Article

Built to Last: Functional and Structural Mechanisms in the Moth Olfactory Network Mitigate Effects of Neural Injury

Charles B. Delahunt [1,*], Pedro D. Maia [2,*] and J. Nathan Kutz [1]

1   Department of Applied Mathematics, University of Washington, Seattle, WA 98195-3925, USA; kutz@uw.edu
2   Department of Mathematics, University of Texas at Arlington, Arlington, TX 76019, USA; pedro.maia@uta.edu
*   Correspondence: delahunt@uw.edu, pedro.maia@uta.edu

Abstract: Most organisms suffer neuronal damage throughout their lives, which can impair performance of core behaviors. Their neural circuits need to maintain function despite injury, which in particular requires preserving key system outputs. In this work, we explore whether and how certain structural and functional neuronal network motifs act as injury mitigation mechanisms. Specifically, we examine how *(i)* Hebbian learning, *(ii)* high levels of noise, and *(iii)* parallel inhibitory and excitatory connections contribute to the robustness of the olfactory system in the *Manduca sexta* moth. We simulate injuries on a detailed computational model of the moth olfactory network calibrated to data. The injuries are modeled on focal axonal swellings, a ubiquitous form of axonal pathology observed in traumatic brain injuries and other brain disorders. Axonal swellings effectively compromise spike train propagation along the axon, reducing the effective neural firing rate delivered to downstream neurons. All three of the network motifs examined significantly mitigate the effects of injury on readout neurons, either by reducing injury's impact on readout neuron responses or by restoring these responses to pre-injury levels. These motifs may thus be partially explained by their value as adaptive mechanisms to minimize the functional effects of neural injury. More generally, robustness to injury is a vital design principle to consider when analyzing neural systems.

Keywords: neuronal injury; injury mitigation; focal axonal swellings (FAS); moth olfactory network

1. Introduction

Injuries are inevitable for most organisms, yet maintaining a satisfactory level of functionality can be decisive for their survival. The progressive wear of a honeybee's wings, for example, challenges the insect to sustain its load lift or face less nourishing foraging trips [1,2]. Functional robustness is desirable for neural systems as well. For example, while computer devices operate in a regime of near-zero tolerance for physical damage, the middle-aged human brain undergoes significant neuronal losses on a daily basis [3]. Robustness to injury is often overlooked when analyzing the purpose and function of neural structures while the transmission of maximum information, high signal-to-noise ratio, and low energy consumption are primarily considered [4]. Analyzing neural information processing in the context of these principles is certainly important, but arguably incomplete. The goal of this work is to examine whether certain neural mechanisms and architectural structures can be understood as adaptive, built-in systems for robustness to brain injury from trauma, aging, and/or other disorders. That is, we examine how biological neural systems are "built to last". In particular, we explore how certain neural architectures can protect the system's key downstream outputs (the "deliverables" of the system) from the effects of damage to upstream regions.

The olfactory system of the *Manduca sexta* moth, though simple, shares many neural structures and mechanisms with higher organisms [5,6]. These include (i) Hebbian plasticity, (ii) reward-triggered stimulation of neural outputs via neuromodulators, (iii) high noise levels, and (iv) inhibitory feed-forward channels running parallel to excitatory

