osmotic mini-pumps implanted containing either NT3 or PBS. Mice were then assessed on a weekly basis having been given treatment.

2.3.7 NT3 is elevated in serum but not in the DRGs of treated mice

To determine the bioavailability of NT3 in blood serum, and the presence of NT3 in the DRGs, a subset of mice (n=4 for NT3 group, n=3 for control) were implanted with osmotic mini-pumps containing either NT3 or vehicle control (PBS) and allowed to incubate for 4 weeks. Blood and cervical DRGs were collected for ELISA analysis (carried out by Dr. Aline Spejo). ELISA results revealed that NT3 was significantly elevated in the blood serum ($p = 0.0006$), but not DRGs ($p = 0.86$) (Figure 13).

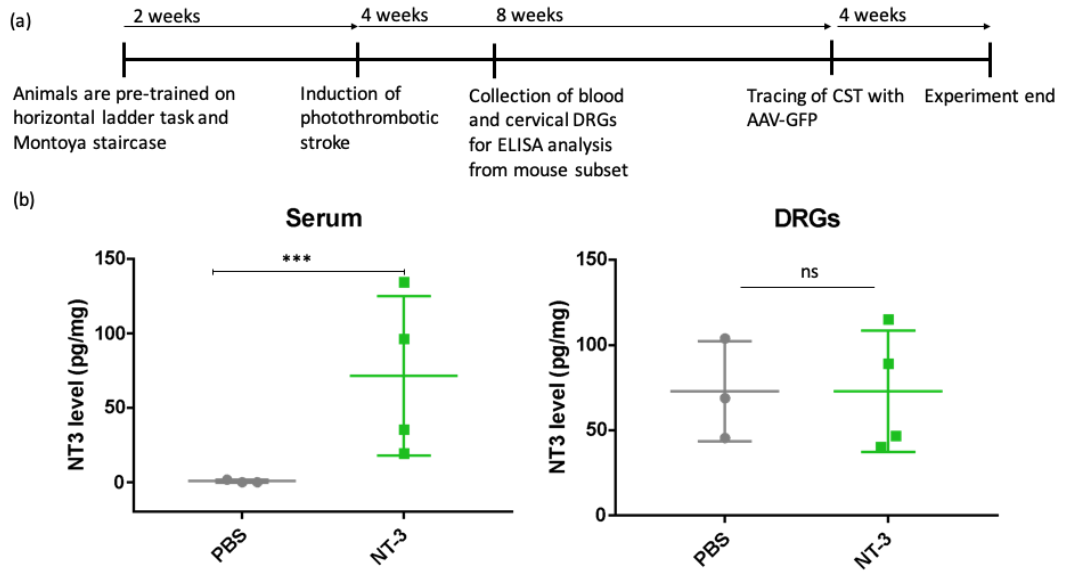


Figure 13 **Subcutaneous NT3 detected in blood serum but not DRGs.**

(a) Mice were first trained on the horizontal ladder and to reach for sucrose pellets using the Montoya staircase for two weeks. Mice then underwent photothrombosis of the motor cortex, and then assessed weekly thereafter until mice underwent a further surgery for the purpose of virally tracing the corticospinal tract from the less-affected hemisphere. Mice reached their experimental endpoint after the viral vector (AAV-UBC-GFP) had incubated for 4 weeks (to allow expression of GFP in the collaterals of the CST). A subset of mice ($n=7$) were terminated early (after 4 weeks of incubation) to measure NT-3 levels in serum. This subset of mice were uninjured.

(b) ELISA results obtained by Dr. Aline Spejo, ($n = 4$ for NT3 group, $n = 3$ for PBS) measuring the NT3 concentration in pg/mg \pm SEM in both the blood serum (left) and cervical DRGs (right) obtained from mice after 28 days with an implanted osmotic mini-pump infusing either NT3 or PBS. NT3 was significantly elevated in serum ($p = 0.0006$), but not DRGs ($p = 0.86$) (Mann Whitney tests).

2.3.8 Staircase performance was affected by subcutaneous NT3 but not the latency to cross the horizontal ladder

Assessment of functional recovery of grasping was done using the Montoya staircase task ($n=7$ for PBS group, $n=8$ for NT-3 Group). We gathered data on four metrics: the total number of pellets eaten by either the affected or non-affected paw, and the furthest step the mice were able to displace a pellet with either paw. Mice

were excluded from all analyses if they were unable to retrieve 8 pellets or more at baseline. Analysis of the data for the total number of pellets eaten using the affected paw after controlling for baseline differences in performance revealed that there was an effect of time ($F_{15,164.086} = 1.868$, $P = 0.030$), no effect of treatment group, ($F_{1,11.043} = 0.287$, $P = 0.603$), and no interaction between time and group ($F_{15,164.086} = 0.862$, $P = 0.608$) (Figure 14a). Mice from both treatment groups recovered to some degree, but the rate of this recovery was not significantly different between groups. Data analysed for pellets retrieved using the less-affected paw (Figure 14b) revealed no effect of time ($F_{15,164.012} = 0.789$, $P = 0.689$), group ($F_{1,11.007} = 2.312$, $P = 0.157$), but an interaction between time and group ($F_{15,164.012} = 2.393$, $P = 0.004$): the control treatment group significantly outperformed the mice treated with NT-3 using their non-affected paw.

Analyses were also run on the maximum distance mice were able to reach following stroke. Analysis of the maximum distance reached using the affected paw (Figure 14c) revealed that there was no effect of time ($F_{15,194.036} = 0.333$, $P = 0.333$), an effect of group ($F_{1,13.015} = 5.292$, $P = 0.039$), and an interaction between time and group ($F_{15,194.036} = 1.910$, $P = 0.024$). Therefore, mice treated with NT-3, on average, displaced sucrose pellets at steps lower down the staircase than the PBS group using the paw affected by photothrombotic stroke. Finally, analyses of the mean distance reached by the non-affected paw (Figure 14d) revealed no effect of time ($F_{15,194.010} = 1.630$, $P = 0.071$), no effect of treatment group ($F_{1,13.003} = 0.026$, $P = 0.874$), but an interaction between time and treatment group ($F_{15,194.010} = 1.927$, $P = 0.023$).




Figure 14 **Subcutaneous NT3 treatment affects the reach of mice but not their ability to grasp for sucrose pellets using their affected paw**

- (a) All measures obtained from the staircase test ($n=8$ for NT3, and $n=7$ for PBS), on both the affected and less affected paw were affected after injury, but most profoundly on the ability for mice to grasp sucrose pellets using their preferred paw, which was approximately 10 pre-injury (5 steps worth of pellets), and 4 pellets 1 week following injury. This impairment was mostly sustained throughout the assessment phase, but there was a significant recovery ($F_{15,164.1} = 1.868$, $P = 0.030$) of mice but no divergence of treatment groups ($F_{15,164.086} = 0.862$, $P = 0.608$).
- (b) The number of pellets eaten using the non-affected paw was impacted by surgery, and, over time, mice seemingly switched to using their non-affected paw to grasp pellets. There was a significant interaction between treatment groups, with an improvement of mice treated with PBS to grasp for more pellets compared with NT-3 treated mice using their non-affected paw ($F_{15,164.1} = 2.393$, $P = 0.004$).
- (c) The furthest step mice were able to displace a pellet was measured on both the affected and non-affected paws. Using the affected paw, NT-3 treated mice consistently reached and displaced sucrose pellets further down the staircase than PBS treated mice and our longitudinal analyses revealed no effect of time, but an

effect of treatment ($F_{1,13.015} = 5.292$, $P = 0.039$), and an interaction between time and group ($F_{15,194.036} = 1.910$, $P = 0.024$).

(d) Using their non-affected paw, NT-3 mice were again consistently more capable of reaching and displacing sucrose pellets further down the Montoya staircase than PBS-treated mice. There was no effect of time or group, but an interaction between the two factors ($F_{15,194.010} = 1.927$, $P = 0.023$) with the PBS group seemingly declining in the distance they reached over time.

The average latency for mice to cross the horizontal ladder was determined using LadderScorerV1. While there was a significant increase in latency for mice to cross the ladder 7 days post injury, there was no effect of treatment group. There was a detectable spontaneous recovery in both treatment groups after the 7 week period (Figure 15).

To summarise, NT-3 did have an impact on reaching, namely that of the distance mice reached, but no effect on the latency for mice to cross the horizontal ladder.




Figure 15 **Subcutaneous NT3 does not impact the latency for mice to cross the horizontal ladder**

Using LadderScorerV1, the average latency for mice to cross a 60 cm horizontal ladder was measured. The latency for mice to cross the horizontal ladder was significantly increased 1 week post injury ($F_{1,35.861} = 6.146$, $p = 0.018$), and, after 11 weeks, mice from both groups (vehicle $n = 11$, NT-3 $n = 10$) spontaneously recovered to some degree ($p < 0.05$, paired T-tests between week 1 and week 11), but there was no effect of treatment, or time.

2.4 Discussion

2.4.1 Photothrombosis

The aim of the experiments outlined in this chapter were to begin understanding how NT-3 restores motor function after stroke. Accordingly, we first wanted to establish an injury model that is well tolerated by mice (because the therapies we are testing are for survivors of stroke) and that the injury induces lasting deficits in grasping function and dexterity to mimic the lasting deficits seen in humans. Photothrombotic stroke is induced by the photoactivation of Rose Bengal, which produces potent free radicals that damage the endothelial walls that line the capillaries exposed to illumination. This damage results in vasoconstriction and platelet aggregation and these events are strictly confined to the location of the light (Watson et al., 1985). We were successfully able to show that, 6 weeks post injury, large infarcts in the motor cortex were consistently present (as determined by cresyl violet staining) of coronal brain sections, which, at their peak, had a difference in cortical volume on the lesioned hemisphere of $150 \text{ mm}^2 \pm 35$. PKC gamma immunolabeling of the cervical CST revealed an 82% reduction in intensity on the contralateral-to-injury CST.

2.4.2 Mice do not spontaneously recover their grip strength or ability to reach and grasp for sucrose pellets 6 weeks following photothrombosis.

With the stroke model established in mice, we next sought to begin understanding the role NT-3-TrkC signalling may have in adult mice, and whether it has a role in endogenous neural repair pathways that lead to spontaneous recoveries. Similar strategies have been adopted to explore the role trkB-BDNF signalling in spinal cord injured mice (Mantilla et al., 2013). The principle aim of this study was to determine whether treatment of $\text{trkC}^{\text{F617A}}$ knock-in mice with 1NMPP1 prevented any

partial, spontaneous recovery of grasping, measured by grip strength and the number of pellets obtained using the Montoya staircase. We hypothesized that NT-3 exerts its rehabilitative effects through its binding to TrkC, and the subsequent signalling cascade that this binding initiates. TrkC's inhibition with 1NMPP1 could sequester the role trkC may play in the spontaneous recovery of grasping function.

Grip strength testing revealed that there was no effect of treatment or time, indicating that 1NMPP1 had no effect on the grip strength of mice, and the initial deficits sustained after unilateral stroke were persistent for the 6 week assessment period with no spontaneous recovery (Figure 12). This was in line with data captured from the staircase task, which also revealed that initial deficits in reaching performance on the preferred paw following photothrombotic stroke were sustained throughout the assessment period (Figure 12a). This was surprising, because rats, on the other hand, do spontaneously modestly recover their ability to reach and grasp for sucrose pellets in a staircase task using their affected paw without explicit training or treatment (Wahl et al., 2014). There are some possibilities to explain why we did not observe a spontaneous recovery in the control group of the 1NMPP1 treatment study: (1) mice are inherently not capable of spontaneously recovering their grasping in the way that rats are using the Montoya staircase and bilateral grip strength meter, or (2) 6 weeks is not enough time for mice to spontaneously recover function of their affected paws on these tasks, (3) the lesion in photothrombosis in mice is comparably more detrimental and severe than that of rats, and the size of the lesion is a known factor in the amount of recovery achievable after injury (Whishaw, 2000).

Motor deficits in mice after stroke have been assessed using a plethora of assays (Balkaya et al., 2013) including those used and outlined in this chapter: grip strength (Kilic et al., 2008, Ferrara et al., 2009), staircase testing (Lee et al., 2007, Bouët et al., 2007, Baird et al., 2001), and the horizontal ladder (Tennant and Jones, 2009). In the studies using the staircase, performance was only assessed up to 30 days after injury, but supports the data we show here: the affected paw does not

spontaneously recover in its ability to retrieve sucrose pellets, beyond 3 or 4 pellets, which we can anecdotally report as being retrievable by the mouse's tongue, not its paw. Another study which monitored the recovery of mice following unilateral stroke showed that a partial recovery was possible, but the tests used (rotarod, grid walking, climbing assays and cylinder task) do not strictly interrogate motor function important for small object manipulation (Bachmann et al., 2014). As we did not observe a spontaneous recovery in our vehicle-control group, any inhibition of this by 1NMPP1 treatment will not have been detected. Notably, 1NMPP1 did not cause a decline in performance of the tasks assessed in this study, either.

2.4.3 Subcutaneous NT3 is detectable in blood serum, not cervical or thoracic level DRGs, but does seem to impact reaching performance

Previous experiments in our lab have shown that subcutaneously delivered NT-3, 24 hours following unilateral stroke, rehabilitates the upper limbs of rats (Duricki et al., 2018b). Accordingly, we sought to expand on data by testing whether subcutaneous NT-3 also rehabilitates the upper limb function of mice. Subcutaneous delivery of NT-3 is advantageous because it is a clinically feasible method of therapy delivery. In this study, we prolonged the assessment phase to 16 weeks (as we did not see a spontaneous recovery in grasping function after 6 weeks), used the horizontal ladder task to measure the latency for mice to cross and quantified additional metrics using the Montoya staircase, namely, reaching distance and the performance of the paw least preferred before injury.

Data captured from the Montoya staircase revealed that, again, initial deficits in reaching performance on the preferred and affected paw following photothrombotic stroke were sustained throughout the testing period (16 weeks). To our knowledge, there is currently no longitudinal data beyond 28 days on the reaching performance of mice using their affected and non-affected paw following cortical stroke. Our data suggests that mice, after unilateral photothrombotic injury, have a sustained deficit in

grasping using their affected paw throughout the assessment phase (up to 16 weeks), that this does not recover spontaneously and that subcutaneously delivered NT-3 does not improve grasping function either. Despite the increased length of assessment in the NT-3 study (from 6 weeks to 16 weeks) we were still unable to detect a recovery in reaching performance. Mice were assessed for 30 minutes each week; possibly with rehabilitation and increased intensity we may have seen a recovery. It has been shown in rats that there is a threshold of activity they must surpass before a post-stroke recovery of upper limb function can be detected (MacLellan et al., 2011), and it is possible that the 30 minute assessment period was not enough to promote a recovery.

2.4.4 The maximum distance reached by mice treated with NT-3 is maintained after injury, while control-treated animals decline with time

It seems that NT-3 has a protective effect on the maximum distance mice are capable of reaching using their affected paw, as mice treated with PBS decline in their ability to reach for the lower steps over the assessment period (Figure 14c and d). Given this, it may be possible that subcutaneous NT-3 has an effect on the range of motion of the shoulder and elbow joints of the affected paw. However, oddly, this also seems to be the case for the non-affected paw; there is a trend for NT-3 treated mice to consistently reach lower steps than PBS treated mice (Figure 14d). At face-value, and if viewed alone without the data collected for the total number of pellets retrieved, the maximum-distance-reached data suggests that NT-3 enhances or protects gross reaching function.

However, viewed in the context of the data collected for the total number of pellets retrieved, the relevance of this metric alone as a readout for a meaningful functional recovery is called into question: the PBS treated mice were able to, with time, presumably because of an effect of training, reach for more pellets using their less-affected paw than they were able to at baseline, despite the comparatively lower

levels of steps they were capable of reaching (Figure 14b). This suggests that the PBS treated group were more accurate and efficient with their reaching attempts than the NT-3 treated group using their non affected paw. The Montoya staircase is advantageous in that the performance of mice is measured objectively; the number of pellets remaining is counted at the end of a session and the more pellets eaten, the better the grip and dexterity of the mice. However, it also has flaws. There are 8 steps in total on each side with a total of 16 possible pellets that can be obtained. The first 4 pellets on either side of the staircase can be eaten by the mouse by simply using its tongue. The average performance of mice shown by us and others (Baird et al., 2001) is around a total of 8 pellets at their peak performance before injury, leaving a small window of 4 pellets to detect any form of dexterous recovery after injury. Therefore, it may be advantageous to train mice so that at baseline they are capable of reaching, on average more than 8 pellets using their preferred paw.

Studies in rats have shown that training the less affected forelimb can interfere with the recovery of the limb affected by the injury (Allred et al., 2005, Allred and Jones, 2008), which could be the case for the mice in this study, whose performance with their less affected forelimb improves over the testing period, possibly to the detriment of the injured forelimb. Baiting both sides of the Montoya staircase may encourage non-use of the affected limb and therefore it may be advantageous to remove the pellets on the side corresponding to the less affected limb. Cross education, or bilateral transference, is the phenomenon whereby acquisition and training of a skill on one limb is transferred to the other, untrained limb. For example, there have been some randomized controlled trials of the potential benefits of bilateral transference after unilateral stroke in humans. One study of 80 patients indicated that training the less-impaired hand improved performance by the more-impaired hand (Iosa et al., 2013) (but not *vice versa*). This is something which in the future could be studied more in animal models of injury.

To summarise, subcutaneously delivered NT-3 does not improve grasping function of the affected limb as determined by the Montoya staircase test and mice do not spontaneously recover their grasping and reaching function even up to 16 weeks of testing. Spontaneous recovery of grasping function can often cloud any benefit of a rehabilitative drug candidate; the fact that this is limited in mice may prove advantageous, as a drug that is rehabilitative could more easily be detected against its control.

Further, NT-3 did not impact the latency for mice to cross the horizontal ladder and our statistical analyses revealed that there was no effect of time in both groups, meaning neither treatment group recovered. Overall, our behavioural data suggests that subcutaneously delivered NT-3 does not improve grasping function in mice following photothrombotic stroke, and have a limited capacity to spontaneously recover. It is possible that explicit and intensive rehabilitation is an important adjunct for mice to recover their reaching function, as it is for rats (Wahl et al., 2014).

2.4.4.1 ELISA assays reveal that exogenous subcutaneous delivered NT-3 does not elevate NT-3 in the DRG

We analysed the bioavailability of NT-3 in blood serum and the cervical DRGs by using an ELISA assay (this work was carried out by Dr. Aline Spejo) from a separate cohort of animals who were uninjured, but had subcutaneously implanted osmotic mini-pumps containing either recombinant NT-3 or PBS. The DRG and blood samples were collected 28 days after implantation. Data from the lab's previous experiments suggest that subcutaneously delivered NT-3 can restore lost motor function following stroke (Duricki et al., 2018b). We hypothesised that NT-3 would be elevated in the DRGs as well as blood serum after infusion from a subcutaneously implanted osmotic mini-pump in our mice (see p. 41 for rationale). To our knowledge, intraganglionic NT-3 levels have not yet been analysed following subcutaneous

delivery by osmotic mini-pump [although we have data showing this is the case after intramuscular AAV1-NT3 (Duricki et al., 2016, Kathe et al., 2016)].

Following ELISA analysis, NT-3 was elevated in serum (76 pg/ml) compared to the PBS control (not detectable) which coincides with our previous findings (Duricki et al., 2018b). NT-3, on the other hand, was not elevated in the DRGs in the NT-3 treated group; both experimental groups had 75 pg/ml of NT-3. As an aside, endogenous NT-3 was greater in the DRGs than in the serum. The lack of elevation of intra-ganglionic NT-3 may explain why we did not see a functional recovery of upper limb function following treatment with NT-3, because we hypothesised that activation of trkC receptors upon DRG neurons is important for an NT-3 induced functional recovery. It is odd that we did not detect an elevation of NT-3 in the DRGs beyond the PBS control as NT-3 can enter the ganglia in a number of ways: by uptake from the terminals of the primary sensory neurons at the muscle spindles or Golgi tendon organs (Wright et al., 1997), and directly through the blood nerve barrier surrounding the cell bodies contained within the DRG, as it is fenestrated and permeable to proteins such as NT-3 (Jimenez-Andrade et al., 2008). In the studies where subcutaneously delivered NT-3 was successful in reversing a loss of motor function, the level of NT-3 within the DRG was not measured, but, comparably, we used a dose 5 times greater (per kg body weight animal) and did not see a reversal in lost motor function.