channels. It is thus an ideal model organism to investigate the injury mitigation effects of these elements. MothNet is a computational model of this olfactory network which incorporates known biophysical parameters and which was calibrated to firing rate data recorded during in vivo learning tasks [7]. MothNet models the food odor-responsive part of the moth's olfactory network, not the sex pheromone-responsive macroglomerular complex, since the MGC has substantially distinct targets and dynamics (despite small overlaps [8]). Also, MothNet models learning-induced Hebbian updates in the synaptic connections of existing neurons, not neurogenesis via for example Juvenile Hormone [9]. See Figure 1 for a system schematic.
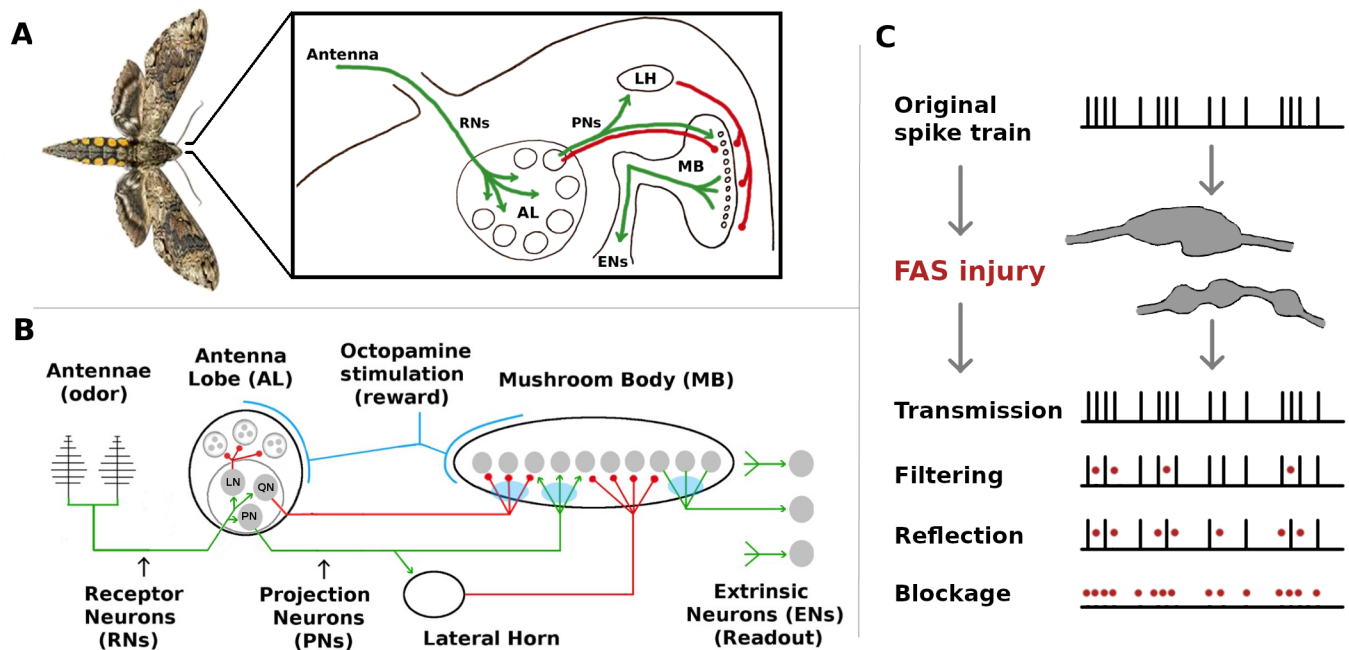

Figure 1. Overview of the Moth Olfactory Network (MON) and axonal injury mechanisms. (**A**,**B**): The MON is organized as a feedforward cascade of three distinct subnetworks and a reward mechanism [10,11]. For a brief summary of the MON, see Section 5.1. For full details, see [7]. (**C**) Focal Axonal Swellings (FAS) are ubiquitous across all severities of traumatic brain injuries and present in other leading brain disorders. They can cause some or all neural spikes in the train to die off in transit, reducing the overall firing rate arriving at the downstream target neuron. Adding FAS-like effects to the MON are the basis of our damage/injury protocols. See Section 5 for details.

The moth olfactory network (MON) also contains well-defined readout neurons (ENs, for Extrinsic Neurons), which are downstream outputs that deliver key actionable encodings to the rest of its body [12,13]. From a functional viewpoint, internal damage is unimportant as long as the key outputs (readouts) of the system are preserved. Thus, to examine injury mitigation effects we ran in silico simulations of neural injury on the MothNet model, and measured how the firing rates (FRs) of readout ENs were affected by injuries and by injury-mitigation mechanisms.

Maia et al. [14] recently introduced a computational model for the cellular level effects that may distort firing rates due to major forms of neuronal injury. They posit phenomenological input/output rules to transform healthy neuronal spike train responses into injured ones, with filters that can be either discrete-time (for spike trains) or continuous-time (for firing rates) signal processors. These filters were derived modeling the effects of demyelination and Focal Axonal Swellings (FAS), which are present in a broad array of neurological disorders [15–18].

Figure 1c exemplifies how a FAS-like injury distorts the propagation of spike trains along the axon, effectively blocking or filtering signals encoded to downstream neurons. In this work, we are agnostic concerning the exact biological underpinnings and pathological

mechanisms that may affect an injured/aging moth. Instead, we simply posit that its neurons might be exposed to detrimental effects that can affect their signaling capabilities. In this sense, and as explained in Maia et al. [14], FAS-based filters provide a more nuanced way to model neuronal malfunction than purely binary ablation which treats a neuron and/or its connections as either fully functional or 100% impaired. Recent computational studies that consider the effects of FAS-like injury in neural networks are providing new insight to decision-making deficits [19], learning impairments [20,21], memory deterioration [22], and motor-function decline [23].

While FAS models effects at the level of spike trains, it also has a meaningful representation in Firing Rate models such as MothNet. In particular, unlike ablation FAS causes reduced but still non-zero FRs. In addition, the low-pass filtering effect of FAS, which impacts closely-bunched clusters of spikes more than sparse spikes, in analogous manner impacts high FRs more strongly than low FRs. Thus FAS, applied in a FR model, results in neuron FRs being reduced but not ablated, with high FR neurons affected more strongly than low FR neurons. For a fixed amount of total damage, FAS results in relatively many partially-damaged neurons, while ablation results in relatively fewer fully-destroyed neurons.

In our simulations, we varied the parameters of each network structure-under-test; applied FAS-type injuries to different subnetworks of the system (simulating the outcome of a traumatic brain injury or concussion); and assessed the net effects on EN outputs. We examined two aspects of FR behavior for a single representative EN: (i) changes in raw FR, and (ii) changes in the ability of the EN to discriminate between a trained and untrained odor.

Our experiments led to four main findings concerning injury-mitigation structures in the moth olfactory system:

1. The learning mechanism, based on the combination of octopamine stimulation and Hebbian growth, can restore both the magnitude and discriminative ability of downstream readout neurons after upstream neurons are injured.
2. The presence of inhibitory neurons parallel to excitatory neurons connecting the subnetworks can mitigate the effects of injury, possibly via a "canceling out" effect.
3. A broad noise envelope on neural firing rates protects EN responses from effects of upstream injury, by enabling the strongest neural signals delivered downstream to still exceed action-triggering thresholds.
4. Simple ablation injury in upstream regions produces downstream effects distinct from the more biologically plausible FAS injury. That is, ablation may be a poor proxy for naturalistic injury in some neural systems.

Concerning item 1 (learning as an injury mitigation mechanism), we note that while it is intuitive that Hebbian plasticity [24,25] might *help* repair an injured network, it does not follow that Hebbian updates *alone* can repair damage: Hebbian "fire together, wire together" updates are proportional to the FRs of both the incoming and the receiving neurons. Hebbian plasticity thus requires that the upstream neurons have sufficiently strong FRs, but precisely these FRs are reduced by injury. Our experiments indicate that octopamine-induced stimulation of the injured upstream neurons is crucial to post-injury plasticity, because it boosts the FRs of injured upstream neurons back to levels that enable non-trivial Hebbian updates. Without such stimulation, the FRs of the injured neurons, as well as the FRs induced by the injured neurons in downstream neurons, are too low to induce gains in synaptic strength via a Hebbian mechanism [7].

Our computational approach allowed us to quantify the mitigating effect of a neural structure-under-test as a function of injury level, injury location, and structure parameters. We recognize that, as in any computational model, specific quantitative outcomes necessarily depend on the particular parameters and assumptions of the MothNet model. However, MothNet's architecture is tethered to known biophysical findings, and its parameter values are calibrated to biophysical findings and in vivo FR data [7]. We believe this approach enables our experimental results to refer back meaningfully to the biological structure.