However, a different source of the recombinant NT-3 was used (Genentech in Duricki et al 2018 versus Peptotech in this thesis) which may account for the lack of behavioural effect seen in this study. On the whole, the results of this study indicate that subcutaneous delivery of NT-3 may not be the most optimal method of delivery, and that it is possible that we are diluting the effect NT-3 has on reversing lost motor function. Instead, a route of delivery that ensures that an elevation of NT-3 occurs in the affected muscles should be used, such as use of a viral vector that induces overexpression of NT-3, or intramuscular injection of NT-3, because we and others

have shown that this robustly induces increased rates of functional recoveries (Kathe et al., 2016, Duricki et al., 2019, Duricki et al., 2016, Duricki et al., 2018a), and that intraganglionic NT-3 is raised following these routes of delivery.

As the goal of this work is to understand how NT-3 has induced a functional recovery, and we have shown more convincingly that intramuscular elevation of NT-3 promotes the recovery of motor function in our previous studies, this type of administration of NT-3 should be deployed in future studies.

To summarise, we have shown that photothrombotic stroke is a good method of modelling upper limb impairments of stroke survivors, and that the deficits it causes are long lasting. We have shown, for the first time, that even after 16 weeks following photothrombotic stroke, mice do not recover their grasping function using their affected paw assessed using the Montoya staircase. Mice do not recover their grip strength completely up to 6 weeks after injury, and after 11 weeks following stroke, they do not recover the speed at which they cross a horizontal ladder. The possible role *trkC*-NT-3 signalling may have in spontaneous recoveries in the absence of NT-3 remains to be seen as we have yet to establish that this occurs in mice. Our data suggest the NT-3 delivered subcutaneously does not assist in the rehabilitation of mice, and that could be because this method of delivery does not elevate NT-3 levels in the DRG. Our focus in experiments beyond this thesis will be to now ensure NT-3 is elevated intraganglionically by intramuscular delivery of NT3.

3 MouseBot: Automating the single pellet reaching task for mice

3.1 Introduction

The single pellet reaching task was first proposed and detailed by Whishaw et al and is now a common-place preclinical tool for researching fine motor function and dexterity (Whishaw, 1996). The goal of the task is the retrieval of a sugar pellet reward. Animals must reach and grasp the sugar pellet through a narrow aperture requiring dextrous skill and fine motor control. The animal is placed in a 4-walled clear plastic box, with a thin vertical slot opening on one side. Beyond the slot are two fixed positions where a sugar pellet can be placed; they are offset laterally on either side of the slot so that they can only be retrieved by the paw contralateral to its position (see introductory section for image, Figure 4). Animals sense the presence of the sugar pellet reward by olfaction (Whishaw, 1996) and begin attempting to reach for the pellet. They target and extend the paw towards the pellet, supinate the forearm, position their digits over the pellet and close the paw around the pellet. The pellet is then brought back through the slot so that it can be eaten.

Research groups use this task to assess whether therapies are rehabilitative (Wahl et al., 2014, Torres-Espín et al., 2018b), however, it is very cumbersome and time consuming. Further, research personnel often carry out training and assessment during the light cycle of animals, when rodents are typically least active. Sessions for each rodent typically last for 20 minutes and requires the constant attention of the assessor, who manually records the success of each trial and replaces pellets removed from the pedestals. Assessment of the task is normally done in real-time, as the trial takes place, where a trial's success is noted by hand. Although some research groups quantify the kinematics of the reach and grasp task (Burnside et al., 2018), often times, many metrics of the task go unrecorded. Namely, the number of reach attempts; the targeting of the pellet, the trajectory of the reach; the velocity of the

reach. Many of these metrics would be very difficult or prone to error if analysed in real-time (such as reach attempts) or impossible to do without first recording the footage of the trial to analyse off-line (such as reach trajectory).

Rats are often the preferred animal for use with this task as they are amenable to training and motivated to continue the task even after injury. Mice are also more difficult to train than rats, typically achieving a peak success rate of 50% (Chen et al., 2014) compared with rats who often surpass 70% performance criterion (Wahl et al., 2014). Therefore, the difference between pre and post injury performance is limited in mice. Performance increases on a behavioural task is often dependent on repeat exposure, therefore, automating the single pellet reaching task for mice such that they are exposed to the task more so than if they were trained manually may increase the performance level possible with mice.

Automating the single pellet reaching task is challenging for several reasons. Dispensing pellets without them crushing while avoiding jamming events is the first hurdle to overcome. The second is presenting them to mice one at a time in the same position and is critical to the reproducibility of the task. The device would need to be safe if placed in the home cages of animals and be durable enough not to deteriorate in their presence (i.e. as a result of chewing). It would also need to have failsafe mechanisms in place so that should a failure occur, the device would automatically switch off and cease to provide power to any of its electronics. It would also need to be able to operate in varying lighting conditions. Mouse cages contain bedding, tunnels and chew toys; a device would need to be able to operate within the cage and handle possible interferences with any of these cage supplements.

As reviewed in the introductory chapter, there are multiple research groups who seek to automate the single pellet reaching task but their desktop solutions only automate the task for one animal at a time. Early work in our lab focused on automating the single pellet reaching task in rats by creating an in-cage solution which can train and assess multiple animals (Gadiagellan, 2018).

To our knowledge, no device exists that automates the single pellet reaching in multi-housed mice. We sought to automate the single pellet reaching task using robotics, computer vision and machine learning for mice such that pellet dispensing and positioning was carried out automatically by the device, with little to no human supervision, and with automatic acquisition of trial footage.

3.3 Methods

3.3.1 Microchipping mice with Radio Frequency Identity tags

Mice (C57BL/6) were anaesthetised with isoflurane (5%) in oxygen (1 L/min), and, using a syringe loaded with a glass capsule RFID tag (FDX-B ISO 11784/11785, supplier: Alibaba) injected into the flank, before the RFID tag was manoeuvred subcutaneously between the shoulder blades of the mouse. All procedures were in accordance with the Animals (Scientific Procedures) Act of 1986. All protocols involving animals received prior approval by the King's College London Animal Welfare Ethical Review Board and were authorized by the UK Home Office Project (license number 70/7865, held by Dr. Lawrence Moon).

3.3.2 Component Design

The MouseBot is designed to fit within the home cages of mice. The hardware for MouseBot was designed using Computer Aided Design (CAD) software (SolidWorks Student Edition 2016-2018). An overview of the device can be viewed in supplementary video 1.

3.3.3 Fused Deposition Modelling (FDM) 3D Printing

Components to be 3D printed were exported as .STL files and the g-code was generated using Ultimaker's slicing software, Cura. The 3D printed components were fabricated using a polylactic acid (PLA) thermoplastic, with a layer height of 0.2 mm, a nozzle diameter of 0.4 mm and an infill percentage of 20. The Ultimaker was configured to extrude PLA as per the manufacturer's default settings. Support material used for components was also in PLA. The printers used for the manufacture of components were the Ultimaker 2+, 3 and 3 Extended.

3.3.4 Laser Cutting

Components to be laser cut were rendered to in 2D at a scale of 1:1, exported as an Adobe Illustrator (.AI) file, imported into InkSpace and had the line widths of the

drawing adjusted to 0.025 mm before being exported as a PDF file (readable by the laser cutter). All of the acrylic components (the face-plates, and subset of the spur gears) were cut using a 60 watt Epilog laser cutter. The thickness of acrylic (supplier: sheetplastics.co.uk) cut was 3 mm; the parameters used to cut these components were: 20% cutting speed, a frequency of 5000, and at 100% power.

3.3.5 Software

The device software was written in Python 3 and executed on a Raspberry Pi 3 with the Raspbian operating system. A flow diagram of the firmware's logic is shown in Figure 16. Firmware was stored in a GitHub repository. The Raspberry Pi was configured to automatically update the firmware from the GitHub upon MouseBot start-up. The MouseBot's software underwent multiple revisions which allowed for a more iterative and agile development process. The output of this development process is a modular, flexible and extensible firmware package. Emphasis was placed on creating modular software packages such that firmware changes could easily be made given new hardware revisions, and also for the purpose of validating different component substitutions.

The MouseBot exists in 4 active states when it is functioning: Pellet Dispensing, Pellet Positioning, Active Trial, Video Processing. Computer vision is essential for the MouseBot's functioning. The device's camera was used extensively in order to minimise the number of electronic components required. In order to make the MouseBot a generic and extensible platform such that others can customise its firmware to cater to their particular research task, fully modular, object-oriented software libraries were created. Libraries were created for data handling, visual data processing, motion control and pellet handling workflows.

When a device is turned on, its firmware automatically initialises the hardware and data connectivity components. This begins by a connected device downloading the latest firmware from a GitHub repository. If the device is not connected to the

internet, it attempts to find a locally stored configuration file. Failure to achieve both means the device needs to be calibrated. It prompts the user to do so by sending an email, and then terminating. If a calibration file is present, the device begins by first turning on the infrared (IR) lights to illuminate the camera's field of view, and spawns a background task which periodically polls the RFID reader to determine if any RFID tags are in range. When a pellet is dispensed from the chute to the spoon, and it is detected by the camera, it indicates to the device that a trial has started.

An active trial ends when the pellet is dislodged from the spoon. Prior to this, image frames generated by the camera at 180 frames per second (FPS) are stored in a circular buffer (with a maximum capacity of 700 frames) until the pellet is no longer detected within a frame. When the pellet is no longer visible, the firmware stops adding frames to the buffer and the frames stored within it are encoded to video (with XVID encoding, in a .avi container) and transmitted to cloud storage (MongoDb) along with the last RFID tag detected during that specific trial. Given that the 700 frames were obtained at 180 FPS, 3.8 seconds of real-time footage are stored at 1/6th the speed of real-time playback. This enables easier reviewing and user validation of the captured data. Trials are pellet centric; if a mouse enters the tunnel to grasp for a pellet and subsequently leaves empty handed, with the pellet still present on the spoon, a trial will not be recorded or persisted to the cloud storage. The flow-chart logic detailing the MouseBot's firmware can be viewed in Figure 16.

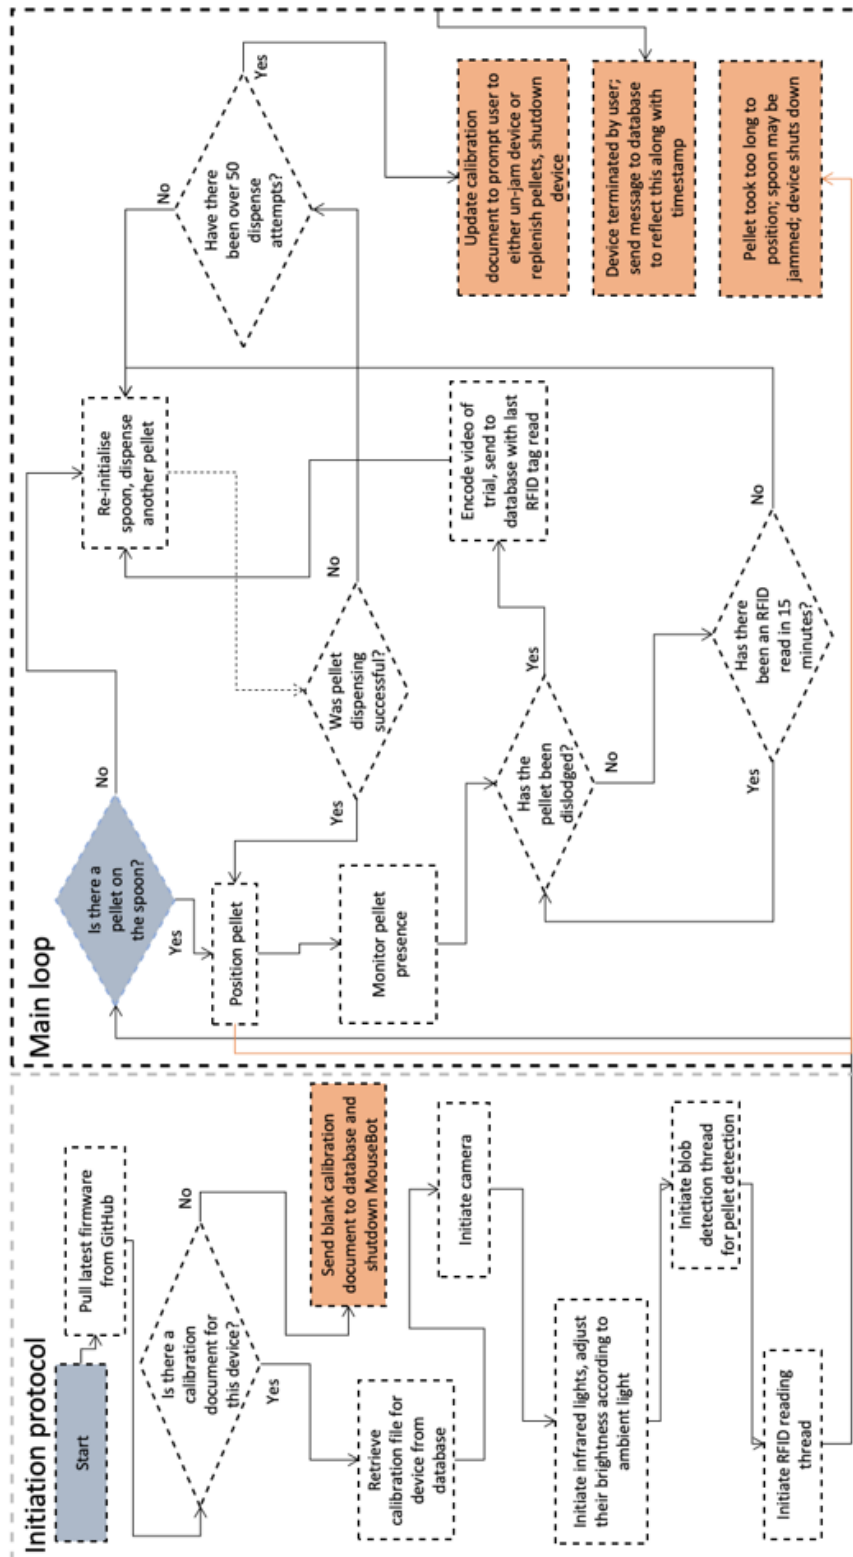


Figure 16 **The MouseBot's python-based firmware which is fully customisable**

The firmware consists of an Initialisation Sequence and a main loop.

3.3.6 Calibration

When a MouseBot is turned on for the first time, the firmware is configured to search for a calibration file unique to its Pi's Media Access Control (MAC) address. Calibration files are stored online, so that they can be altered remotely. A local copy is also stored in the event an internet connection cannot be established. All data generated by the MouseBots are stored in MongoDB. MongoDB is a document-oriented, NoSQL database. Calibration documents are stored in a cloud instance of MongoDB and contains various variables for each individual device. If a calibration file for the device does not exist (i.e. it is the first time it has been initiated), a calibration document is created, and the calibration state is set to Incomplete. The user will be prompted for the calibration procedure to be executed. For example, to assign a name to the device (such that data querying can be confined to individual cages) or to specify the region of interest to contain blob detection to within the frames generated by the cameras. Calibration files contain the fields required for MouseBots to function. Each calibration document contains the metrics in table 1 and table 2.

3.3.7 Motor Control

Actuation of the hopper drum is achieved with a stepper motor and gear configuration (Figure 17). It is controlled using a Tic Stepper Motor Controller T834 board. Commands sent from the Raspberry Pi 3 to these motor controllers using the Inter-integrated Circuit (I²C) communication protocol. The firmware invokes the method `dispenseOne(duration=4)` which causes the dispenser's motor to turn clockwise for 4 seconds. By default this method's duration argument is set to that specified in the device's calibration file (see later for more details on the calibration file).

3.3.8 Pellet dispensing

Singular pellet dispensing is a complex engineering challenge. The off-the-shelf dispensers are often expensive and large. Designing and fabricating them from

scratch using 3D printed technologies introduces further challenges. 3D printed components are not smooth due to a relatively low layer resolution (0.2mm), and the pellets are brittle, meaning they are susceptible to crushing. The RatBot (Gadiagellan, 2018) used a hyper-smooth and polished augur, which drove pellets up the hopper and into the chute (Figure 8). While this approach worked reliably, the polished augurs were costly. For this reason, we sought to create a dispenser which would be both low-cost and 3d-printable. Ideally, a device would be left unattended in-cage for multiple days, and as such, it was imperative that the dispenser operated reliably, minimised dust generation and did not crush pellets. The final dispensing module contained a custom-built hopper with the capacity to fit approximately 1200 x 20 mg sucrose pellets (Datesand, #F05301). These pellets are guided towards a slot by small dividers surrounding the inner diameter of the bottom most part of the hopper (Figure 17, supplementary video 1 and supplementary video 2). A flow control mechanism was designed to prevent aggregation at the bottom of the drum which otherwise would have resulted in jams and subsequent overload of the stepper motors. The bottom most part of the hopper is drum-like and rotates in order to facilitate the motion of pellets towards the dispensing chute.




Figure 17 A rendered schematic of the dispenser module

The pellet drum is driven by a motor whose drive axis is oriented normal to the drum's centre axis and power transfer from the motor to the drum is achieved by a spur gear mechanism geared to increase torque (9:31) (Figure 17). Bearings ensure correct seating and low friction of the rotating drum. A garbage chute was designed to feed into the lowest most part of the hopper, to ensure the removal of crushed pellets or debris. Pellets are driven through the hopper's destination chute by gravity and the rotation of the drum towards the bottom most part of the hopper. The lowest section of the drum contains dividers in a circular pattern (Figure 17). These dividers drive pellets up and around the drum toward a chute that channels pellets to the spoon. The rate of dispensing is modulated so that typically a single pellet enters the chute and is directed towards the spoon, however in the case that more than one pellet enters, the spoon is designed such that only one pellet can be held. Surplus pellets fall into a reservoir at the bottom of the device which are inaccessible to the mice, and thus can be reused.

3.3.9 Onboard Camera

The MouseBot's onboard camera (Raspberry Pi Camera NOIR v2, connected to a Raspberry Pi) is positioned to obtain an elevated view of the reaching slot, the

dispensed pellet and the forepaw of the mice carrying out the task (Figure 18). The camera facilitates the following tasks: pellet detection, pellet positioning, pellet monitoring, and the recording of the active reaching trial. These are all achieved by using the Python OpenCV library and executed locally on the Raspberry Pi. When active, the camera captures 180 frames per second, with dimensions 640(w) x 480(h) pixels. It typically takes a mouse less than 0.3 seconds to reach, grasp and retrieve a pellet. Accordingly, the firmware of the device was designed such that the highest frame rates and resolutions were maintained, while sustaining low CPU usage for the Raspberry Pi, and a scalable file size (< 1MB), so that many videos can be saved to a server.