2. Results

Throughout this work, we targeted two distinct regions with our injury protocols: *(i)* the Antennae and *(ii)* the channel between the Antennal Lobe (AL) and the Mushroom Body (MB). Their specific locations are shown in Figure 2.
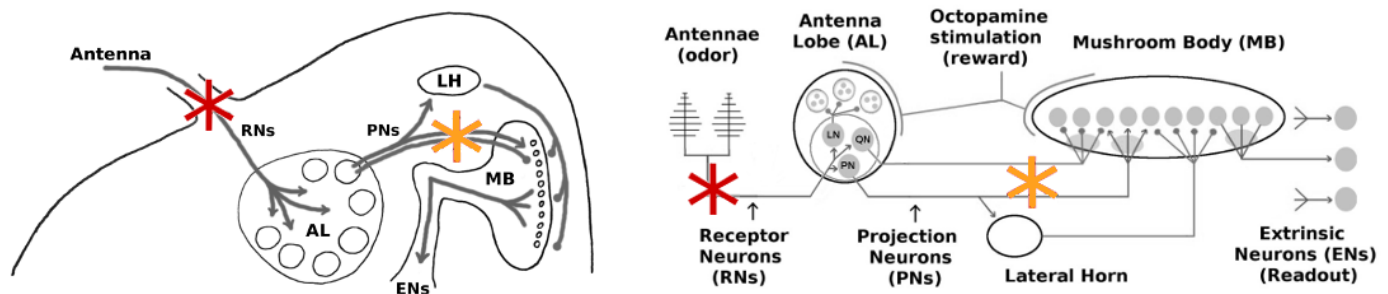

Figure 2. Location of injuries in experiments. Damage to Antennae affects the Receptor Neurons (RNs) (red stars) and reduces the overall input to the AL. Damage to the AL→MB channel (orange stars) weakens the signals passed by both excitatory projection neurons (PNs) and inhibitory projection neurons (QNs).

(i) The antennae comprise the outermost region of the olfactory system and are arguably the most exposed to external environmental shocks. Damage in this location should affect primarily the Receptor Neuron (RN) subpopulation (∼30,000). We note that the hundreds of RNs responsive to a given odor are spread throughout the antennae, ensuring that localized damage to an antenna does not disproportionately reduce the response to a particular odor.

(ii) The AL→MB channel is internal and is a center for signal transfer in the network. Damage in this location would affect both excitatory projection neurons (PNs) and inhibitory projection neurons (QNs) that link the AL to the MB.

2.1. Plasticity-Induced Recovery from Injury

The Moth Olfactory Network (MON) contains some plastic synaptic connections, and it can learn [26]: In response to reward (sugar at the proboscis), a large neuron releases octopamine widely in the AL and MB. Octopamine stimulation of FRs, combined with Hebbian-like synaptic updates, strengthens the plastic synaptic connections in the AL→MB and MB→EN channels, resulting in stronger EN responses to stimuli.