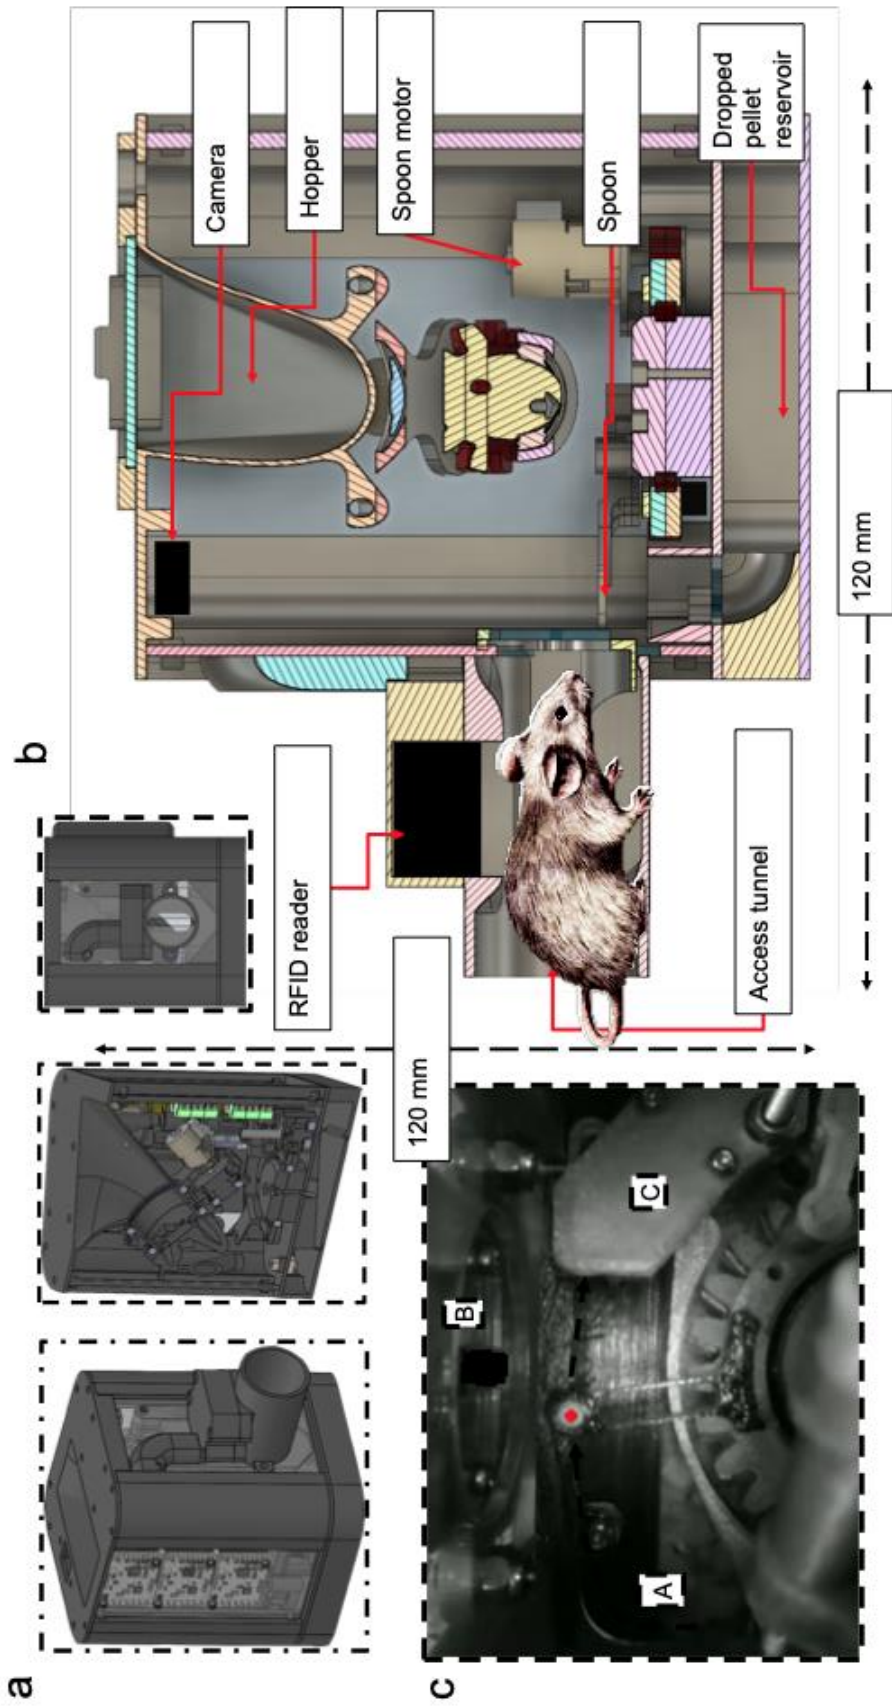


Figure 18 **MouseBot: a 3d-printed and laser cut unit controlled using a Raspberry Pi**

- a. Three separate views of the MouseBot, from left to right: isometric view, a cross-section, showing how the hopper is situated, and the front view (for a more detailed view, see supplementary video 1).
- b. A cross section of the Mousebot shows how a mouse interacts with the device, which climbs and walks through an elevated tunnel (its elevation limits the amount of bedding and other debris from entering the device).
- c. The camera's view is top-down and shows the spoon with a pellet highlighted in red. This red highlight was automatically generated by the blob detection algorithm used to position and detect the presence of the pellet. The pellet arrives from the left hand side (A), rotates into the view of the camera and is positioned relative to the slot (B), and then once the trial is finished the spoon goes to (C), a brush, where the spoon is cleaned, completing its process.

3.3.10 Encoding Video

Videos are encoded locally on the Raspberry Pi, saved locally to disk, and, if there's a connection to the cloud-hosted database, converted to a byte stream and stored in the MongoDB instance with gridfs. Gridfs is a protocol for storing and retrieving files that exceed the 16MB document threshold of MongoDB (<https://docs.mongodb.com/manual/core/gridfs/>). This enables videos and large documents, greater than 16 MB, to be stored in chunks within the Mongo database, and later stitched back together when called upon. If the video is sent to the cloud successfully, the local copy is deleted, otherwise it is retained until a connection is re-established and subsequent cloud upload is successful. While the camera stream captures frames at 180 frames per second, with frame dimensions of 640 by 480, videos are cropped to only include regions of interest (i.e. the area of the frame containing only the spoon and slot) before being stored online. This reduces encoding time, bandwidth usage and storage requirements.

3.3.11 Pellet Detection using Blob Detection

In the RATBOT by Gadiagellan et al 2018, the presence of a pellet was detected using proximity sensors embedded within the spoon. In this version, the need for the sensor is eliminated and instead, computer vision techniques are employed. This reduces the overall electronic footprint of the device requiring few electronic components. We found that infrared through-beam sensors (monitoring the chute) were not reliable for pellet detection due to dust accumulation which ultimately resulted in permanent beam obstruction.

Successful pellet dispenses are assessed by having the spoon move such that it is just in view of the camera, and using a blob detection algorithm to determine whether pellets are present within the frame. Blob detection can be optimised to find uniform areas of pixels that are within a given pixel intensity range and fall within a specified surface area. Empirical testing of these parameters, given the frame size and position of the camera, found that pellets could be identified uniquely (even if other similarly sized objects entered the field of view, such as bedding) by setting the minimum and maximum intensity thresholds to 50 and 255 respectively, and filtering by blob size not to include any blobs which are smaller than 100 pixels in area (Figure 18 and supplementary video 3). Pellets have high RGB pixel values (close to 255) given their light colour, and the background hardware on which the pellet resides is black to ensure a high contrast between the two. An infrared ring light surrounds the camera lens such that the illumination of the field of view is maintained during the dark cycle. The drop zone is designed such that dropped pellets are guided away from the field of view after they are dropped so that pellets cease to be detected.

The dispenser is programmed to rotate for 4 seconds before moving the spoon into the field of view of the camera to detect whether a pellet is present. Next, the firmware analyses a sliding window of 30 frames from the camera's live feed using the blob detection algorithm, taking approximately 1 second to complete. If there are

more than 6 frames with a blob detected; the function returns True, indicating a pellet is present, otherwise False is returned.

3.3.12 Pellet positioning

Positioning the pellet relative to the reaching slot is an integral part of the single pellet reaching task (supplementary video 3). Modulating the position of the pellet allows us to tailor the task to individual mice as they learn. The traditional Wishaw window task features two divots of fixed distance relative to the slot on either of its sides. We sought to create a method by which pellet positioning could be made highly customisable with varying distances such that the distance from the slot could be modulated automatically depending on the performance of the animal. For example, we could in principle automatically modulate the spoon's position depending on the performance rating of the animal which may result in faster training times. The slot is directly in the middle of the frame (because of the location of the camera mount), and as such, the middle position is found by halving the dimensions of the x axis of the images returned. The width of the cropped frame is 450, and therefore, the middle is 225. Pellets are presented to the mouse with a "spoon" (Figure 18c).

The spoon is an arm with a concave cut-out at its most distal part where pellets are caught from the chute before being presented to mice. The spoon and its driving mechanism are centrally positioned in the MouseBot (Figure 18). It is driven from the dispenser's chute towards the slot in a circular/arc motion by a stepper motor positioned parallel to its axis of rotation (supplementary video 2). Power transfer from the stepper motor to the spoon is applied by a spur gear. The spoon is actuated during pellet detection and pellet positioning states.

Once blob detection confirms the presence of a pellet's presence, the spoon is rotated to its target location. A proportional closed-loop control system is implemented for pellet positioning. The spoon is moved in the direction that minimises the error between the current coordinates of the detected blob, and the target

coordinates of the presentation position (supplementary video 3). Once the pellet's x coordinate matches that of the target coordinate, the spoon motor is powered down and the device enters the monitoring state. This signifies an active trial is in progress and the circular buffer containing 700 frames begins population. This method is run as a background task using a multi-threaded approach so that in-memory frame storage and previous-task-encoding can take place concurrently. The firmware polls the background task for the presence of a pellet at a frequency of 0.5 Hz.

When a pellet is no longer detected, the MouseBot goes through an initialisation sequence whereby the spoon is moved towards a cleaning brush (supplementary video 1) and movement terminates when it hits a homing switch. Once the homing switch is engaged, the motors switch direction until the spoon reaches the second switch signifying the spoon has arrived at the chute and is ready to receive another pellet.

3.3.13 MouseBot identifies each mouse using RFID technology

To ensure the correct identification of mice carrying out the task, the MouseBot uses radio frequency identity (RFID) technology. Mice are implanted subcutaneously with glass capsule RFID tags (Figure 19b) under anaesthesia such that the capsule is positioned between the shoulder blades of the animal. Access to the MouseBot's slot is gained when mice walk through a tunnel measuring 4.5 cm in diameter, which confines access to one mouse at a time. RFID reading is handled via a serial connection between the Raspberry Pi and an RFID reader (manufacturer: Priority1Design, model number: RFIDREAD-mRW-125) and the function handling its operation is contained within its own thread. Multithreaded programming enables a computer to run different software tasks concurrently, such as the collection of image frames from the camera, while the RFID reader scans for in-range tags. The thread tasked with returning the RFID reader output polls the reader to determine if it has detected an in-range RFID tag. Polling takes place at a rate of 2Hz. When an RFID

chip is in range of the reader, the tag ID is added to a circular buffer. The RFID reader allows us to know how long a mouse has spent in the tunnel, when the trial began and whether it left the tunnel empty-handed.




Figure 19 **Identifying individual mice in a multi-housed cage with a retrofitted MouseBot**

(a) The histogram of trial activity over 4 hours, between 6pm and 10pm with 24 bins (each bar is a 10 minute interval). At its peak, the MouseBot can run 9 trials every 10 minutes with 5 mice housed in the same cage.

(b) An RFID capsule beside a quarter for size comparison and the RFID reader to the right (not to scale). The RFID reader is integrated within the MouseBot's access tunnel such that individual mice can be identified.

(c) A summary of the number of trials executed among a sample of 5 mice within the same cage with intermittent access to a MouseBot over the course of 2 months during prototyping. Each row represents a mouse with the first column being their respective RFID tag.

3.3.14 Data Storage

The dataset containing all of the trial videos and its metadata is stored in a MongoDB instance run on a cloud server (Digital Ocean) and local backups are made on a weekly basis. Videos are stored using GridFs.

3.3.15 Validation

3.3.15.1 Pellet dispensing and detection

To validate that blob detection was a viable method of pellet detection, a software routine (`DispenserPerformance.py`) was written to have the MouseBot present 70 pellets to the camera, and a determination of pellet presence was outputted by the device. This output was confirmed manually and a percentage of correct and incorrect predictions was generated.

3.3.15.2 Pellet positioning

Validation of the accuracy of pellet positioning was achieved by using a convolutional neural network trained for key point detection to locate the position of the pellet in a frame. 80 trials where the pellet's target position was set to "middle" (which is the x coordinate that is half of the width of the frame) were extracted from the database for the purpose of determining the consistency of pellet positioning. These videos underwent key-point detection to detect the position of the pellets and the likelihood (a value between 0 to 1) that the coordinate was accurate according to the trained network (see below for more details). Videos are event synchronised, because they are encoded and stored when the pellet has left the frame. Therefore, we found that pellets were typically present at the 450th frame of each video and absent at the 650th (because the likelihood value as determined by the pose-estimation algorithm was consistently either close to 1 or 0 at these frame numbers, respectively (see Figure 23a for a representative trace).

3.3.16 The firmware implements timeouts in the case of failure events

Incorporated into the device's firmware are safety checks that result in powering down the motion control system. That is, automatic fail-safe mechanisms terminate the trials in the event that a) the pellet takes too long to position, b) there has not been a pellet detected in 50 dispense attempts or c) the spoon has not

reached the chute or cleaning station. In the event that any one of these scenarios occur, an email is also sent to the user along with the appropriate error message so that the device can be investigated. Re-initialisation of the device will not occur until the safety-error has been cleared. The motors that drive the spoon and dispenser are controlled via stepper motor controllers which are configured with a “watchdog” facility, that is, if the stepper controllers do not receive an “I’m alive” signal every 1 second from the firmware, the motors will power down. This is implemented in the case the firmware crashes during operation of the motor; which would otherwise continue running until explicitly sent a command to de-energize. These fail-safe mechanisms were measures taken to ensure the safety of the animals within the cage given that it is an in-cage device.

3.3.17 Training neural network for Pose-Estimation

A neural network was trained to estimate the pose of mice reaching for pellets during trial periods. The neural network was trained using the DeepLabCut toolbox (version 2.06) (Mathis et al., 2018). To create the training dataset, frames were selected from four 23 second videos of a mouse reaching using the MouseBot. A total of 153 frames were extracted manually and labelled using DeepLabCut’s graphical user interface executed within a Conda environment on Mac OS X. Frames that were extracted typically contained all or most of the features of interest; the pellet, the mouse’s snout, and 1st, 3rd and 5th digits (Figure 20 and supplementary video 6). Once frames were extracted and labelled, the neural network was trained using Google Colab with a runtime utilising a GPU. The network was trained for 110300 iterations.

3.3.18 Training a Long Short Term Memory (LSTM) recurrent neural network for trial classification

Training a trial outcome classifier using the raw pixel values from each frame from each video would be computationally expensive. We therefore used the

DeepLabCut toolbox to create a feature detector for the snout, the pellet and the leading, middle finger of the mouse's paw such that their location within each frame could be determined. With the model trained, we then created a pipeline such that videos in the database were analysed with our feature detector and persisted the results as metadata along with each video. The x, y coordinates of each detected feature and their likelihood parameters were persisted. Each video contains 700 frames, which were captured at around 180 FPS, and were encoded to playback at 30 frames per second. Each video terminates shortly after a pellet leaves the frame. For this reason, videos are event synchronised because the removal of the pellet occurs at the same time in each video. This means the feature detector's output data require minimal pre-processing in order to be used as a training dataset for machine learned models. Initial attempts to automate the outcome of the trial (miss, or a hit, see supplementary video 4 and supplementary video 5 for examples) using a decision tree (a decision tree being a series of if/else statements that collectively predict an outcome) did not generalise well across the dataset; in contrast, the temporal structure of the pose-estimation data makes it suitable to be classified using a neural network of type Long Short Term Memory (model), which takes into consideration the order of the input data.

LSTM models are typically used to predict the next item in a sequence, for example the last word in a sentence, or the price action of a stock given its historic price data. They can also be used to classify time-series data too, given a sequence. As each video is stored in the database along with a sequence of coordinates of the key features in each frame, we opted to use an LSTM model to classify trial outcomes using Google Colab, written in Python 3, which leverages the free GPU resources (NVIDIA Tesla K80) in Google's cloud. To gather the data to train the recurrent neural network for trial classification, the database containing all of the saved videos and metadata was queried to return all videos which had been manually scored, and that had already been analysed by our trained feature detector. This returned 1796 videos,

of which 511 were manually classified as “hits” and “1285” as misses. Training a model with an unequal number of class examples may lead to biases in the eventual model, and therefore, the video dataset to train the model was reduced to 1022 to contain an equal number of videos manually classified as “hits” and “misses”. The data used to train the eventual model were the coordinates and likelihood values, generated by our previously described convnet, of the pellet, the snout of the mouse and the leading middle finger. The shape of the data fed into the network was a 250 by 9 matrix, which was the number of frames used from the video and the 9 features just described, respectively. We found empirically that using just the last 250 frames from each video was most predictive of the trial outcome. The dataset was split into validation, test and training subsets by the following ratio 159:225:638. A sequential model of type LSTM was created (using the high-level neural network API “Keras” running on top of Tensorflow) with 256 nodes and a sigmoidal activation layer. Network weights were iteratively adjusted using the Adam optimisation algorithm, with a learning rate of 0.001. The model was trained for 200 epochs. See Figure 20 for overall workflow.

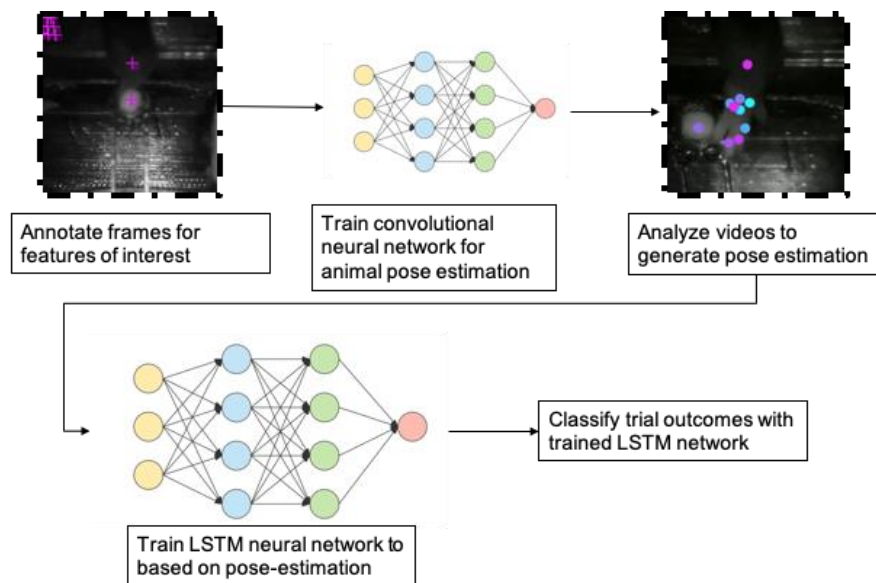


Figure 20 **A schematic of the pipeline used to generate a trial outcome classifier.**

Key frames are extracted from a sample of videos containing a mouse reaching for a sugar pellet generated by the MouseBot. These annotated features are used to train a convolutional neural network for pose-estimation. The data generated by the pose-estimation is then used to train a Long Short Term Memory network for trial classification.

3.4 Results

3.4.1 Overview of device

MouseBot is an automated version of the single pellet reaching task and consists of five main hardware modules; (1) a 20mg sugar pellet dispenser, (2) a “spoon” for positioning the sugar pellet to the left or right of a slit, (3) camera with infrared light array, (4) Radio Frequency Identification (RFID) reader [see figure 19] and (5) distribution board and Raspberry Pi 3. All five modules are contained within a unit measuring 120 x 120 x 120 mm (Figure 18a and b). Briefly, the MouseBot is a device that fits into the home cages of mice. Pellets are loaded into a hopper situated at the top of the device (with a maximum capacity of 1200 20mg sucrose pellets), and are fed through a rotating dispenser mechanism one-by-one where they fall onto a spoon via a chute. The spoon contains a concave cut-out for the seating of pellets (Figure 18c and Figure 22). The pellet is detected on the spoon using computer vision algorithms. Our data provides the feedback in a closed loop positioning sequence (see supplementary video 3). The camera video records trials which are subsequently stored in the cloud. Mice, which are RFID chipped in the nape of their neck, access the reaching window through a small tunnel. The tunnel has an integrated RFID reader that identifies the mouse interacting with the device (Figure 18d). All of the components were either purchased, 3D printed or laser cut.

3.4.2 Pellet Handling

3.4.2.1 Using blob detection algorithms to determine the presence of pellets

We found blob detection to be a good determinant of pellet presence (supplementary video 3, Figure 21). This was manually validated and was found to be 97% successful over 70 trials. The two failures were due to damaged pellets, not quite spherical, seated in the spoon such that it did not appear circular in the view of the camera; i.e., our pellet detection method wrongly predicted the presence of a pellet as absent twice out of the 70 trials.




Figure 21 **Validating blob detection as a method of pellet detection**

A confusion matrix outlining the predicted outcome pellet dispenses and whether they were detected with the blob detection algorithm against the actual outcome across 70 trials. The overall accuracy of the algorithm was found to 97%, with the errors (2 out of the total 70 trials) being due to slightly damaged pellets not being picked up by the blob detector.

3.4.2.2 Pellet positioning was found to be consistent

Pellets are 2.5 mm in diameter which equates to approximately 50 pixels within the frame of a video when it is seated on the spoon (Figure 22 and Figure 23a). Pellets were consistently present at the 450th frame of each video, and also consistently removed from the frame at the 650th as determined by using a pose-estimation model to predict the presence of a pellet in a given frame (Figure 23b). Distribution plots of the predicted coordinates of a pellet from the 450th and 650th frames of 80 videos found that pellets were consistently positioned very close to the “middle position” with a mean x coordinate of 230 px (5 px greater than the target, or, 0.25 mm) with a standard deviation of 25px, which equates to an error of +/- 1.25 mm (Figure 23c and d).




Figure 22 **Stills taken from a reaching trial where a pellet is dislodged and missed**

A pellet is positioned in the "middle" position which is directly in front of the slot. "Middle" corresponds to an x coordinate value of 225. From top to bottom: (i) the pellet is placed on the spoon, and positioned just in front of the slot. (ii) The mouse reaches through the slot and grasps at the pellet. (iii) The pellet is dislodged as the paw of the mouse retracts backwards.




Figure 23 **Pellet positioning is consistent**

(a) A still taken from a MouseBot-generated video where a pellet is positioned on the spoon. Pellets have a diameter of 2.5 mm which equates to approximately 50 pixels.

(b) A representative trace of a pellet's likelihood of being in a video's frame as determined by a key-point detection model used to pin-point the whereabouts of a pellet. In this particular video, the pellet's likelihood remains as 1 until approximately the 500th frame, where it drops to close to zero for the remaining frames.

(c) A density plot created from the coordinates of the pellet as determined by our key-point detection model, at the 450th frame from 80 MouseBot-generated videos, pellets are aggregated around the red dashed line depicting the pellet's target location. The mean x coordinate was 230 px (5 px greater than the target, or, 0.25 mm) with a standard deviation of 25px, which equates to an error of +/- 1.25 mm.

(d) We also took the pellet's co-ordinates at the 650th frame; the distribution was much more broad, consistent with a variety of successful and unsuccessful retrievals.

3.4.4 Using a neural network for pose-estimation and a recurrent neural network to classify the outcomes of trials

A total of 42 mice have interacted with a MouseBot during its prototyping stages, and they have collectively carried out over 9153 trials. We found that the number of trials executed by group housed mice varied between the animals. Often, the majority of the trials carried out with the device will be done by one or two mice (Figure 19c) regardless of the number of mice housed together (which ranged between 5 -10). Despite this, a MouseBot will typically record around 200-300 videos of trials in a given night when housed with mice. These videos can be manually classified as a hit (i.e. a successful grasp) (supplementary video 5) or a miss (supplementary video 4) (a failed retrieval) very quickly as they are only seconds long and the outcome of the trial can be classified by playing back the last 250 frames only. However, with time and more devices, manual assessment will quickly become unsustainable. Given this, we trained a simple LSTM neural network to classify the outcomes of trials. We used 1022 examples of trials, containing 511 cases of 'hits' and 'misses' that had been manually validated, and trained a model for 200 epochs. Initial attempts using just the predicted location of the pellet as the LSTM model's training data looked promising: the pellet-only trained model had an overall accuracy of 77%. We then iteratively improved the model by confining training data to the final segments of the video, increasing the input features of the training dataset to include both the snout and middle finger of the animal, and compressing the input data to integers. The model classified trials correctly for hits and misses 91% and 80% of the time, respectively, using the coordinates and likelihood parameters of the snout, middle finger and pellet (Figure 24).




Figure 24 **An LSTM neural network is used to automatically classify the outcomes of trials**

A confusion matrix outlining the performance of our trained binary classifier for the purpose of automatically scoring the outcome of videos generated by the MouseBot. The model correctly classified 'hits' 91% of the time, and 'misses', 80% of the time.

3.6 Discussion

Our results show that with open-source software languages (Python), single board computers like the Raspberry Pi, and fusion deposition modelling (FDM) 3D printing, bespoke, custom devices can be developed to automate preclinical animal behavioural assays. We show that we can successfully dispense pellets in single file, detect their presence using computer vision algorithms, position the pellets relative to the slot, guided by the integrated camera, and exclusively record the video frames leading up to the removal of the pellet with a frame capture of 180 Hz. MouseBots can transmit data to cloud storage wirelessly and all data generated can be queried from any device with an internet connection using Google Colab. While this chapter details the last iteration of the MouseBot, both the software and hardware throughout the undertaking of this PhD, underwent major revisions. For the software, this meant re-writing the firmware so that it became more modular (making it easier to change by me or by others in the future). For the hardware, the module responsible for pellet dispensing also went through major revisions. Flow control mechanisms were introduced to limit pellet aggregation and crushing for example. We found that, while 3D printing, laser cutting and Python development enables academics to automate their behavioural assays, there are also some bottlenecks. The following paragraphs will describe problems that were encountered during prototyping and solution that were developed

3.6.1.1 3D printing and Computer Aided Design

Fusion Deposition Morphometry (FDM) is the leading consumer level 3D-printing technology, owing to its relative ease of use and increasingly lower prices to purchase and maintain. New users of 3D printing will quickly discover its benefits. An idea can go from concept to prototype in as little as a day. However, with more complicated problems, the bottlenecks of these technologies, and design practices in general, quickly become apparent. The MouseBot, for example, has many components and

moving parts which can be prone to jamming, failure, or other errors. Users who are motivated to automate a task that is laborious and time consuming may find that the solution they propose may also require a level of maintenance that rivals the time required to do the task that they try to automate.

Devices with many moving parts are more likely to require maintenance and can be difficult to assemble and scale up. For that reason, users are advised to spend as much time as possible at the conceptual and design stages of prototyping in order to minimise the number of components and moving parts that are required. This may solve two issues: (1) less time will be required maintaining devices, and (2) it simplifies both the manufacturing and scaling up of these devices.

If the academic wishes to automate a behavioural task, it is advisable that they spend as much time as possible doing the task manually. This may bring to light simple solutions to the problem, but also give the designer insight about the animal behaviour which may cause the device problems later on. Events that may occur rarely, such as bedding falling into the device, may result in lots of time spent maintaining the devices later on.

Users are advised to familiarise themselves with their printer; to understand its tolerances (for FDM printing, this is typically +/- 0.2 mm), to understand how to optimise the slicing settings, when and when not to use support material, which orientation to print the model and also how to design a model such that minimal printing support material is required for a successful print. This may mean printing a component in parts, which then raises the question: how best to adhere them? An adhesive glue is permanent; will the device need to be disassembled for cleaning? Can the two components be adhered merely by push-fit? These are example questions one must ask in order to save time in future.

3.6.1.2 MouseBots emit noise and this may contribute to altered behaviour

Each MouseBot is equipped with multiple moving parts, actuated by two stepper motors. While every effort is made to reduce the noise these devices make, namely, not to startle or alarm animals, the devices, particularly when dispensing a pellet, emit a volume of around 50 decibels from a distance of approximately 10 cm from the device. In future, we may make efforts to dampen this noise. However, it may also be advantageous: the noise the device emits may act as a cue for mice, informing them when a pellet has been dispensed. While have not formally been tested, efforts will be made in future to understand how the sound these devices make may affect the behaviour of the mice.

3.6.1.3 The cost of the components and to manufacture a MouseBot is approximately £150 not including the 3 hours of labour for their assembly

The cost to manufacture a MouseBot is relatively low in comparison to typical behavioural equipment, with the most expensive components being the electrical boards; the Raspberry Pi (£30), the two stepper motor controllers (£25 each), the two stepper motors they control (£10 each), and the RFID reader (£25). The total cost of these components is £125. The 3D printed components weigh approximately 400 grams. A typical spool of 3D printed PLA material is £30 for 1kg, meaning the printed hardware is approximately £12. The laser cut components are a negligible cost, with the approximate size of a sheet required to cut all of the components costing under £5. In total, one MouseBot costs under £150 for the materials. If a 3D printer or laser cutter are not available in-house to the researcher, many third party companies exist for which users can outsource their manufacturing. It takes around 3 hours of labour to assemble one MouseBot.

3.6.1.4 MouseBots can be used in-cage with multiple animals

Other devices created for the purpose of automating the single pellet reaching task are desktop solutions which require that animals executing the task (normally

rats) are individually housed (Fenrich et al., 2016, Wong et al., 2015, Ellens et al., 2016). Our device solves this issue by using radio frequency identification to distinguish between mice. However, we did find that the number of trials registered for each mouse housed in the same cage varied considerably. Work is now underway to introduce a physical barrier that will be programmed to limit the number of trials per mouse. Another strategy is to persist with the current design, free of any gate or physical barrier, allowing mice free access to grasp for pellets, and having the device alert users when a pre-specified number of mice have reached a pre-defined performance criteria for a preclinical trial, while also taking advantage of the variability in trial success. For example, mice who perform more trials after an injury may have a greater recovery. The variability will allow experimenters to undertake correlative and causative studies, between trial number and success rates.

3.6.1.5 Validating pellet detection and positioning

Detection of pellets in other devices has typically been achieved using non-camera physical sensors. For example, the use of infrared beams, which are interrupted by a pellet passing through it. These methods will work out-of-the-box and are easy to start up and require minimal firmware configuration, but are liable to malfunction over time because build-up of pellet and cage debris can block the signal and incorrectly classify the presence of pellets. Instead, the MouseBot has an integrated overhead camera (which avoids debris) and uses blob detection to determine where pellets are present in a frame. We omitted the beam-breakers to simplify the hardware configuration and used the camera to determine whether a pellet has been dispensed or not. Our pellet detection method, using blob detection, was found to be accurate 97% of the time. With optimisations this can be brought nearer to 100%, but is a promising starting point.

Pellet positioning is achieved using computer vision algorithms, whereby the position of the pellet in a frame, as determined by blob detection, directs the position

of the spoon until the pellet's target location is reached. To validate the positioning of the pellet, we used a feature detector trained using the DeepLabCut toolbox to return, independent of the blob detection method, the position of the pellet within the frame. Using the key-point detector trained with the DeepLabCut toolbox we show that the pellets are consistently positioned within 0.5 mm of its target location.

3.6.2 Classification of trial outcomes

Classifying the outcome of each trial automatically is a difficult engineering problem. Our previous device, the RatBot, used the combinatorial outputs of various physical sensors to classify trial outcomes (Gadiagellan, 2018), while others have created devices which output short video files that can be scored manually in a way that is quicker than traditional video scoring (Ellens et al., 2016). The videos created by the MouseBot are event synchronized: the pellet is removed at the same point. Long and short term memory neural networks are often used to forecast events, like the weather (Hewage et al., 2019). For example, in a sentence, each of the previous words and their sequence to some extent predict the next word. LSTM neural networks are a type of recurrent neural network which “remembers” the input data at the beginning of the sequence. As the data of pose-estimation are structured in a temporal sequence, with each line of data corresponding to the whereabouts of the trained features in each frame, we trained a simple LSTM neural network to attempt to classify the outcomes of trials. Initial attempts using just the location of the pellets looked promising. The model trained had an overall accuracy of 77%. We iteratively improved the model by confining its training data to the final 250 frames of the video, increasing the input features to include both the snout and middle finger of the animal, and compressing the input data to integers. The final model correctly classified 91% of trials which were manually scored as “hits” and 80% of trials which were scored as “misses”. It is possible to review the trials which were misclassified to try and gain

insight into the reasons why. However, this is a good starting point to continue iteratively improving the model. Others (data unpublished) have also used random forest classifiers for this purpose with promising results, which also classifies the number of reach attempts too. Given the simplicity of the LSTM model we trained, and the accuracy at which trials were classified based on the pellet location alone, and with the advent of TensorFlow Lite, (a version of Tensorflow designed especially for the use of machine learning on devices with limited computation power), it may be possible to classify the outcomes of trials within the Raspberry Pi itself, using just the coordinates generated from the blob detection algorithm. This would offload the computational power required to run pose-estimation.

3.6.2.1 Conclusions

The development of MouseBots demonstrates that prototyping technologies in academic environments its accessible, with 3D printing and low cost electronics. We have shown that in-cage devices can be built, and that the beginnings of a platform that will simplify the testing of reaching and grasping is nearing its final stages of development. The MouseBot has limitations, namely, controlling for the variability of the frequency of trials carried out by mice on a daily basis housed in the same cage.

However, we have shown that using low-cost prototyping technology and open-source software, bespoke devices can be manufactured in academic settings for the purpose of automating preclinical behavioural assays. In the future, we hope to share the dataset of videos such that other users can query the videos with their own scripts and leverage their own machine learning models to answer new experimental questions.

3.7 Supplementary Materials

3.7.1 Figure legends for each supplementary video

Supplementary Video 1 **A walkthrough of the MouseBot and its main mechanical components**

Supplementary Video 2 A pellet being dispensed from the hopper, onto the spoon and finally into the view of the frame.

This movie shows a top-down view to within the MouseBot's hopper. The flow control mechanism has been removed for clarity and the dividers that pattern the lower edge of the hopper can be seen shuttling sucrose pellets up and around into a chute situated at the top edge of the hopper. A sugar pellet, can then be seen caught by the spoon, having been channelled down through the dispenser's chute. It is then positioned in view of the camera, in order to determine the presence of the pellet. It is then taken to the brush, where the pellet is removed and the cycle begins again. The firmware being run in this video is a test firmware to monitor the performance of the dispenser and the sugar pellet detection.

Supplementary Video 3 A movie showing the pellet being positioned from the perspective of the camera

This movie shows a computer terminal/console and a view of a pellet being positioned. The firmware prints various outputs to the terminal. Initially there are two pellet-dispense trials where no pellet is detected, followed by a pellet that is dispensed on the spoon. A live view of a section of the camera (where blob detection is confined) can be seen in the lower right-hand corner, in which blob detection is visibly shown by the black circle. Once the pellet is dispensed and detected, the spoon undergoes the positioning loop, where the pellet can be seen moving across the field of view, then stops slightly offset to the right of the reaching slot. It is at this point frames are added to a circular buffer to later be stored into video.

Supplementary Video 4 An example video of a trial where the pellet is missed by the mouse

This movie shows a reaching trial captured by a MouseBot, where a mouse paw can be seen reaching for a pellet, twice, before dropping it. The video was captured at

approximately 180 frames per second and encoded to play back at 30 FPS (6 times slower than real time).

Supplementary Video 5 An example video of a trial where the mouse successfully reaches the pellet

This movie shows a trial where a mouse successfully reaches for a pellet in one attempt. The pellet is removed from the frame at roughly the same time point in every reaching video, meaning automatic analysis is more straightforward.

Supplementary Video 6 Animal Pose-Estimation applied to reaching videos (part 1)

To automatically analyse and classify video trials, we manually annotated features of interest in frames generated by the MouseBot.

Supplementary Video 7 Animal Pose-Estimation applied to reaching videos (part 2)

Using the trained convolutional neural network with the DeepLabCut toolbox, an example video of its output can be seen.

Table 1. Calibration document contained within an online mongoDB server which each MouseBot calls upon before initiating

FIELD	DESCRIPTION
MACADDRESS	The unique MACAddress for the MouseBot
ROI	The region of interest (ROI) field, is a list containing four elements. These are the upper left and lower right coordinates of the rectangle used to define the field of view in which to apply the blob detection algorithm.

	Note, this is not the field of view or image size of the final videos saved to the cloud.
CALIBRATED	This is the field the MouseBot checks upon start up; if this is “no”, it signifies that the MouseBot was started and is waiting for the user to calibrate the device, i.e. define the ROI. If “yes”, the MouseBot can initiate and function as normal. If the document does not exist at all, the MouseBot creates a blank calibration document with the full field of view of the camera (for the user to specify the ROI for blob detection).
UPLOADDATE	This is the timestamp of when the calibration image and its metadata were uploaded to the MongoDB instance.
MOTOR_LIMIT	A list containing 6 elements. In this order, these elements are the general purpose input/output pin allocations for each of the home switches, the I2C addresses for the Raspberry Pi to send and receive data from the stepper motor controllers, the number of seconds for the MouseBot to run the dispenser motor, and the number of seconds for the MouseBot to run the spoon motor so that it presents itself in frame of the camera.
DEVICENAME	The name of the device given by the user.
SPOONPOSITION	Defines the spoon position, and takes one of four inputs; “left”, “right”, “middle” and “smart”. The “smart”

	option is still in development stage but will optimise the test for individual mice.
DEVICESTATUS	The DeviceStatus is either “online” or “offline” and its purpose is to prevent the re-initiation of a device’s firmware when an instance is already running. Also, when a device switches status, users can opt-in to an email service which updates them of this change, along with the last message received from the device, which may include an error message. It is updated based on a script run on a remote server, separate from the Raspberry Pi which monitors the time elapsed since the last message from the device was sent to the database. When a device is running, this is updated every half a second. A status will be updated from offline to online when the last message received from the device is longer than 5 minutes.
LASTMESSAGEFROMDEV —	Last message received from the device can take any one of these messages: "No More Pellets or Jammed" "Unit closed by user" "DeviceWorking" "Pellet took too long to position, shutting down."
MESSAGERECEIVED	Messages sent to the database are time-stamped and are saved in this field.

Table 2. The metadata contained along with each video when uploaded to a MongoDB server following each trial.

Field	Description
MACAddress	The unique MACAddress for the MouseBot from which the video was captured.
Trial_Time	The date and time the trial was recorded.
Computer Result	The success of a trial as determined by an automated pipeline; this can either be “CSV Uploaded”, to indicate that pose-estimation has been executed on the footage and is ready to be analysed, “Hit”, to indicate a successful reach, or “miss”, to indicate a drop.
filename	The filename of the trial’s video.
ManualResult	The success of a trial as determined by the user. This field is for the validation of the Computer Result.
uploadDate	The upload date of the video sent to the database.
FPS	The frames per second of the recorded video.
RFID_Tag	The RFID tag associated with the video. This is the last RFID tag to have been read by the MouseBot’s RFID reader before the pellet was no longer detected.
VidError_	Vid Error is a manually updated field, it is default set to ‘n/a’, it is updated when a user notices any oddities about a video, for example, whether the pellet was positioned correctly prior to it being dislodged.

CSV_FILE	This field contains a serialised CSV file. The CSV file contains pose-estimation coordinates for each of the 700 frames generated by a convnet trained to identify the location of pellets, digits of the reaching paw and the mouse's snout.
-----------------	---

4 Using a deep convolutional neural network to automate the scoring of walking accuracy on a horizontal ladder by mice

4.1 Introduction

The horizontal ladder is a behavioural task used in pre-clinical research to assess deficits in coordination and balance (Metz and Wishaw, 2009). Briefly, a mouse or rat is placed on a horizontal ladder with irregularly spaced rungs and is video recorded as it walks across. Following stroke or spinal cord injury, animals crossing the ladder make errors in forepaw and/or hindpaw placement (Duricki et al., 2019) such as a complete slip through the rungs or incorrect placement of the paw. If the initial neurological injury is large enough then any spontaneous recovery is partial, leaving a medium-sized, persistent deficit. These features make the horizontal ladder task a good outcome measure for rehabilitative drug candidates screened in preclinical studies. Recording footage of traversing animals requires minimal manual labour from technicians and animals learn to cross the ladder in as little as two sessions, lasting ten minutes each.

The rate limiting step is the manual analysis and scoring of recorded videos. Many hours of footage are generated in a typical preclinical study where large cohorts of animals are assessed on a weekly basis; typically (in our laboratory and that of other labs) each rodent is video-recorded performing three trials per week (Duricki et al., 2016). Scoring these videos manually (for each limb separately) is cumbersome, time consuming and (because it is mundane) potentially prone to error. Automation of this task could allow researchers to record more ladder traversals per week, enabling a more precise estimate of limb function and thus a more powerful study to detect differences between groups. Commercially available devices exist to automate this analysis; these include the use of pressure sensitive rungs to count correct steps, and infrared beam-breakers positioned between rungs to detect slips such as with the

Erasmus Ladder. These solutions are expensive, require additional hardware, and cannot discern which body part breached the rungs. Automatically tracking the limbs of a mouse traversing across a horizontal ladder in a video is a difficult computer vision task. Research groups have used reflective markers on key mouse limb locations (Takeoka et al., 2014) so that they can be tracked in video footage for gait analysis; however this is a time consuming method of pose-estimation, because marker application is required for all mice in the study, which may need additional re-application.