The goal of this set of experiments was to examine how far the Hebbian learning mechanism can compensate for neural injury, i.e., how far it can operate as a mitigation (as opposed to a learning) mechanism. In the first experiment, RNs in the Antennae→AL channel were injured (Figure 2, red stars). In the second experiment, PNs in the AL→MB channel were targeted (Figure 2 orange stars). AL noise was set to naturalistic levels (calibrated per in vivo data [7]) and FAS-like injury levels ranged from 0% to 60%. The MON was subsequently retrained with 5 odor puffs, close to sufficient to max out the allowable synaptic weights. The average EN readout response was recorded, as a key measure of the actionable output of the system. A typical timecourse is shown in Figure 3. In each experiment, over 30 (n = 34–38, mean = 36) MothNet instances were generated from template (i.e., a specification of network parameters used to randomly generate MothNet instances) and tested at each injury level.
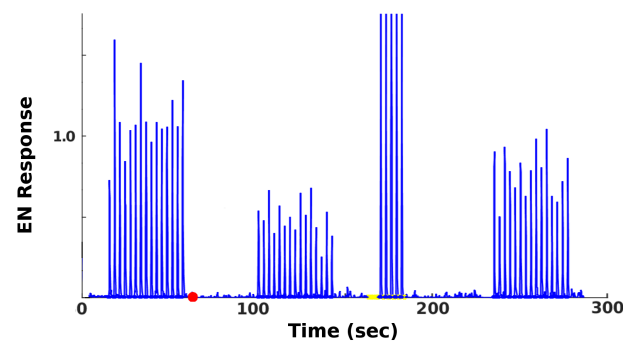
Figure 3. Typical EN timecourse. Readouts from the EN in a typical experiment, in which injury attenuated the EN odor response, and training partly restored it. *x*-axis is time. *y*-axis gives magnitude of EN response (dimensionless units), with pre-injury response to the odor ≈ 1 (absent odor, EN response ≈ 0). Events (with times in parentheses) are: Naive response (20–55); injury (red dot at 60); injured response (100–150); 5 puffs training (170–190), with the very strong (cropped) responses due to octopamine; post-training response (240–280).

We examined two properties of the EN readout:

(i) Magnitude of EN output, a basic property relevant to triggering behavioral response to an odor. Examining effects on EN FR magnitude required no pre-training of the network prior to injury.

(ii) Discriminative ability between two odors (one trained and one control), measured as the Fisher linear discriminant

$$Fd = \frac{\mu_{train} - \mu_{control}}{0.5(\sigma_{train} + \sigma_{control})} \tag{1}$$

where μ, σ are the mean and std. dev. of EN responses to trained and control odors. In these simulations, prior to injury one odor was trained so the system could discriminate the trained odor vs. control odor ($Fd \approx 5$). We either used two randomly-generated odor profiles with broad, overlapping projections onto the AL; or {odor + noise} vs. noise, with mean noise magnitude set between 0.2 and 1.0 times the odor magnitude (in MothNet, odor magnitude is controlled by a scalar which multiplies unit-length odor vectors before they are inputted to the RNs).

Results were as follows:

(i) EN magnitude: As expected, injury reduced raw EN response magnitude, and training restored some of this loss. The MON was much more robust to RN (antennae) damage than to PN damage. Complete restoration was achieved (on average) for injury levels below 25% for RN damage and below 8% for PN damage. At these injury values, injury reduced EN odor responses to approximately 70% of the naive baseline, and training restored them to baseline. See the plots in Figure 4A,B. We note that at low injury levels, the system was able to boost EN output by about 140% to 150%, a value constrained by the model's saturation parameter for the synaptic connection weights. At high levels of injury to PNs, however, the learning mechanism's ability to recover EN performance decreased (see green curves in Figure 4).