A toolbox, DeepLabCut, is now revolutionising research in animal behaviour (Mathis et al., 2018). DeepLabCut is a toolbox that allows researchers to generate artificial neural networks that estimate the poses of animals in video footage. DeepLabCut provides the tools for researchers to create their own dataset to train an artificial neural network to detect features of interest. A dataset is created by extracting key frames from a video and manually annotating the location of features. As little as 200 annotated frames are enough to generate a dataset to train accurate models for feature detection because the convolutional neural network used (ResNet50) has already been pre-trained on a large dataset of images. We used DeepLabCut to generate a model that would detect the pose of mice crossing a horizontal ladder. We then sought to explore whether it was possible to automatically score the number of slips (for each limb separately) made by traversing mice before and after large photothrombotic cortical stroke. An ideal software package would automatically generate a table of scores from a directory of videos (each containing multiple ladder traversals) from a given mouse and also provide a method to manually validate the automatically scored videos.

Ordinarily, videos of mice or rats traversing the horizontal ladder are recorded from a camera on one side of the ladder, positioned slightly lower than the rungs so that all limbs are viewable from one point of view. We therefore sought to create software that could be applied to this common recording angle but also generalise to

other recording angles and distances from the horizontal ladder such that it could be retrospectively applied to previously recorded videos.

Our goal was to create a web-based software package that is portable and deployable on any operating system and platform and for these reasons we opted to use Google Colab. Google Colab is a cloud and browser-based Jupyter notebook that allows users to execute both Python 2 and 3 code in runtimes which leverage powerful Graphics Processing Units (NVIDIA Tesla K80s). Google Colab allows users to create short compartmentalised cells of code, which can be executed individually (rather than all at once) making it a useful tool for prototyping and data scientists. LadderScorer software, based on Google Colab, uses neural networks to estimate the pose of mice crossing the horizontal ladder which is a computationally expensive process. The use of Google Colab negates the need for users to have expensive hardware. In this chapter I describe the development and validation of this software, which we refer to as “LadderScorer”.

4.2 Methods

4.2.1 Training of artificial neural network to monitor mice crossing a horizontal ladder

A C57BL/6 mouse was video recorded using a tripod-mounted Samsung Galaxy S7 edge camera crossing a horizontal ladder (Maze Engineers) kindly provided by Prof. Elizabeth Bradbury. The video was cropped to minimise the space below and above the rungs and mouse (smaller frames reduces training and inference time). Using the DeepLabCut 2.0 toolbox in a Conda environment on Windows 10, 150 frames were manually extracted from the recorded video and labelled as outlined in Figure 25. The features annotated were the nose, the most distal part of each visible limb within a frame, the base of the tail and its tip. Once frames were annotated to generate the neural network’s training data, a PC with a dedicated graphics processing unit (GPU), an NVIDIA GTX 980 Ti, was used to train the network for

256,000 iterations. During analysis, features are labelled by the network on new frames, not included in the training dataset, if their p-cutoff values (or likelihood) surpasses 0.9. The model was refined with a further 280 frames extracted from three additional videos of the same mouse recorded with a GoPro Black 7, in linear shooting mode, at 50 frames per second, and at 1080p resolution. These additional frames contained limbs that had slipped through the rungs. Refinements to the models generated were trained using the weights of the previously trained model, rather than from scratch.



Figure 25 **A screenshot taken from a video, having been analysed by our trained pose-estimation model**

4.2.2 Development of analysis software

The software was architected such that the scores of a mouse crossing a horizontal ladder could be generated by running a Python library called “LadderAnalysis” with the path to the project’s configuration file and the file path to the video wished to be analysed. For example:

```
LadderAnalysis(config_path=config_path,full_vid_filename='/content/drive/MyDrive/MouseHorizontalLadder-Sotiris-2019-01-08/videos/00187.MTS')
```

This can be used within a loop to generate, automatically, the scores (which are

the number of errors made by each limb, and the latency for the mouse to cross, for each run) for each video within a directory.

4.2.3 The configuration file tells LadderScorer the file paths of the trained feature detector's model, so that the appropriate model is used to analyse videos.

4.2.3.1 Detecting the presence of an animal in a given frame and segmenting video footage into individual trials

Footage of mice crossing a horizontal ladder typically contains multiple traversals. Thus it was first necessary to split each video into individual runs so that metrics for each can be output separately. To do this, each video was analysed with our pose-estimation model to generate coordinates and likelihood values for a given feature. Likelihood values are a probabilistic non-zero floating point number with a maximum value of 1, where 1 means the model is confident the given feature is within a frame at the given location, and a value close to 0 means there is little confidence. A k-means clustering algorithm was used on the rolling average of each of the likelihood values generated for each feature. The rolling average of 100 frames (or 2 seconds worth of video) was used to reduce the influence of occasional spurious values generated by our pose-estimation model. Each frame was subsequently allocated to one of two different clusters (0, meaning the mouse was not in frame, and 1, meaning it was) and can be seen plot over the duration of a given video in Figure 28b. The frame numbers where a switch in the mouse's presence in the video are taken (i.e. the frame number where a 0>1 transition occurs, or vice versa), and are used to split the video into individual trials.

4.2.3.2 Calculating the latency for mice to cross the ladder

The latency for mice to cross was calculated for each run by multiplying the number of consecutive frames clustered into group 1 (meaning the mouse was likely to be in frame) by the frames per second the footage was captured (50).

4.2.3.3 Determining the average body length of the mouse

The body length was determined by calculating the Euclidean distance, $d(t, s)$, between the predicted coordinates of the snout (s) and the base of the tail (t) when their likelihood values were both greater than 0.99 in a given frame, and taking the arithmetic mean of these values across the video.

$$d(t, s) = \sqrt{(t_1 - s_1)^2 + (t_2 - s_2)^2}$$

Equation 1. **Function used to determine the distance between the tail base and the snout in a given frame**

Where t is the predicted coordinates of tail base and s is the predicted coordinates of the snout, while t_1 and t_2 are the tail base's x and y coordinates, respectively, and s_1 and s_2 are the snout's x and y coordinates, respectively.

4.2.3.4 Determining the location of the ladder's rungs

To determine a slip made by a mouse, it is necessary to know where the rungs are in a given video. Initially, this was done by having the user manually specify the two ends of the horizontal ladder's rungs within the first frame of the video. However, it was found that by plotting a line of best fit through the predicted coordinates of the mouse's snout, which is close to the rungs throughout the video, an estimate of the rung's locations could be determined (denoted by the green line in Figure 26).

4.2.3.5 Determining the number of slips made by each paw

The number of slips for a given paw was determined by counting the number of times its predicted coordinate surpassed the "slip threshold". The slip threshold was determined for an animal by dividing its average body length by 10, which was empirically found to be roughly half the length of its forelimb. If a limb's y coordinate

surpassed the line of the rungs by the slip threshold it is registered as a slip. Animals that slip through the rungs will often make subsequent slips through the same rungs as they try to regain their footing which are not scored as additional errors when scored manually. For this reason, a variable called “slipped state” was created within the software which was set to True when a given paw surpassed the slip threshold for the first time. No additional slips would be registered until the slipped state was reset to False. A reset occurs when the x coordinate of the slipped limb has moved in either direction in the x axis by half of the animal’s body length from the x coordinate of the first registered coordinate to have surpassed the slip-threshold.




Figure 26 The predicted coordinates of the back left paw overlaid on the first frame of a video

In this plot, the blue dots are the predicted coordinates of the back left paw across the length of one trial. The green line is the line of best fit through the predicted coordinates of the mouse's snout across the entire video and is used to give an approximation of the location of the rungs. Finally, the red dots denote the coordinates which have surpassed the rung line by the slip threshold.

4.2.3.6 Horizontal Ladder

Mice were habituated on a horizontal ladder by allowing them to roam freely across its rungs in either direction with their cage-mates for 15 minute sessions on 5 consecutive days. One end of the horizontal ladder was modestly aversive in that it was open and contained no dark zones, whereas the other end mimicked a home-cage environment (Figure 27), with bedding, nesting material and tunnels. During habituation, mice learned where the home-cage environment was and routinely traversed the ladder to this zone. After habituation, baseline videos were recorded. Mice were positioned at the aversive end and allowed to traverse to the home-cage environment. Once they reached the home-cage environment, we allowed them to roam freely (in this zone) for approximately 10 seconds, before placing them once again at the aversive end. Video footage of mice traversing the ladder was captured using a GoPro HD Black 7, shot in linear mode, with a resolution of 1920 x 1080 at 50 frames per second. Each video contained footage of three ladder traversals by a given mouse on a given day. The horizontal ladder was kindly loaned to us by Professor

Elizabeth Bradbury. It was obtained originally from Maze Engineers. For each run, mice were placed on the right-hand side of the horizontal ladder, and travelled from right to left for each run, in one direction.




Figure 27 **The horizontal ladder used in this study**

The box on the left side was filled with bedding and animal husbandry where animals would traverse to from the right side.

4.2.4 Unilateral photothrombotic stroke

Photothrombotic stroke was carried out as previously described in 1985 by Watson et al (Watson et al., 1985) and in Results Chapter 2 of this thesis.

4.2.5 Correlation Analysis and Agreement between manual and automated methods

To validate LadderScorer on its ability to accurately score videos of mice crossing the horizontal ladder in terms of both the latency for mice to cross, and the number of full slips made by each paw, an independent technician scored a subset of videos manually (from both pre and post-injury time points), so correlations could be drawn between automatic (LadderScorer) and manual analyses. Errors made by each limb were scored by watching back recorded footage at half speed and tallying the number of times limbs fully breached the rungs, where a slip is defined as the paw breaching below the rungs by at least approximately 50% of its full length, to coincide with the automatically generated slip-threshold. The latency for mice to cross the ladder was also calculated manually by an independent technician who used a stopwatch which was started and stopped when a mouse's entire body entered and then subsequently left the horizontal ladders' rungs, respectively. Agreement between the manual and automatic scoring of videos was determined by using Pearson's correlation analysis. For the latency for mice to cross, the values of each traversal outputted by LadderScorer were directly compared with the value determined by the human scorer, whereas, for the correlation analysis of slips, the arithmetic mean of the number of slips made for each limb across the number of runs made was compared.

4.2.7 Labelling videos with LadderScorer

In order to manually validate the predictions of LadderScorer, a function in Python was written to overlay its outputs upon the video being assessed. These outputs include the limb being assessed, and whether it is in a slipped-state or not (the user can, at present, only monitor the outputs of one limb at a time), the frame number, the predicted position of the rungs, and the slip-threshold. This script can also be executed within Google Colab and the outputted video is stored within the same directory of the video analysed within the user's Google Drive.

4.3 Results

4.3.1 Deep neural network predicts the pose of mice

We trained a neural network to recognise each paw, head, tail and tail tip of mice crossing a horizontal ladder using frames extracted from a singular video of a C57Bl/6 mouse. Video footage of a mouse traversing a ladder can be seen in supplementary video 8 with its limbs tracked as they move through the frames with coloured markers depicting their estimated location as predicted by our trained neural network. With a likelihood, or p cut-off value, set to 0.9 the trained model predicted features with an average error (relative to the human annotator) of 4.67 pixels on the test dataset (see table 3) which equates to around 1.5 mm. By comparison, within each video, the length of a mouse was between 250 - 300 pixels long.

Table 3. The evaluation results of the trained neural network used to analyse the videos for pose-estimation with a p-cut off value of 0.9

Training iterations:	Train error(px)	Test error(px)	p-cutoff used	Train error with p-cutoff (px)	Test error with p-cutoff (px)
256000	2.07	8.76	0.9	2.07	4.67

4.3.2 LadderScorer segments a video with multiple ladder crossings by a mouse into individual runs and calculates its latency to cross

With the model trained and capable of feature detection for our needs, we then sought to develop a method to accurately determine the latency for mice to cross as well as the number of mis-steps made by mice traversing the ladder.

In our laboratory and that of other groups, multiple horizontal ladder traversals from a single animal are video recorded consecutively and stored as one file. We wanted to configure a way for our software to separate each video into individual runs. To achieve this, the LadderScorer analysis pipeline begins by detecting the presence of a mouse in the video. The software is configured to do this by clustering the likelihood parameters for multiple features (snout, paws, tail-tip) outputted by the pose-estimation model into two clusters with the k-means clustering algorithm. K-means clustering is an unsupervised learning strategy which will cluster data into a pre-specified number of groups (here, two) and find the centroid within each, and categorise and label new data into the group whose centroid is closest (i.e., has a lesser Euclidean distance). In this case, one cluster represents frames in which the snout, paws and tail-tip are “all absent” whereas the other cluster represents frames in which the snout, paws and/or tail-tip are visible; a given frame is allocated a value of 0 or 1. These values, plotted against the frame number generate plots that can be seen in Figure 28b. The length of time a mouse is detectable upon the horizontal

ladder within the left and right hand limits of the camera view can be determined by the number of consecutive frames where the likelihood values are clustered within the group associated with higher likelihood values (indicating that the mouse is present within the frame); and from this the latency for mice to cross can be extrapolated.

Given a video containing multiple traversals of the horizontal ladder, LadderScorer segments the video and produces latency data for each run. Using LadderScorer to generate latency outputs, we show the software identified an increase in the time taken to traverse 7 days after injury (Figure 28c).

A correlation analysis of latencies recorded from 4 animals (each traversing the ladder multiple times) before and after photothrombotic stroke by both methods of analysis; manually and with LadderScorer was then done to validate the software. A pair-wise comparison of 25 runs, between manual and automatic scoring, yielded a Pearson's correlation coefficient of 0.99 (Figure 28d).

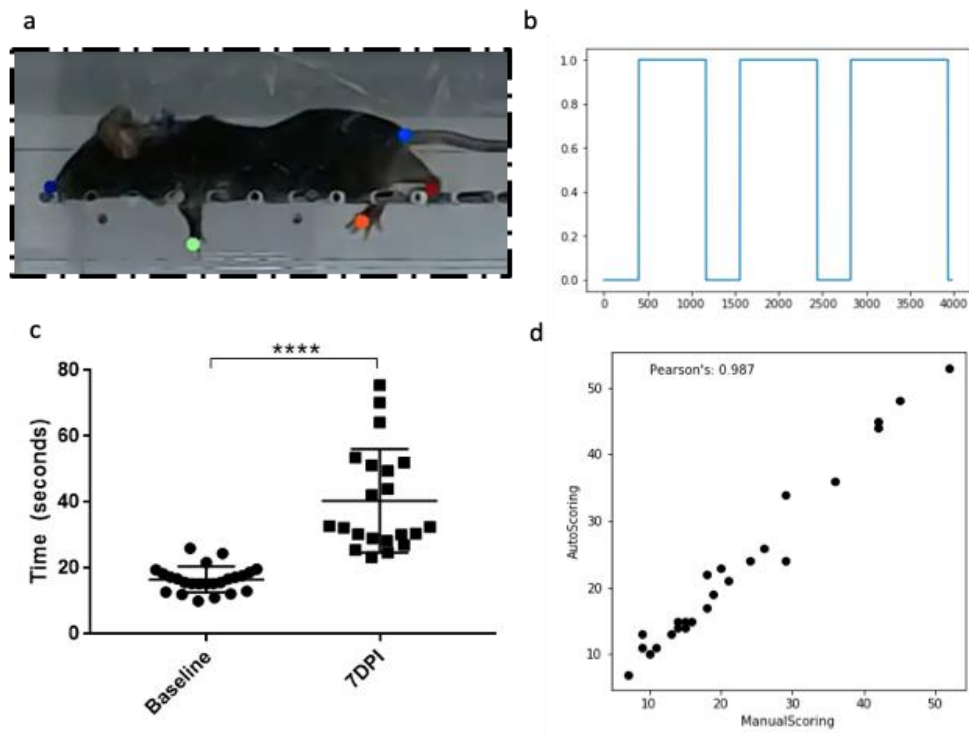


Figure 28 **Data from the feature detector can be used to segment videos into individual runs and generate latency data**

a. Photograph shows a mouse seven days after photothrombotic stroke crossing the horizontal ladder with its paws, tail base and snout automatically labelled using a trained convolutional neural network. The features labelled were the most visible distal part of each limb, the tip of the tail, the tail base and the snout.

b. A “presence plot” which indicates whether a mouse is on the horizontal ladder is present using a k-means clustering approach on the likelihood values generated by the trained feature-detector where y axis is the clustered group (0, or 1, where cluster 1 is associated with the higher likelihood values) and the x axis is the frame number.

c. The latency for mice to cross the horizontal ladder before and 7 days following photothrombotic stroke as determined by LadderScorer. Each dot represents a mouse’s average time taken to cross the ladder across three traversals (p < 0.05, two-tailed t-test.)

d. There was good agreement between the automated measurement of latency and those generated manually (Pearson’s r = 0.99).

4.3.3 LadderScorer can be used to determine the number of slips made by the affected limbs of mice crossing the horizontal ladder

In order to determine whether a mouse's limbs have slipped through the rungs, it is first necessary to determine where the horizontal plane containing the rungs is in a given video. We found that plotting a line of best fit through the predicted x and y coordinates generated for the nose (whose likelihood score is above 0.9) demarcates the horizontal plane containing the rungs in each frame (denoted by the green line, Figure 26). This meant that it was therefore not necessary for the user to specify where the rungs were in each video, or to train a neural network that would recognise the edges of each lab's separate apparatus.

With accurate rung detection, attempts were made to begin quantifying the number of slips made. It was necessary to set a "slip-threshold" (Figure 29) whereby, if a predicted paw-coordinate was detected below the rung-line, a slip was counted. We determined this threshold on a mouse-by-mouse basis (to seek to control for body size) by calculating the arithmetic mean of the mouse's body length as it traverses across the horizontal ladder from its nose to the base of its tail (equation 1), and dividing it by 10 which we empirically determined to be approximately 50% the length of a mouse's forelimb. Any predicted paw coordinates which fell below the rung line by this slip-threshold or more was therefore considered in a slipped-state. Determining the slip-threshold relative to the mouse length means that the software may cope when applied to videos recorded from different angles and distances from the ladder, with differently sized mice or rats. When mice or rats make slips through rungs, regaining footing on the rungs often leads to many subsequent slips between the same two rungs by the same paw in that trial (see supplementary video 8). It was important that the software excludes these subsequent corrective footing attempts and not classify them as additional slips. An analysis of the slips made by 21 mice in

this study by LadderScorer detected an increase in the number of slips made by the affected limbs of mice after injury (Figure 30). LadderScorer also determined that the number of slips made before injury was low.



Figure 29 **LadderScorer predicts the position of the rungs and establishes a slip-threshold**

A screenshot from a video analysed by LadderScorer: the outputs generated by LadderScorer can be overlaid upon the video from which the data was extracted such that its predictions can be manually validated. In the video, the predicted horizontal plane containing the rungs can be seen in green. This is generated by passing a line of best fit through the x and y coordinates of the mouse's snout across all the frames in the video. The blue line is an automatically generated "slip threshold". Metrics appear in the top left of the frame which outlines the traversal number (typically out of three), the limb being monitored, whether the given limb is in a slip state and the current frame number (see also supplementary video 9).




Figure 30 **LadderScorer detects a deficit in motor function after injury**

The total, combined number of slips made by forelimbs and hindlimbs were calculated per mouse with (n = 26) before and 7 days post injury using LadderScorer (P < 0.05; two-tailed t test)

4.3.4 Comparison of automated with a manual method of assessment revealed good agreement between methods for paws closest to the camera, but not those furthest away

To determine the agreement between manual scoring and automatic scoring by LadderScorer software, we used Pearson's correlation analysis on the average number of slips made for each paw per run in a given video as determined by both methods. Pearson's correlation analysis revealed a strong correlation of outputted values for the left forelimb and hindlimbs (r=0.82, and r=0.9, respectively) between manual and automatic analysis compared with the r values for the right forelimbs and hindlimbs, which were 0.4 and 0.5 respectively (Figure 31).

These results suggest a better performance by LadderScorer on the limbs in closest view of the camera. Each mouse is positioned on the right hand side of the ladder for each new traversal, and moves from right to left within each video, meaning the mouse's left side is always closest to the camera. Accordingly, a subgroup analysis was performed by confining the analysis of LadderScorer to mice who had undergone photothrombosis on the right cortical hemisphere, such that the affected limbs always appeared on the side of the horizontal ladder closest to the camera during the traverses from right to left.

Pearson's Correlation coefficient of outputted values from the manual and automatic analysis revealed an increase in correlation from 0.9 to 0.95 for the left hindlimb, and 0.82 to 0.86 for the affected forelimb (Figure 32).