(ii) Discrimination: Injury affected EN discriminative ability much less than it did EN magnitude, because it impacted EN responses to both trained odor and control. Given two odors, injury reduced discrimination, while retraining readily restored all losses (Fisher discriminant plots in Figure 4C,D. Between {odor + noise} *vs* noise, injury had no effect on discriminative ability at any noise level, likely because the sparsely-firing MB is an effective noise filter [7]. Post-injury training served to further increase discrimination between {odor + noise} vs. noise (results not shown).
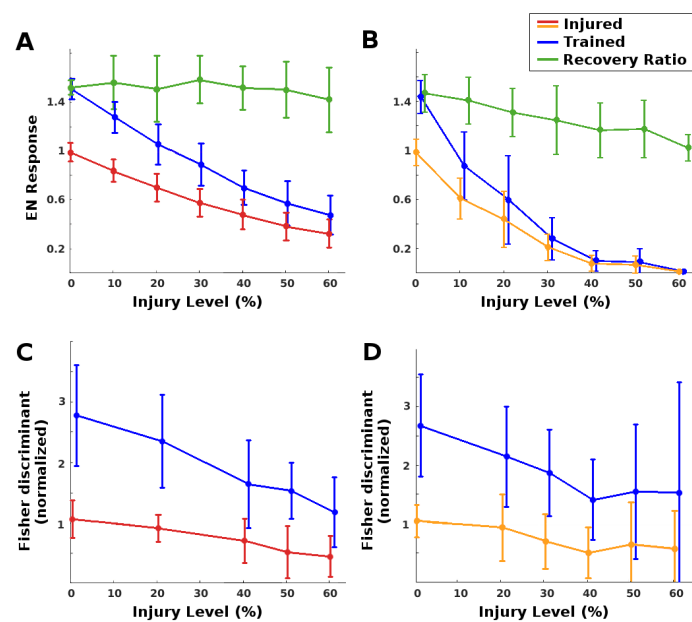
Figure 4. Learning as injury compensation mechanism. Red/orange: Post-injury EN odor response, normalized by naive healthy odor response. Blue: Post-training EN response, normalized by naive healthy odor response. Green: Relative increase from post-injury response due to training. Each datapoint shows the mean and std dev, over n = 31 to 62 (mean = 40) moth instances. A moth's EN response was defined as its mean response to 15 odor exposures. (**A**): Injury to RNs: Trained EN responses (blue) fully regained their pre-injury levels (black line) from injured levels (red) if injury was on average ≤25%. The ability of training to recover lost ground was fairly steady vs. injury level (green). (**B**): Injury to PNs was more traumatic: Post-injury EN response (orange) was lower, and trained responses (blue) fully regained pre-injury levels if injury was on average ≤8%. Also, the ability of training to recover lost ground decreased as injury level increased (green). (**C,D**): Changes in Fisher discriminant between pre-trained odor and control odor, due to injury and subsequent additional training (values are normalized to the initial healthy Fisher discriminant). Injury reduced the ability to discriminate (red/orange curves), while post-injury training fully restored it to above baseline (blue curves). For 0% injury, slight deviations from 1 are due to variations in responses to two groups of odor puffs, pre- and post-injury. (**C**): Injury to RNs. (**D**): Injury to PNs.

2.2. Inhibitory Neurons and Protective Canceling Out Effect

Each glomerulus in the AL has ≈5 excitatory PNs that feed forward to the MB. The moth also has a smaller number of inhibitory neurons (here called QNs) that also feed forward to the MB, analogous to and in parallel with the PNs. We note that these feed-forward QNs are one of three inhibitory networks in the AL-MB. The other two, viz. lateral inhibitory neurons within the AL [27] and global sparsity-inducing inhibition onto the MB from the Lateral Horn [28] (or global self-inhibition by the MB as in *drosophila* [29]), have different functions and are assumed to be non-plastic. Our experiments target the QNs, which innervate only a subset of MB neurons, and are presumed in MothNet to be plastic like PNs.

The goal of this set of experiments was to test whether the existence of QNs parallel to PNs might mitigate the effect of injuries applied to this region (orange stars in Figure 2). We varied the QN:PN ratio (0, 2, 4, 5, and 7 QNs/glomerulus, for 5 PNs/glomerulus) while injuring the AL→MB channel. Each parameter pair (e.g., "4 QNs, 50% FAS injury") had at least 30 moth instances (31 to 40, mean = 35). We found that higher numbers of QNs correlated strongly with reduced effects on EN output magnitudes from upstream injury, but had no clear effect on discriminative ability (results not shown). Results reported here are for effects on EN response magnitude.

Moths with high QN counts had stronger post-injury EN odor responses (Figure 5A), and post-injury training sessions allowed them to fully recover from much higher levels of injury than moths with few or no QNs ($\approx$8% injury for QNs = 0, $\approx$ 15% injury when QNs = 4, and $\approx$30% injury when QNs = 7; Figure 5B, C). High QN counts had another, unexpected advantage regarding the Signal-to-Noise Ratio,

$$SNR = \mu(F)/\sigma(F), \tag{2}$$

where μ and σ correspond to the mean and standard deviation, and $F = \{f_i, i = 1 \ldots n\}$ is the set of discrete EN responses (peak FR) to a series of odor puffs. Naive SNR values (i.e., pre-injury, pre-training) were similar for all QN counts (Figure 5E). While post-injury SNR always dropped proportionally to the severity of the injury, high QN counts substantially reduced these losses to SNR (Figure 5F).
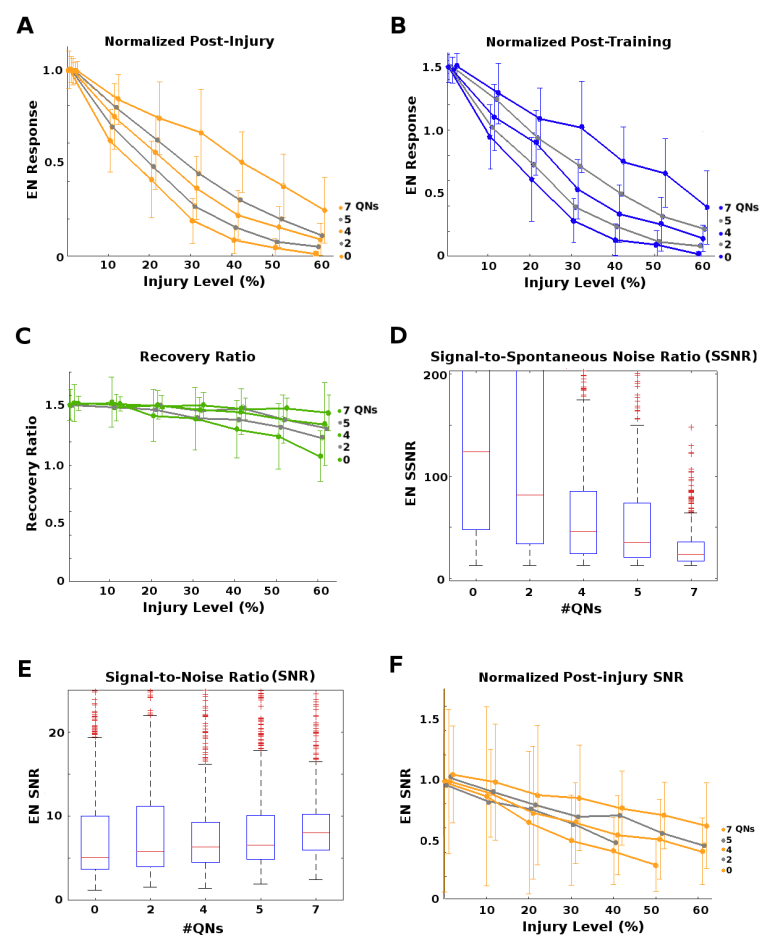
Figure 5. Effects of parallel inhibitory channels. (**A**): Post-injury EN odor responses normalized by naive healthy odor responses, vs. injury level. Each curve corresponds to a number of QNs per 5 PNs, from 0 to 7. Higher QN:PN ratios resulted in much lower impact on EN responses for a given level of injury. (**B**): Post-training EN odor responses normalized by naive healthy odor responses, vs. injury level. Each curve corresponds to a number of QNs per 5 PNs, from 0 to 7. Higher QN:PN ratios resulted in stronger recovery. (**C**): Ratio of post-training to post-injury EN odor responses vs. injury level. Recovery rate dropped off at injury levels $\geq$20% for #QN = 0, but higher numbers of QNs reduced this drop-off, i.e., ensured better recovery. (**D**): Box-whisker plots (showing 25 and 75%iles as a blue box and median as a red line) of the ratio of naive healthy EN odor responses to

spontaneous EN noise (SSNR). This measure of signal clarity was much lower in moths with high QN counts. (**E**): Raw Signal-to-Noise Ratio (SNR) of naive healthy EN responses was fairly uniform across #QNs. (**F**): Post-injury SNR normalized by pre-injury SNR. High QN counts gave strong protection against injury-induced degradation of SNR.