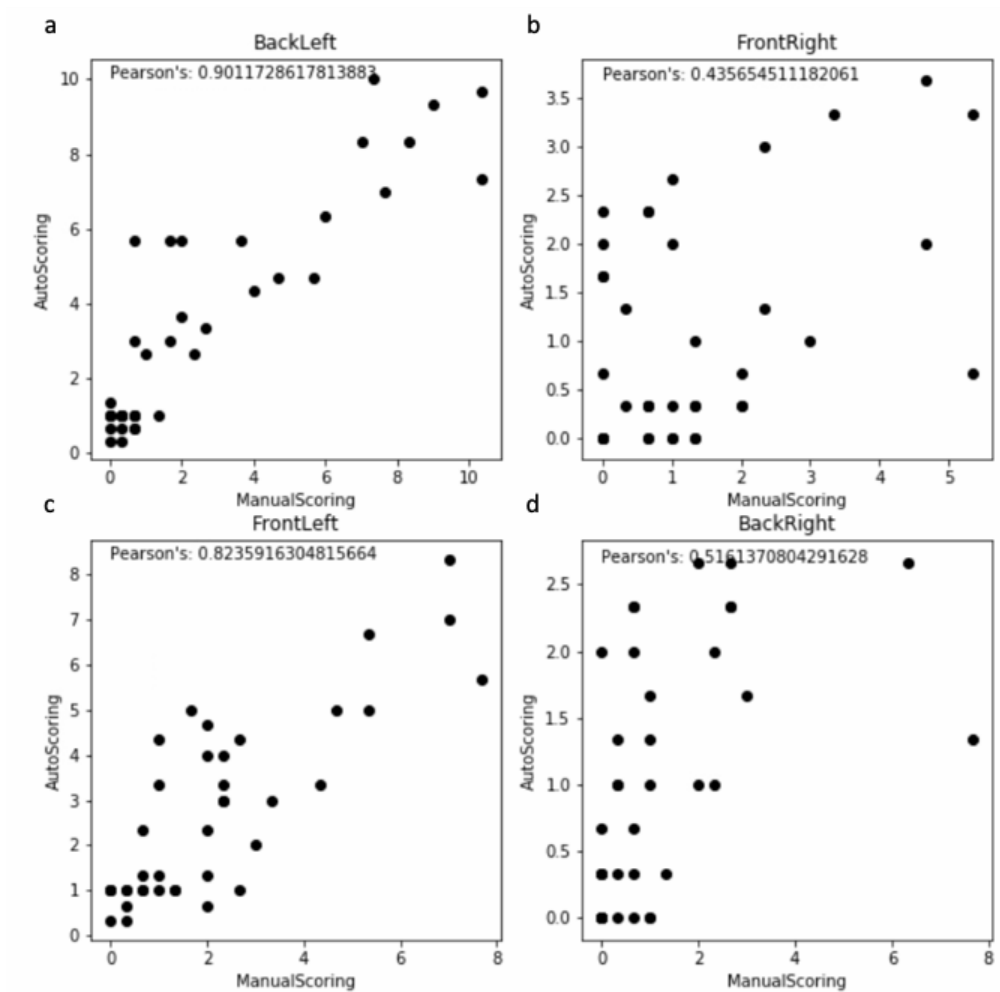


Figure 31 Agreements between manual assessment of foot-slips and LadderScorer-generated trial outcomes

The number of slips as determined between manual and automatically generated scores of (a) the left hindlimb (Pearson $r=0.90$) (b) right forelimb (Pearson $r=0.43$), (c) left forelimb (Pearson $r=0.51$) and (d) right hindlimb (Pearson $r=0.82$). To generate these plots, the average runs from 39 videos (containing 21 different mice at both pre-injury and post-injury timepoints) were computed by each method which included both pre and post-injury mouse. Each dot is the average of the number of runs (three or four) made by the mouse in a given video.




Figure 32 **LadderScorer determines slips of the affected hindlimbs and forelimbs of mice traversing the horizontal ladder**

Scatter plots with the generated results from manual analysis (x axis) vs automatic scoring (y axis), for the affected hindlimb (left) and forelimb (right), where the Pearson's correlation coefficient for each was determined to be 0.95 and 0.86, respectively. Each dot is the average number of slips made by a singular mouse (n=21, from a total of 39 videos).

4.5 Discussion

LadderScorer is a Google Colab-based software package that can be used to automatically score mice traversing a horizontal ladder. LadderScorer can segment videos into separate trials and can quantify the latency of a mouse to cross the ladder. The software automatically determines the location of the rungs and determines a 'slip-threshold' based on the body length of the animal such that the number of slips made by a mouse traversing a horizontal ladder can be quantified, which the software shows increases 7 days after photothrombotic stroke. Our data shows that LadderScorer can be used to automate the analysis of horizontal ladder footage of mice traversing ladder comparably to a human scorer before and after the onset of a unilateral photothrombotic stroke.

4.5.1 Pose-estimation is a good candidate for feature reduction for the quantification of locomotion metrics of mice traversing a horizontal ladder

A challenge in machine learning is designing good features from which to generate models. To use all of the pixels in a given frame as the input values to infer whether mice have slipped through the rungs would be computationally wasteful and expensive, and it is therefore necessary to reduce each frame into meaningful data that might help predict the incidence of slips in a more computationally manageable way. DeepLabCut and pose-estimation solves this problem by reducing the data of each frame to the predicted coordinates of key features as well as a confidence, or likelihood score, which, together, provide a map of a mouse's limbs, snout and tail. Furthermore, models can be trained with as little as 200 annotated frames, because transfer learning is leveraged, which uses pre-trained models, trained for object detection on hundreds of thousands of images of static objects including animals and humans, as a starting point. Our model accurately predicted the location of the features it was trained to detect with a pixel error of 4.67, which, in the videos used

in this study, equates to around 1.5 mm. The accuracy of this model provides us with the basis to begin extrapolating more information from the videos.

4.5.2 LadderScorer individually segments videos containing multiple ladder traversals

Video segmentation was achieved using k-means clustering on the likelihood values outputted by the trained feature detector. Our lab and others often record the multiple traversals of a given mouse at a particular experimental timepoint in one video, thus it was necessary for us to confine the analysis of slips to these sections of the video to minimise error, and to output the quantities of slips and latencies for each run. Determining the presence of a mouse within the video also gave us a means to quantify the time taken for each mouse to cross. Correlation analysis of the values outputted by our software compared to those outputted by a human scorer showed a strongly positive correlation with an r value of 0.99, which validated the software's ability to segment video footage and output accurate latency data for each traversal. While not analysed in this study, we will also be able to quantify the speed of animals with this algorithm.

4.5.3 LadderScorer automatically predicts the location of the ladder's rungs

To determine whether a limb had slipped through the rungs, we predicted the location of the latter by plotting a line of best fit through the coordinates of the mouse's snout as it tracked across the video. This line of best fit was empirically determined to be a good estimate of the rung's locations within the video. This method has yet to be formally validated against the rung's ground truth location and work is underway to do this. While the line of best fit generally superimposed well on rungs on the video footage when visually inspected, rung prediction will only be accurate so long as the camera position does not change dramatically over the course of a video, because a line of best fit is generated from the coordinates of the snout across the entire video.

After the establishment of an automatic line which demarcates the horizontal plane of the rungs in the 2D video, we then sought to determine a “slip-threshold”, which when a coordinate was detected below, a slip was counted. A goal of this software is to be able to generalise well to videos generated from different centres, whose camera, camera angle and distance from the camera will be different to those used in this study. For this reason, we calculated a slip threshold such that it was relative to the body length of the animal within the video.

4.5.4 LadderScorer predicts more accurately the number of slips made by the limbs in closest view of the camera

LadderScorer is a “hard coded” algorithm, meaning the rules which dictate whether a slip occurred were designed by a human such that the quantified slips outputted by LadderScorer correlated more so with the video scores generated by the human. This approach seems to have worked well in creating rules to determine the slips made by the affected limbs that are in full view of the one, side-on positioned camera, with correlation coefficients (r values) of 0.96 and 0.86, respectively. However, it should be known that “hard coded” rules, or, decision trees, do not often generalise well to data outside of that used to facilitate the programming of the software. For this reason, these values will need further validation with more data. These rules do not seem to work well in predicting the number of slips made by the paws furthest away from the camera, presumably because these limbs are often occluded by the body of the animal as it traverses across the ladder. To mitigate this, one may be able to use a mirror positioned such that both sides of the animal are viewable to the camera, or by using a second camera positioned on the other side of the ladder. While this approach may improve the predictions of slips used by this software on all of the limbs, it will not be possible to apply this software retrospectively to videos of horizontal ladders recorded in the classical way.

Recently, groups have begun using unsupervised learning approaches to classify different animal behaviours based on their predicted pose generated using using DeepLabCut (Hsu and Yttri, 2019, Bourached and Nachev, 2019). This strategy requires no prior data annotation further from those required for DeepLabCut, and work by clustering temporal pose-estimation data into groups which can be recognised as distinct behaviours. It may be possible to apply a similar strategy to the pose-estimation data generated in this study such that slips can be quantified in an unbiased way.

Another approach may be to use a supervised learning strategy, where instances of slips of each paw are manually annotated, and models are generated to recognise patterns in the pose-estimation data which predict more accurately when slips occur.

4.5.5 Conclusions

To conclude, we show that it is now possible to begin extracting data from video footage of mice traversing a horizontal ladder and reduce the amount of time required to manually score videos from weeks to minutes. LadderScorer requires minimal computational power to run, and can be used on basic PCs and laptops as it is Google Colab based. While the software shows promise, there are still areas which require improvement, and the next stages will be refining the software such that slips made from limbs of both sides can be scored more accurately.

4.6 Supplementary Materials

4.6.1 Figure legends for supplementary videos

Supplementary Video 8 **Markerless tracking of a mouse traversing a horizontal ladder**

Features of a mouse are tracked using a convolutional neural network trained using the DeepLabCut toolbox. The video field of view also tracks automatically according

to the position of the mouse and this is done using the estimated position of the snout and tailbase. The right hindlimb can be seen tracked with a red dot only when it comes into view and is not obscured.

Supplementary Video 9 Validation video generated by LadderScorerV1

Footage generated by LadderScorer, tracking the left hindlimb as it goes in and out of “slip state” as it passes below and above the automatically generated slip threshold based on the length of the animal within the video.

5 Conclusions and General Discussion

There exists a great need for new therapies for the growing number of stroke survivors. Upper limb function is very often affected after stroke but the options for recovery are limited. There are multiple streams of ongoing research which aim to resolve this. Exoskeletal devices, which take advantage of the remaining muscle activity to drive its actuators have shown promise in both patients who have suffered from stroke or spinal cord injury (Louie and Eng, 2016, Ho et al., 2011). Stem cell therapies, which aim to replace the lost neurons after ischaemic stroke (Borlongan, 2019), is also an area of research which is being explored with great interest. Another approach is to promote neuroplasticity, and promote the capacity for the CNS to re-wire. There are multiple strategies which aim to achieve this. A notable method, referenced throughout this text, is the use of anti-NOGO-A therapy in conjunction with rehabilitation in rats after large cortical stroke (Wahl et al., 2014), another, is the use of lipopolysaccharide (LPS), a pro-inflammatory stimulus, which reignites the CNS to become neuroplastic in chronic injury settings, again in conjunction with rehabilitation in spinal cord injured rats (Torres-Espín et al., 2018a). The strategy we are adopting makes use of an endogenous neurotrophin, NT-3, to coax corticospinal axons, spared from injury, to re-wire in such a way that re-animates limbs whose function is lost after stroke. We have shown that NT-3 has the capacity to do this in rats, when delivered using an adeno associated viral vector but there are still unknowns as to how it brings about this change.

My PhD had two principal aims: The first was to create new automated methods of evaluating the efficacy of rehabilitative drug therapies, and the second was to use these methods to begin answering how NT-3 restores the loss of motor function in mice who had undergone photothrombosis.

The analysis of animal behaviour has classically been cumbersome and time consuming for research personnel, and devices that automate these tasks are needed. With the advent of low-cost electronics, single board computers, and technologies like 3D-printing, creating bespoke devices for the purpose of automating preclinical behaviour has become more widely accessible. The resources available online and elsewhere to learn and use programming languages like Python are also ever-expanding, accessible and free. This gives access to people in academic settings who are not formally trained engineers the ability to use methods such as machine learning and computer vision to supplement the hardware users can now build. Indeed, many such devices are being made including alternative solutions to the MouseBot, detailed further in the introduction of this text. Machine learning is also increasingly being used and made more accessible to biologists. DeepLabCut is revolutionising the quantification of animal behaviour, with users now able to quantify behaviours previously not possible based on footage captured of the behaviour. For example: pupil dilation, *Drosophila* movement and limb kinematics during reaching and grasping tasks (see: <http://www.mousemotorlab.org/deeplabcut>). The novel methods I have been developing are promising, but still flawed (see below), meaning they cannot yet be used reliably for a longitudinal preclinical study.

Former lab member Dr. Dhiresan Gadiagellan sought to automate the single pellet reaching and grasping task in rats, and successfully showed that he could train, assess and detect deficits in the reaching performance of rats before and after injury. Dr Gadiagellan used CAD software and 3D-printing technology to build an Internet of Things device that could be retrofitted into standard rat housing and whose collected data could be accessed from anywhere. The RatBot used an array of hardware sensors to classify trials online. My goal was to adapt the RatBot such that it could be used with mice, but with some differences. I wanted computer vision to be the cornerstone of its success. The reason for this was to reduce the hardware required (we would no longer require infrared beams placed in various locations around the

device to detect the presence or absence of a pellet, for example) and so that we could capture and store footage of each trial. By automating the capture of and the storing of 1000s of trials organised with metadata in a database, we are able to retrospectively apply analyses to these videos, as well as share data with others such that animal use (the repetition of studies) can be reduced.

Using the OpenCV Python library, I was able to successfully dispense and position pellets (as the camera was used to detect their presence and live position), monitor the pellet's presence through a trial while simultaneously recording the footage of an active trial at high speed (180 frames per second) thousands of times. All of this was possible using a very low cost computer [Raspberry Pi 3 (£30)] and we were able to automatically run thousands of trials and store these videos in cloud storage. We were then able to show that with machine learning, we could automatically classify the trials of these videos as successful or unsuccessful reaches: first, we trained a convolutional neural network, which is the gold standard statistical modelling method for image classification, to predict and identify the positions of features such as the snout, pellet, and middle fingers of mice executing a reaching trial, then used these predictions to train a second neural network, Long Short Term Memory, to classify the overall outcome of the trial. Our proof-of-concept study showed that we were able to successfully classify 91% and 80% of hits and misses respectively, which is a strong starting point to continue optimising (e.g., by including additional input features [for example, the whereabouts of the pellet as determined by blob detection], and by fine tuning the hyperparameters of the LSTM).

The groundworks for the automation of the single pellet reaching task have been implemented, but work is still required to make the device more reliable for use in a longitudinal preclinical trial. An issue we encountered is that mice housed together interacted with the MouseBot to varying degrees. We will need to implement a physical barrier such that the number of trials carried out by mice is kept as consistent as possible. The RFID reader integrated within the device tells us which mouse

interacts with the device and can be used to count the daily trials executed by each one. This information can be used to limit the number of trials each mouse can do.

DeepLabCut also inspired us to begin automating the analysis and scoring of the footage captured of mice and rats traversing the horizontal ladder. The horizontal ladder is a tool typically used to assess, fine motor control and proprioception, which are all affected after stroke. We used DeepLabCut to train a convolutional neural network to predict the whereabouts of the limbs of mice traversing a horizontal ladder, with a view to quantifying the number of slips made by each paw. This is a task that typically takes hundreds of hours to manually carry out, and its automation could save researchers' time, increase the throughput of experiments and make their scoring more objective. We created a Google Colab notebook (LadderScorerV1), to begin making sense of the pose-estimation of a convolutional neural network trained using the DeepLabCut toolbox. For example, could we use it to quantify locomotion metrics of an injured mouse traversing a ladder?

The latency data outputted seemed promising, and correlated positively and strongly (Pearson's correlation = 0.99) against manual latency measurement. We were able to use this metric to quantify the latency of mice traversing the horizontal ladder in our longitudinal NT-3 study. Where LadderScorerV1 currently falls short is its ability to quantify the number of slips made by mice, particularly on the paws contralateral the camera (because they can be occluded by the animal's body when they are above the rungs). We used a decision tree to determine the number of slips made by the animal, which worked well for the limbs in full view of the camera, but not on the limbs on the other side. We found that using an additional machine learned model on the footage gathered by the MouseBot was a good way of classifying the outcomes of trials, and could be a good strategy moving forwards to count the number of slips made by animal for each individual paw. As LadderScorerV1 is implemented on a Google Colab notebook (i.e., with its interface placed on a website and with its data analysed using GPUs in the cloud), it carries out its analysis using the freely

available GPU resources, and is completely portable. This negates the need for research personnel needing expensive GPU resources of their own. With additional refinements and optimisations, it may be possible to fully automate the analysis of footage recorded of mice traversing a horizontal ladder.

With these refinements made, to both the MouseBot and LadderScorer, we hope to hasten and make more objective the collection and analysis of data collected from preclinical trials of novel rehabilitative drug candidates. We also seek to use them to help us understand how NT-3 reverses the loss of motor function.

5.1.1.1 Future Directions concerning MouseBot

Work is underway to improve the MouseBot. Namely, and for reasons mentioned in previous sections, to simplify the device. The MouseBot has many components which take a significant amount of time to source, manufacture, and assemble. This is problematic for a number of reasons. While prototyping has become feasible in academic settings, without specialist equipment it is difficult to manufacture multi-part devices at scale. Our vision is to have high-throughput, large-scale training of many tens or hundreds of mice simultaneously. Simplifying the device will enable the process of production and deployment. Also, as there are multiple moving components within the device, the likelihood of wear and tear, as well as jams, are likely. With a system containing many tens, perhaps hundreds, of devices working simultaneously, there will likely be a need for maintenance as it exists in its current form.

Currently, the interface to interact with the MouseBot is primitive, and requires that users are familiar with Linux terminal commands and Python to alter key parameters concerning how the devices operate. A graphic user interface will also be made for users to more easily interact with the devices. This will include a method to change the experimental design, the difficulty of the task, and methods to download the videos for manual inspection and the data that the implemented algorithms produce.

We are also implementing mirrors within the device so that we can extract the 3D pose of the animal from one camera; this will enable us to analyse the 3D kinematics of the reach-and-grasp movement. We plan to use Anipose, a Python package that has been written to automatically extract the 3D pose of animals from a variety of 2D cameras (Karashchuk et al., 2020).

5.1.1.2 Future Directions concerning LadderScorer

One of the key issues that came to light during the development of the LadderScorer is that the performance of the algorithm is far better at detecting foot slips made by the limbs closest to the camera. We are addressing this in a number of ways. The first is to improve the predictive model which detects the whereabouts of the limb key points in a given frame, with the hope that the transient slips made by the limbs furthest away from the camera are detected with more accuracy.

Ultimately, we would like refine this algorithm to ensure that it is as generalised as possible; for users of other labs to be able to use this algorithm retrospectively, with their previously recorded videos. However, we concede that this may not be possible, so are exploring options that may not make this algorithm backwards compatible. One such solution is the use of a rail, upon which a camera can be mounted, and then moved, either automatically or manually, parallel to and at the same velocity as the mouse or rat as it crosses the ladder. The benefit of this is that the camera can be positioned closer to the animal, zoomed-in, with more of its field of view dedicated to the animal (rather than the full length of the ladder), and therefore the resolution of the mouse within the frame is far higher, and thus lessening the challenge of creating a model that performs feature detection well.

Mentioned previously in the text is the exploration of using multiple cameras with the algorithm as is. As it performs well detecting slips on the limbs closest to the camera, by placing a second camera on the other side of the ladder, ensuring that both sides are recorded, we should be able to accurately count the slips on both sides.

An alternative to this is the use of a mirror, positioned along the ladder in such a way which enables that both sides of the mouse are captured from one view of the camera. This will require some algorithmic adjustments but may be a simpler solution to having two cameras, and requires less space.

Work to date has also focused on counting foot slips on mice traversing the horizontal ladder, however, arguably, the horizontal ladder is more often used with rats. For this reason we have also begun work automating the quantification of foot slips of rats crossing the horizontal ladder, which are larger animals and therefore an easier subject for a model to predict its features.

To summarise, we are adopting many different strategies to address the flaws the algorithm currently has.

5.1.1.3 Future directions concerning the elucidation of NT-3's rehabilitative effects

Work is also planned to further the mechanistic understanding of NT-3's rehabilitative properties. The experiment we are prioritising at present aims to determine whether functionally meaningful neuroplasticity occurs in the CSTs of mice after unilateral stroke, having been treated with NT-3. We were unable to show a functional benefit of NT-3 when administered subcutaneously and in the future we will explore whether rehabilitation can be induced by way of a gene-therapy approach, such as with our previous rat studies (Duricki et al., 2016). Briefly, this will involve injecting AAV1-CMV-NT3 into the affected forelimb and hindlimb of mice 24 hours after induction of cortical stroke by way of photothrombosis.

With a mouse model of NT-3-induced functional recovery after photothrombotic stroke established, we will then begin the dissection of its pathway by way of using our previously outlined genetic strategies.

In conclusion, work is well underway to address the need for solutions that automate behavioural tasks that interrogate the motor function of mice, particularly that of their upper limbs. The MouseBot is showing promise as a viable alternative solution to manually training and assessing mice at the single pellet reaching task. LadderScorer, with refinements, will also offer researchers the option to automate the analysis and scoring of horizontal ladder videos. With these two solutions, we hope to continue elucidating the mechanism through which NT-3 promotes a functional recovery after unilateral cortical stroke so that it can be exploited in a clinical setting.

6 References

- AHMAD, M., J DAR, N., S BHAT, Z., HUSSAIN, A., SHAH, A., LIU, H. & H GRAHAM, S. 2014. Inflammation in ischemic stroke: mechanisms, consequences and possible drug targets. *CNS & Neurological Disorders-Drug Targets (Formerly Current Drug Targets-CNS & Neurological Disorders)*, 13, 1378-1396.
- ALBERS, G. W., MARKS, M. P., KEMP, S., CHRISTENSEN, S., TSAI, J. P., ORTEGA-GUTIERREZ, S., MCTAGGART, R. A., TORBEY, M. T., KIM-TENSER, M. & LESLIE-MAZWI, T. 2018. Thrombectomy for stroke at 6 to 16 hours with selection by perfusion imaging. *New England Journal of Medicine*, 378, 708-718.
- ALLRED, R. P. & JONES, T. A. 2008. Maladaptive effects of learning with the less-affected forelimb after focal cortical infarcts in rats. *Experimental neurology*, 210, 172-181.
- ALLRED, R. P., MALDONADO, M. A., HSU AND, J. E. & JONES, T. A. 2005. Training the "less-affected" forelimb after unilateral cortical infarcts interferes with functional recovery of the impaired forelimb in rats. *Restor. Neurol. Neurosci.*, 23, 297-302.
- ALTAR, C. A., ANTHONY ALTAR, C., SIUCIAK, J. A., WRIGHT, P., IP, N. Y., LINDSAY, R. M. & WIEGAND, S. J. 1994. In Situ Hybridization of trkB and trkC Receptor mRNA in Rat Forebrain and Association with High-affinity Binding of [125I]BDNF, [125I]NT-4/5 and [125I]NT-3. *European Journal of Neuroscience*, 6, 1389-1405.
- AMANTEA, D., MICIELI, G., TASSORELLI, C., CUARTERO, M. I., BALLESTEROS, I., CERTO, M., MORO, M. A., LIZASOAIN, I. & BAGETTA, G. 2015. Rational modulation of the innate immune system for neuroprotection in ischemic stroke. *Frontiers in neuroscience*, 9, 147.
- ARDESCH, D. J., BALBI, M. & MURPHY, T. H. 2017. Automated touch sensing in the mouse tapered beam test using Raspberry Pi. *J. Neurosci. Methods*, 291, 221-226.
- BACHMANN, L. C., LINDAU, N. T., FELDER, P. & SCHWAB, M. E. 2014. Sprouting of brainstem-spinal tracts in response to unilateral motor cortex stroke in mice. *J. Neurosci.*, 34, 3378-3389.
- BAIRD, A. L., MELDRUM, A. & DUNNETT, S. B. 2001. The staircase test of skilled reaching in mice. *Brain Res. Bull.*, 54, 243-250.
- BALKAYA, M., KRÖBER, J. M., REX, A. & ENDRES, M. 2013. Assessing post-stroke behavior in mouse models of focal ischemia. *J. Cereb. Blood Flow Metab.*, 33, 330-338.
- BEJOT, Y., BENATRU, I., ROUAUD, O., FROMONT, A., BESANCENOT, J. P., MOREAU, T. & GIROUD, M. 2007. Epidemiology of stroke in Europe: geographic and environmental differences. *J. Neurol. Sci.*, 262, 85-88.
- BIERNASKIE, J., CHERNENKO, G. & CORBETT, D. 2004. Efficacy of rehabilitative experience declines with time after focal ischemic brain injury. *J. Neurosci.*, 24, 1245-1254.
- BLACKWELL, A. A., BANOVETZ, M. T., QANDEEL, WHISHAW, I. Q. & WALLACE, D. G. 2018. The structure of arm and hand movements in a spontaneous and food rewarded on-line string-pulling task by the mouse. *Behavioural Brain Research*, 345, 49-58.
- BORLONGAN, C. V. 2019. Concise review: stem cell therapy for stroke patients: are we there yet? *Stem cells translational medicine*, 8, 983-988.
- BOUËT, V., FRERET, T., TOUTAIN, J., DIVOUX, D., BOULOUARD, M. & SCHUMANN-BARD, P. 2007. Sensorimotor and cognitive deficits after

- transient middle cerebral artery occlusion in the mouse. *Exp. Neurol.*, 203, 555-567.
- BOURACHED, A. & NACHEV, P. 2019. Unsupervised Videographic Analysis of Rodent Behaviour. *arXiv [cs.CV]*.
- BRADBURY, E. J., KHEMANI, S., VON R, KING, PRIESTLEY, J. V. & MCMAHON, S. B. 1999. NT-3 promotes growth of lesioned adult rat sensory axons ascending in the dorsal columns of the spinal cord. *European Journal of Neuroscience*, 11, 3873-3883.
- BROEKS, J. G., LANKHORST, G. J., RUMPING, K. & PREVO, A. J. H. 1999. The long-term outcome of arm function after stroke: results of a follow-up study. *Disability and Rehabilitation*, 21, 357-364.
- BRUNO, V., BATTAGLIA, G., COPANI, A., D'ONOFRIO, M., DI IORIO, P., DE BLASI, A., MELCHIORRI, D., FLOR, P. J. & NICOLETTI, F. 2001. Metabotropic glutamate receptor subtypes as targets for neuroprotective drugs. *J. Cereb. Blood Flow Metab.*, 21, 1013-1033.
- BURNSIDE, E. R., DE WINTER, F., DIDANGELOS, A., JAMES, N. D., ANDREICA, E.-C., LAYARD-HORSFALL, H., MUIR, E. M., VERHAAGEN, J. & BRADBURY, E. J. 2018. Immune-evasive gene switch enables regulated delivery of chondroitinase after spinal cord injury. *Brain*, 141, 2362-2381.
- CANDELARIO-JALIL, E. 2009. Injury and repair mechanisms in ischemic stroke: considerations for the development of novel neurotherapeutics. *Curr. Opin. Investig. Drugs*, 10, 644-654.
- CHAUDHRY, V., GIULIANI, M., PETTY, B. G., LEE, D., SEYEDSADR, M., HILT, D. & CORNBATH, D. R. 2000. Tolerability of recombinant-methionyl human neurotrophin-3 (r-metHuNT3) in healthy subjects. *Muscle Nerve*, 23, 189-192.
- CHEN, C.-C., GILMORE, A. & ZUO, Y. 2014. Study motor skill learning by single-pellet reaching tasks in mice. *J. Vis. Exp.*
- CHEN, H.-H., TOURTELLOTTE, W. G. & FRANK, E. 2002. Muscle spindle-derived neurotrophin 3 regulates synaptic connectivity between muscle sensory and motor neurons. *J. Neurosci.*, 22, 3512-3519.
- CHEN, X., YE, H., KURUVILLA, R., RAMANAN, N., SCANGOS, K. W., ZHANG, C., JOHNSON, N. M., ENGLAND, P. M., SHOKAT, K. M. & GINTY, D. D. 2005. A chemical-genetic approach to studying neurotrophin signaling. *Neuron*, 46, 13-21.
- CÔTÉ, M.-P., AZZAM, G. A., LEMAY, M. A., ZHUKAREVA, V. & HOULÉ, J. D. 2011. Activity-dependent increase in neurotrophic factors is associated with an enhanced modulation of spinal reflexes after spinal cord injury. *J. Neurotrauma*, 28, 299-309.
- CRAMER, S. C. 2008. Repairing the human brain after stroke: I. Mechanisms of spontaneous recovery. *Annals of Neurology*, 63, 272-287.
- CRAMER, S. C. & CHOPP, M. 2000. Recovery recapitulates ontogeny. *Trends in Neurosciences*, 23, 265-271.
- DEL ZOPPO, G. J., MILNER, R., MABUCHI, T., HUNG, S., WANG, X., BERG, G. I. & KOZIOL, J. A. 2007. Microglial activation and matrix protease generation during focal cerebral ischemia. *Stroke*, 38, 646-651.
- DENG, J., DONG, W., SOCHER, R., LI, L.-J., LI, K. & FEI-FEI, L. 2009. ImageNet: A large-scale hierarchical image database. *2009 IEEE Conference on Computer Vision and Pattern Recognition*.
- DMITRIEVA, V., STAVCHANSKY, V., POVAROVA, O., SKVORTSOVA, V., LIMBORSKA, S. & DERGUNOVA, L. 2016. Effects of ischemia on the expression of neurotrophins and their receptors in rat brain structures outside the lesion site, including on the opposite hemisphere. *Molecular Biology*, 50, 684-692.
- DONG, Y. & BENVENISTE, E. N. 2001. Immune function of astrocytes. *Glia*, 36, 180-190.

- DREIER, J. P. 2011. The role of spreading depression, spreading depolarization and spreading ischemia in neurological disease. *Nat. Med.*, 17, 439-447.
- DRIEU, A., LEVARD, D., VIVIEN, D. & RUBIO, M. 2018. Anti-inflammatory treatments for stroke: from bench to bedside. *Therapeutic advances in neurological disorders*, 11, 1756286418789854.
- DURICKI, D., DRNDARSKI, S., BERNANOS, M., WOOD, T., BOSCH, K., CHEN, Q., DAVID SHINE, H., SIMMONS, C., WILLIAMS, S. C. R., MCMAHON, S. B., BEGLEY, D. J., CASH, D. & MOON, L. D. F. 2018a. Corticospinal neuroplasticity and sensorimotor recovery in rats treated by infusion of neurotrophin-3 into disabled forelimb muscles started 24 h after stroke.
- DURICKI, D., DRNDARSKI, S., BERNANOS, M., WOOD, T., BOSCH, K., CHEN, Q., SHINE, H. D., SIMMONS, C., WILLIAMS, S. C. R., MCMAHON, S. B., BEGLEY, D. J., CASH, D. & MOON, L. D. F. 2019. Stroke Recovery in Rats after 24-Hour-Delayed Intramuscular Neurotrophin-3 Infusion. *Ann. Neurol.*, 85, 32-46.
- DURICKI, D., HUTSON, T. H., KATHE, C., SOLEMAN, S., GONZALEZ-CARTER, D., PETRUSKA, J. C., SHINE, H. D., CHEN, Q., WOOD, T. C., BERNANOS, M., CASH, D., WILLIAMS, S. C. R., GAGE, F. H. & MOON, L. D. F. 2016. Delayed intramuscular human neurotrophin-3 improves recovery in adult and elderly rats after stroke. *Brain*, 139, 259-275.
- DURICKI, D., KAKANOS, S., HAENZI, B., CHRISTINA, W., CASH, D. & MOON, L. 2018b. Subcutaneous neurotrophin-3 infusion induces corticospinal neuroplasticity and improvements in dexterity and walking in elderly rats after large cortical strokes. *bioRxiv*.
- ELLENS, D. J., GAIDICA, M., TOADER, A., PENG, S., SHUE, S., JOHN, T., BOVA, A. & LEVENTHAL, D. K. 2016. An automated rat single pellet reaching system with high-speed video capture. *J. Neurosci. Methods*, 271, 119-127.
- ERNFORS, P., LEE, K. F., KUCERA, J. & JAENISCH, R. 1994. Lack of neurotrophin-3 leads to deficiencies in the peripheral nervous system and loss of limb proprioceptive afferents. *Cell*, 77, 503-512.
- FARIÑAS, I., WILKINSON, G. A., BACKUS, C., REICHARDT, L. F. & PATAPOUTIAN, A. 1998. Characterization of neurotrophin and Trk receptor functions in developing sensory ganglia: direct NT-3 activation of TrkB neurons in vivo. *Neuron*, 21, 325-334.
- FARR, T. D. & WHISHAW, I. Q. 2002. Quantitative and qualitative impairments in skilled reaching in the mouse (*Mus musculus*) after a focal motor cortex stroke. *Stroke*, 33, 1869-1875.
- FEIGIN, V. L., FOROUZANFAR, M. H., KRISHNAMURTHI, R., MENSAH, G. A., CONNOR, M., BENNETT, D. A., MORAN, A. E., SACCO, R. L., ANDERSON, L., TRUELSEN, T., O'DONNELL, M., VENKETASUBRAMANIAN, N., BARKER-COLLO, S., LAWES, C. M. M., WANG, W., SHINOHARA, Y., WITT, E., EZZATI, M., NAGHAVI, M., MURRAY, C., GLOBAL BURDEN OF DISEASES, I., RISK FACTORS, S. & THE, G. B. D. S. E. G. 2014. Global and regional burden of stroke during 1990-2010: findings from the Global Burden of Disease Study 2010. *Lancet*, 383, 245-254.
- FENRICH, K. K., MAY, Z., TORRES-ESPÍN, A., FORERO, J., BENNETT, D. J. & FOUAD, K. 2016. Single pellet grasping following cervical spinal cord injury in adult rat using an automated full-time training robot. *Behav. Brain Res.*, 299, 59-71.
- FERRARA, A., EL BEJAOU, S., SEYEN, S., TIRELLI, E. & PLUMIER, J.-C. 2009. The usefulness of operant conditioning procedures to assess long-lasting deficits following transient focal ischemia in mice. *Behav. Brain Res.*, 205, 525-534.

- FOX, G., GALLACHER, D., SHEVDE, S., LOFTUS, J. & SWAYNE, G. 1993. Anatomic variation of the middle cerebral artery in the Sprague-Dawley rat. *Stroke*, 24, 2087-92; discussion 2092-3.
- FRANCESCHINI, M., FUGAZZARO, S., AGOSTI, M., SOLA, C., DI CARLO, A., CECCONI, L., FERRO, S. & ITALIAN STUDY GROUP ON IMPLEMENTATION OF STROKE, C. 2018. Acute Phase Predictors of 6-Month Functional Outcome in Italian Stroke Patients Eligible for In-Hospital Rehabilitation. *Am. J. Phys. Med. Rehabil.*, 97, 467-475.
- GADIAGELLAN, D. 2018. *Neural repair and robotics: the in-cage automation of a test for dexterity after stroke*. King's College London.
- GELDERBLUM, M., LEYPOLDT, F., STEINBACH, K., BEHRENS, D., CHOE, C.-U., SILER, D. A., ARUMUGAM, T. V., ORTHEY, E., GERLOFF, C. & TOLOSA, E. 2009. Temporal and spatial dynamics of cerebral immune cell accumulation in stroke. *Stroke*, 40, 1849-1857.
- GELDERBLUM, M., WEYMAR, A., BERNREUTHER, C., VELDEN, J., ARUNACHALAM, P., STEINBACH, K., ORTHEY, E., ARUMUGAM, T. V., LEYPOLDT, F. & SIMOVA, O. 2012. Neutralization of the IL-17 axis diminishes neutrophil invasion and protects from ischemic stroke. *Blood, The Journal of the American Society of Hematology*, 120, 3793-3802.
- GILL, D. & VELTKAMP, R. 2016. Dynamics of T cell responses after stroke. *Current opinion in pharmacology*, 26, 26-32.
- GILL, L. C., GRANSEE, H. M., SIECK, G. C. & MANTILLA, C. B. 2016. Functional recovery after cervical spinal cord injury: Role of neurotrophin and glutamatergic signaling in phrenic motoneurons. *Respir. Physiol. Neurobiol.*, 226, 128-136.
- GLIEM, M., MAUSBERG, A. K., LEE, J. I., SIMIANTONAKIS, I., VAN ROOIJEN, N., HARTUNG, H. P. & JANDER, S. 2012. Macrophages prevent hemorrhagic infarct transformation in murine stroke models. *Annals of neurology*, 71, 743-752.
- GRILL, R., MURAI, K., BLESCH, A., GAGE, F. H. & TUSZYNSKI, M. H. 1997. Cellular delivery of neurotrophin-3 promotes corticospinal axonal growth and partial functional recovery after spinal cord injury. *J. Neurosci.*, 17, 5560-5572.
- HAN, H. S. & YENARI, M. A. 2003. Cellular targets of brain inflammation in stroke. *Current opinion in investigational drugs (London, England: 2000)*, 4, 522-529.
- HAO, X.-Z., YIN, L.-K., ZHANG, X.-X., TIAN, J.-Q., LI, C.-C., FENG, X.-Y., JIANG, M. & YANG, Y.-M. 2016. Combining systemic and stereotactic MEMRI to detect the correlation between gliosis and neuronal connective pathway at the chronic stage after stroke. *Journal of neuroinflammation*, 13, 156.
- HATTIANGADY, B., RAO, M. S., SHETTY, G. A. & SHETTY, A. K. 2005. Brain-derived neurotrophic factor, phosphorylated cyclic AMP response element binding protein and neuropeptide Y decline as early as middle age in the dentate gyrus and CA1 and CA3 subfields of the hippocampus. *Experimental Neurology*, 195, 353-371.
- HEWAGE, P., BEHERA, A., TROVATI, M. & PEREIRA, E. Long-Short Term Memory for an Effective Short-Term Weather Forecasting Model Using Surface Weather Data. IFIP International Conference on Artificial Intelligence Applications and Innovations, 2019. Springer, 382-390.
- HO, N., TONG, K., HU, X., FUNG, K., WEI, X., RONG, W. & SUSANTO, E. An EMG-driven exoskeleton hand robotic training device on chronic stroke subjects: task training system for stroke rehabilitation. 2011 IEEE international conference on rehabilitation robotics, 2011. IEEE, 1-5.
- HSU, A. I. & YTTTRI, E. A. 2019. B-SOiD: An Open Source Unsupervised Algorithm for Discovery of Spontaneous Behaviors.

- HUNTER, D. V., SMAILA, B. D., LOPES, D. M., TAKATOH, J., DENK, F. & RAMER, M. S. 2018. Advillin Is Expressed in All Adult Neural Crest-Derived Neurons. *eNeuro*, 5.
- HURST, J. L. & WEST, R. S. 2010. Taming anxiety in laboratory mice. *Nat. Methods*, 7, 825-826.
- HUTCHINSON, K. J., GÓMEZ-PINILLA, F., CROWE, M. J., YING, Z. & BASSO, D. M. 2004. Three exercise paradigms differentially improve sensory recovery after spinal cord contusion in rats. *Brain*, 127, 1403-1414.
- IKONOMIDOU, C. & TURSKI, L. 2002. Why did NMDA receptor antagonists fail clinical trials for stroke and traumatic brain injury? *Lancet Neurol.*, 1, 383-386.
- INSAFUTDINOV, E., PISHCHULIN, L., ANDRES, B., ANDRILUKA, M. & SCHIELE, B. 2016. DeeperCut: A Deeper, Stronger, and Faster Multi-person Pose Estimation Model. *Computer Vision – ECCV 2016*, 34-50.
- IOSA, M., MORONE, G., RAGAGLINI, M., FUSCO, A. & PAOLUCCI, S. 2013. Motor strategies and bilateral transfer in sensorimotor learning of patients with subacute stroke and healthy subjects. A randomized controlled trial. *Eur J Phys Rehabil Med*, 49, 291-9.
- JIMENEZ-ANDRADE, J. M., HERRERA, M. B., GHILARDI, J. R., VARDANYAN, M., MELEMEDJIAN, O. K. & MANTYH, P. W. 2008. Vascularization of the Dorsal Root Ganglia and Peripheral Nerve of the Mouse: Implications for Chemical-Induced Peripheral Sensory Neuropathies. *Molecular Pain*, 4, 1744-8069.
- KARASHCHUK, P., RUPP, K. L., DICKINSON, E. S., SANDERS, E., AZIM, E., BRUNTON, B. W. & TUTHILL, J. C. 2020. Anipose: a toolkit for robust markerless 3D pose estimation. *bioRxiv*, 2020.05.26.117325.
- KATHE, C., HUTSON, T. H., MCMAHON, S. B. & MOON, L. D. F. 2016. Intramuscular Neurotrophin-3 normalizes low threshold spinal reflexes, reduces spasms and improves mobility after bilateral corticospinal tract injury in rats. *Elife*, 5.
- KEEFE, K. M., SHEIKH, I. S. & SMITH, G. M. 2017. Targeting Neurotrophins to Specific Populations of Neurons: NGF, BDNF, and NT-3 and Their Relevance for Treatment of Spinal Cord Injury. *Int. J. Mol. Sci.*, 18.
- KILIC, E., KILIC, U., BACIGALUPPI, M., GUO, Z., ABDALLAH, N. B., WOLFER, D. P., REITER, R. J., HERMANN, D. M. & BASSETTI, C. L. 2008. Delayed melatonin administration promotes neuronal survival, neurogenesis and motor recovery, and attenuates hyperactivity and anxiety after mild focal cerebral ischemia in mice. *J. Pineal Res.*, 45, 142-148.
- KLEIN, A., SACREY, L.-A. R., WHISHAW, I. Q. & DUNNETT, S. B. 2012. The use of rodent skilled reaching as a translational model for investigating brain damage and disease. *Neurosci. Biobehav. Rev.*, 36, 1030-1042.
- KLEIN, R., SILOS-SANTIAGO, I., SMEYNE, R. J., LIRA, S. A., BRAMBILLA, R., BRYANT, S., ZHANG, L., SNIDER, W. D. & BARBACID, M. 1994. Disruption of the neurotrophin-3 receptor gene *trkC* eliminates Ia muscle afferents and results in abnormal movements. *Nature*, 368, 249-251.
- KWAKKEL, G., KOLLEN, B. J., VAN DER GROND, J. & PREVO, A. J. H. 2003. Probability of regaining dexterity in the flaccid upper limb: impact of severity of paresis and time since onset in acute stroke. *Stroke*, 34, 2181-2186.
- LABAT-GEST, V. & TOMASI, S. 2013. Photothrombotic ischemia: a minimally invasive and reproducible photochemical cortical lesion model for mouse stroke studies. *JoVE (Journal of Visualized Experiments)*, e50370.
- LAMBALLE, F., KLEIN, R. & BARBACID, M. 1991. *trkC*, a new member of the *trk* family of tyrosine protein kinases, is a receptor for neurotrophin-3. *Cell*, 66, 967-979.
- LANGHORNE, P., COUPAR, F. & POLLOCK, A. 2009. Motor recovery after stroke: a systematic review. *The Lancet Neurology*, 8, 741-754.
- LAU, J., MINETT, M. S., ZHAO, J., DENNEHY, U., WANG, F., WOOD, J. N. & BOGDANOV, Y. D. 2011. Temporal control of gene deletion in sensory ganglia

- using a tamoxifen-inducible Advillin-Cre-ERT2 recombinase mouse. *Mol. Pain*, 7, 100.
- LAWRENCE, E. S., COSHALL, C., DUNDAS, R., STEWART, J., RUDD, A. G., HOWARD, R. & WOLFE, C. D. A. 2001. Estimates of the Prevalence of Acute Stroke Impairments and Disability in a Multiethnic Population. *Stroke*, 32, 1279-1284.
- LEE, J., FRIESE, A., MIELICH, M., SIGRIST, M. & ARBER, S. 2012. Scaling proprioceptor gene transcription by retrograde NT3 signaling. *PLoS One*, 7, e45551.
- LEE, J.-K., PARK, M.-S., KIM, Y.-S., MOON, K.-S., JOO, S.-P., KIM, T.-S., KIM, J.-H. & KIM, S.-H. 2007. Photochemically induced cerebral ischemia in a mouse model. *Surgical Neurology*, 67, 620-625.
- LEONARD, A. S., SOREN LEONARD, A., PURANAM, R. S., HELGAGER, J., LIU, G. & MCNAMARA, J. O. 2012. Conditional deletion of TrkC does not modify limbic epileptogenesis. *Epilepsy Research*, 102, 126-130.
- LI, H., ZHANG, N., LIN, H.-Y., YU, Y., CAI, Q.-Y., MA, L. & DING, S. 2014. Histological, cellular and behavioral assessments of stroke outcomes after photothrombosis-induced ischemia in adult mice. *BMC neuroscience*, 15, 58.
- LIESZ, A., SURI-PAYER, E., VELTKAMP, C., DOERR, H., SOMMER, C., RIVEST, S., GIESE, T. & VELTKAMP, R. 2009. Regulatory T cells are key cerebroprotective immunomodulators in acute experimental stroke. *Nature medicine*, 15, 192.
- LIPTON, P. 1999. Ischemic Cell Death in Brain Neurons. *Physiological Reviews*, 79, 1431-1568.
- LO, E. H., DALKARA, T. & MOSKOWITZ, M. A. 2003. Mechanisms, challenges and opportunities in stroke. *Nature Reviews Neuroscience*, 4, 399-414.
- LONGA, E. Z., WEINSTEIN, P. R., CARLSON, S. & CUMMINS, R. 1989. Reversible middle cerebral artery occlusion without craniectomy in rats. *stroke*, 20, 84-91.
- LONGO, F. M. & MASSA, S. M. 2013. Small-molecule modulation of neurotrophin receptors: a strategy for the treatment of neurological disease. *Nat Rev Drug Discov*, 12, 507-25.
- LOPES, G., NOGUEIRA, J., DIMITRIADIS, G., MENENDEZ, J. A., PATON, J. J. & KAMPFF, A. R. 2017. A robust role for motor cortex. *bioRxiv*.
- LOUIE, D. R. & ENG, J. J. 2016. Powered robotic exoskeletons in post-stroke rehabilitation of gait: a scoping review. *Journal of neuroengineering and rehabilitation*, 13, 53.
- MACLELLAN, C. L., KEOUGH, M. B., GRANTER-BUTTON, S., CHERNENKO, G. A., BUTT, S. & CORBETT, D. 2011. A critical threshold of rehabilitation involving brain-derived neurotrophic factor is required for poststroke recovery. *Neurorehabil. Neural Repair*, 25, 740-748.
- MANTILLA, C. B., GRANSEE, H. M., ZHAN, W.-Z. & SIECK, G. C. 2013. Motoneuron BDNF/TrkB signaling enhances functional recovery after cervical spinal cord injury. *Experimental Neurology*, 247, 101-109.
- MANTILLA, C. B., GREISING, S. M., STOWE, J. M., ZHAN, W.-Z. & SIECK, G. C. 2014. TrkB kinase activity is critical for recovery of respiratory function after cervical spinal cord hemisection. *Exp. Neurol.*, 261, 190-195.
- MANZANARES, G., BRITO-DA-SILVA, G. & GANDRA, P. G. 2018. Voluntary wheel running: patterns and physiological effects in mice. *Braz. J. Med. Biol. Res.*, 52, e7830.
- MATHIS, A., MAMIDANNA, P., CURY, K. M., ABE, T., MURTHY, V. N., MATHIS, M. W. & BETHGE, M. 2018. DeepLabCut: markerless pose estimation of user-defined body parts with deep learning. *Nat. Neurosci.*, 21, 1281-1289.
- MAZNYCZKA, A., SEN, S., COOK, C. & FRANCIS, D. P. 2015. The ischaemic constellation: an alternative to the ischaemic cascade-implications for the validation of new ischaemic tests. *Open Heart*, 2, e000178.

- MCMAHON, S. B., ARMANINI, M. P., LING, L. H. & PHILLIPS, H. S. 1994. Expression and coexpression of Trk receptors in subpopulations of adult primary sensory neurons projecting to identified peripheral targets. *Neuron*, 12, 1161-1171.
- METZ, G. A. & WHISHAW, I. Q. 2009. The ladder rung walking task: a scoring system and its practical application. *J. Vis. Exp.*
- MOON, S.-K., ALAVERDASHVILI, M., CROSS, A. R. & WHISHAW, I. Q. 2009. Both compensation and recovery of skilled reaching following small photothrombotic stroke to motor cortex in the rat. *Exp. Neurol.*, 218, 145-153.
- MU, X., SILOS-SANTIAGO, I., CARROLL, S. L. & SNIDER, W. D. 1993. Neurotrophin receptor genes are expressed in distinct patterns in developing dorsal root ganglia. *J. Neurosci.*, 13, 4029-4041.
- MURAGAKI, Y., TIMOTHY, N., LEIGHT, S., HEMPSTEAD, B. L., CHAO, M. V., TROJANOWSKI, J. Q. & LEE, V. M. 1995. Expression of trk receptors in the developing and adult human central and peripheral nervous system. *J. Comp. Neurol.*, 356, 387-397.
- MURASE, K., IGARASHI, K. & HAYASHI, K. 1994. Neurotrophin-3 (NT-3) levels in the developing rat nervous system and in human samples. *Clinica Chimica Acta*, 227, 23-36.
- MURPHY, T. H., BOYD, J. D., BOLAÑOS, F., VANNI, M. P., SILASI, G., HAUPT, D. & LEDUE, J. M. 2016. High-throughput automated home-cage mesoscopic functional imaging of mouse cortex. *Nat. Commun.*, 7, 11611.
- NESTO, R. W. & KOWALCHUK, G. J. 1987. The ischemic cascade: temporal sequence of hemodynamic, electrocardiographic and symptomatic expressions of ischemia. *Am. J. Cardiol.*, 59, 23C-30C.
- NOORSHAMS, O., BOYD, J. D. & MURPHY, T. H. 2017. Automating mouse weighing in group homecages with Raspberry Pi micro-computers. *J. Neurosci. Methods*, 285, 1-5.
- O'COLLINS, V. E., MACLEOD, M. R., DONNAN, G. A., HORKY, L. L., VAN DER WORP, B. H. & HOWELLS, D. W. 2006. 1,026 experimental treatments in acute stroke. *Ann. Neurol.*, 59, 467-477.
- PARKMAN, H. P., RAO, S. S. C., REYNOLDS, J. C., SCHILLER, L. R., WALD, A., MINER, P. B., LEMBO, A. J., GORDON, J. M., DROSSMAN, D. A., WALTZMAN, L., STAMBLER, N., CEDARBAUM, J. M. & FUNCTIONAL CONSTIPATION STUDY, I. 2003. Neurotrophin-3 improves functional constipation. *Am. J. Gastroenterol.*, 98, 1338-1347.
- PICASCIA, A., GRIMALDI, V., IANNONE, C., SORICELLI, A. & NAPOLI, C. 2015. Innate and adaptive immune response in stroke: Focus on epigenetic regulation. *Journal of neuroimmunology*, 289, 111-120.
- PODUSLO, J. F. & CURRAN, G. L. 1996. Permeability at the blood-brain and blood-nerve barriers of the neurotrophic factors: NGF, CNTF, NT-3, BDNF. *Brain Res. Mol. Brain Res.*, 36, 280-286.
- PRABHAKARAN, S., ZARAHN, E., RILEY, C., SPEIZER, A., CHONG, J. Y., LAZAR, R. M., MARSHALL, R. S. & KRAKAUER, J. W. 2008. Inter-individual variability in the capacity for motor recovery after ischemic stroke. *Neurorehabil. Neural Repair*, 22, 64-71.
- PRICE, C., MENON, D., PETERS, A., BALLINGER, J., BARBER, R., BALAN, K., LYNCH, A., XUEREBO, J., FRYER, T. & GUADAGNO, J. 2004. Cerebral neutrophil recruitment, histology, and outcome in acute ischemic stroke: an imaging-based study. *Stroke*, 35, 1659-1664.
- SAHENK, Z. 2007. Pilot Clinical Trial of NT-3 in CMT1A Patients. *Progress in Neurotherapeutics and Neuropsychopharmacology*, 2, 97-108.
- SCHNELL, L., SCHNEIDER, R., KOLBECK, R., BARDE, Y. A. & SCHWAB, M. E. 1994. Neurotrophin-3 enhances sprouting of corticospinal tract during development and after adult spinal cord lesion. *Nature*, 367, 170-173.

- SHANBHAG, N. C., HENNING, R. H. & SCHILLING, L. 2016. Long-term survival in permanent middle cerebral artery occlusion: a model of malignant stroke in rats. *Sci. Rep.*, 6, 28401.
- SHERRINGTON, C. S. 1913. Reflex inhibition as a factor in the co-ordination of movements and postures. *Quarterly Journal of Experimental Physiology: Translation and Integration*, 6, 251-310.
- SHIMAZU, K., ZHAO, M., SAKATA, K., AKBARIAN, S., BATES, B., JAENISCH, R. & LU, B. 2006a. NT-3 facilitates hippocampal plasticity and learning and memory by regulating neurogenesis. *Learning & Memory*, 13, 307-315.
- SHIMAZU, K., ZHAO, M., SAKATA, K., AKBARIAN, S., BATES, B., JAENISCH, R. & LU, B. 2006b. NT-3 facilitates hippocampal plasticity and learning and memory by regulating neurogenesis. *Learn. Mem.*, 13, 307-315.
- SKAPER, S. D. 2012. The neurotrophin family of neurotrophic factors: an overview. *Methods Mol. Biol.*, 846, 1-12.
- SOLEMAN, S., YIP, P. K., DURICKI, D. A. & MOON, L. D. F. 2012. Delayed treatment with chondroitinase ABC promotes sensorimotor recovery and plasticity after stroke in aged rats. *Brain*, 135, 1210-1223.
- SPINK, A. J., TEGELENBOSCH, R. A. J., BUMA, M. O. S. & J., L. P. 2001. The EthoVision video tracking system—A tool for behavioral phenotyping of transgenic mice. *Physiology & Behavior*, 73, 731-744.
- TAKEO, C., NAKAMURA, S., TANAKA, T., UCHIDA, D., NOGUCHI, Y., NAGAO, T., SAITO, Y. & TATSUNO, I. 2003. Rat cerebral endothelial cells express trk C and are regulated by neurotrophin-3. *Biochem. Biophys. Res. Commun.*, 305, 400-406.
- TAKEOKA, A. & ARBER, S. 2019. Functional Local Proprioceptive Feedback Circuits Initiate and Maintain Locomotor Recovery after Spinal Cord Injury. *Cell Rep.*, 27, 71-85.e3.
- TAKEOKA, A., VOLLENWEIDER, I., COURTINE, G. & ARBER, S. 2014. Muscle spindle feedback directs locomotor recovery and circuit reorganization after spinal cord injury. *Cell*, 159, 1626-1639.
- TENNANT, K. A. 2014. Thinking outside the brain: structural plasticity in the spinal cord promotes recovery from cortical stroke. *Exp Neurol*, 254, 195-9.
- TENNANT, K. A. & JONES, T. A. 2009. Sensorimotor behavioral effects of endothelin-1 induced small cortical infarcts in C57BL/6 mice. *J. Neurosci. Methods*, 181, 18-26.
- TESSAROLLO, L. 1998. Pleiotropic Functions of Neurotrophins in Development. *Cytokine & Growth Factor Reviews*, 9, 125-137.
- TESSAROLLO, L., TSOUFAS, P., DONOVAN, M. J., PALKO, M. E., BLAIR-FLYNN, J., HEMPSTEAD, B. L. & PARADA, L. F. 1997. Targeted deletion of all isoforms of the trkC gene suggests the use of alternate receptors by its ligand neurotrophin-3 in neuronal development and implicates trkC in normal cardiogenesis. *Proceedings of the National Academy of Sciences*, 94, 14776-14781.
- TESSAROLLO, L., TSOUFAS, P., MARTIN-ZANCA, D., GILBERT, D. J., JENKINS, N. A., COPELAND, N. G. & PARADA, L. F. 1993. trkC, a receptor for neurotrophin-3, is widely expressed in the developing nervous system and in non-neuronal tissues. *Development*, 118, 463-475.
- TESSAROLLO, L., VOGEL, K. S., PALKO, M. E., REID, S. W. & PARADA, L. F. 1994. Targeted mutation in the neurotrophin-3 gene results in loss of muscle sensory neurons. *Proc. Natl. Acad. Sci. U. S. A.*, 91, 11844-11848.
- TORRES-ESPÍN, A., FORERO, J., FENRICH, K. K., LUCAS-OSMA, A. M., KRAJACIC, A., SCHMIDT, E., VAVREK, R., RAPOSO, P., BENNETT, D. J. & POPOVICH, P. G. 2018a. Eliciting inflammation enables successful rehabilitative training in chronic spinal cord injury. *Brain*, 141, 1946-1962.

- TORRES-ESPÍN, A., FORERO, J., FENRICH, K. K., LUCAS-OSMA, A. M., KRAJACIC, A., SCHMIDT, E., VAVREK, R., RAPOSO, P., BENNETT, D. J., POPOVICH, P. G. & FOUAD, K. 2018b. Eliciting inflammation enables successful rehabilitative training in chronic spinal cord injury. *Brain*, 141, 1946-1962.
- WAHL, A. S., OMLOR, W., RUBIO, J. C., CHEN, J. L., ZHENG, H., SCHROTER, A., GULLO, M., WEINMANN, O., KOBAYASHI, K., HELMCHEN, F., OMMER, B. & SCHWAB, M. E. 2014. Asynchronous therapy restores motor control by rewiring of the rat corticospinal tract after stroke. *Science*, 344, 1250-1255.
- WANG, Q., TANG, X. N. & YENARI, M. A. 2007. The inflammatory response in stroke. *Journal of neuroimmunology*, 184, 53-68.
- WANG, X., LIU, Y., LI, X., ZHANG, Z., YANG, H., ZHANG, Y., WILLIAMS, P. R., ALWAHAB, N. S. A., KAPUR, K., YU, B., ZHANG, Y., CHEN, M., DING, H., GERFEN, C. R., WANG, K. H. & HE, Z. 2017. Deconstruction of Corticospinal Circuits for Goal-Directed Motor Skills. *Cell*, 171, 440-455.e14.
- WARD, N. S. 2017. Restoring brain function after stroke — bridging the gap between animals and humans. *Nature Reviews Neurology*, 13, 244-255.
- WATSON, B. D., DIETRICH, W. D., BUSTO, R., WACHTEL, M. S. & GINSBERG, M. D. 1985. Induction of reproducible brain infarction by photochemically initiated thrombosis. *Ann. Neurol.*, 17, 497-504.
- WHISHAW, I. Q. 1996. An endpoint, descriptive, and kinematic comparison of skilled reaching in mice (*Mus musculus*) with rats (*Rattus norvegicus*). *Behav. Brain Res.*, 78, 101-111.
- WHISHAW, I. Q. 2000. Loss of the innate cortical engram for action patterns used in skilled reaching and the development of behavioral compensation following motor cortex lesions in the rat. *Neuropharmacology*, 39, 788-805.
- WONG, C. C., RAMANATHAN, D. S., GULATI, T., WON, S. J. & GANGULY, K. 2015. An automated behavioral box to assess forelimb function in rats. *J. Neurosci. Methods*, 246, 30-37.
- WRIGHT, D. E., ZHOU, L., KUCERA, J. & SNIDER, W. D. 1997. Introduction of a neurotrophin-3 transgene into muscle selectively rescues proprioceptive neurons in mice lacking endogenous neurotrophin-3. *Neuron*, 19, 503-517.
- XING, C., ARAI, K., LO, E. H. & HOMMEL, M. 2012. Pathophysiologic cascades in ischemic stroke. *Int. J. Stroke*, 7, 378-385.
- YING, Z., ROY, R. R., EDGERTON, V. R. & GÓMEZ-PINILLA, F. 2003. Voluntary exercise increases neurotrophin-3 and its receptor TrkC in the spinal cord. *Brain Res.*, 987, 93-99.
- ZHOU, L., BAUMGARTNER, B. J., HILL-FELBERG, S. J., MCGOWEN, L. R. & SHINE, H. D. 2003. Neurotrophin-3 expressed in situ induces axonal plasticity in the adult injured spinal cord. *J. Neurosci.*, 23, 1424-1431.