However, high QNs counts also carried a downside. They had a much lower EN Signal-to-Spontaneous Noise Ratio,

$$SSNR = \mu(F)/\mu(Sp), \tag{3}$$

where F is defined as above and Sp is the spontaneous EN firing rate. The SSNR measures the clarity of the signal with respect to background noise, and their values for different QN counts are shown in Figure 5D. Many moth instances from high QN count templates were rejected due to untenably high naive spontaneous noise.

p-values corresponding to Figure 5A,B are given in Tables A7 and A8 in Appendix B. They indicate that the injury-mitigating effect of high QN ratios was meaningful (we avoid the term "significant" in association with p-values, following the arguments in [30]).

Our results demonstrate that the presence of parallel inhibitory neurons in a channel help protect that channel's signal from the effects of injury (at the cost of decreased SSNR). We hypothesize that QNs achieve this by a "canceling out" mechanism: When inhibitory QNs are injured, the overall transmitted signal increases, offsetting the decreases caused by injury to excitatory PNs.

2.3. AL Noise Preserves the Highest EN Responses

The AL is a noisy network. We ask whether this neural noise has injury-mitigation benefits. We suppose that vital odor-related behavior is triggered when a discrete EN response f exceeds some threshold, and that due to AL noise these responses f to a particular odor (at a given concentration) vary as if drawn from a distribution. The moth gets n exposures to a given odor plume, and thus has n discrete responses $F = \{f_i, i = 1 \ldots n\}$. Then to induce the behavior, a triggering response (i.e., $f_i >$ threshold) is needed for only some, not all, f_i.

In this case, it suffices for the system to protect only the strongest (top-scoring) EN responses from injury-induced attenuation in order to maintain its behavioral response. The goal of this experiment was to examine whether higher AL noise levels might preferentially protect the top-scoring EN responses from injury-induced attenuation. Noting that F parametrizes a Gaussian

$$\mathcal{N}(\mu(F), \sigma(F)) = \mathcal{N}(mean(F), std\ dev(F)) \tag{4}$$

we define this top-scoring tranche as those responses at the top end of the distribution: $\{f_i \in F \mid f_i > \mu(F) + \sigma(F)\}$. This corresponds to f_i boosted by fortuitous noise effects and thus most likely to exceed the triggering threshold.

The AL noise level is controlled by a single parameter in MothNet. We adjusted neural noise in the AL to different multiplicative factors of the "natural" AL noise level (i.e., the level calibrated to in vivo data). Factors were 0, 0.33, 0.67, 1.0, and 1.33, where 1.0 is the natural level. Various severities of FAS-like injury were applied to RNs in the Antennae→AL channel (Figure 2A). Over 30 (31 to 62, mean = 40) moth instances were generated from template for each {AL noise, injury level} datapoint. To measure attenuation in top-scoring responses, we defined the Top-End Preservation P as:

$$P(F_j) = \frac{\mu(F_j) + \sigma(F_j)}{\mu(F_h) + \sigma(F_h)}, \tag{5}$$

where F_h is the set of pre-injury (healthy) responses to repeat applications of some stimulus, and F_j is the set of discrete responses to the same stimulus post-injury at level j. $P(F_j)$ measures how much an injury affects the top-scoring responses (represented by $\mu(F) + \sigma(F)$) when it shrinks the entire response distribution from $\mathcal{N}(\mu(F_h), \sigma(F_h))$ to $\mathcal{N}(\mu(F_j), \sigma(F_j))$.

$P(F, j)$ ranges between 1 and 0, where 1 implies no injury-induced attenuation, and 0 implies total attenuation.

Higher AL noise increased the top-end preservation P of EN responses caused by a given level of injury (Figure 6A). It also increased the post-training recovery possible: For example, full recovery occurred for injury $\leq 28\%$ when AL noise was greater than natural level, vs. $\leq 20\%$ when AL noise = 0 (Figure 6B).
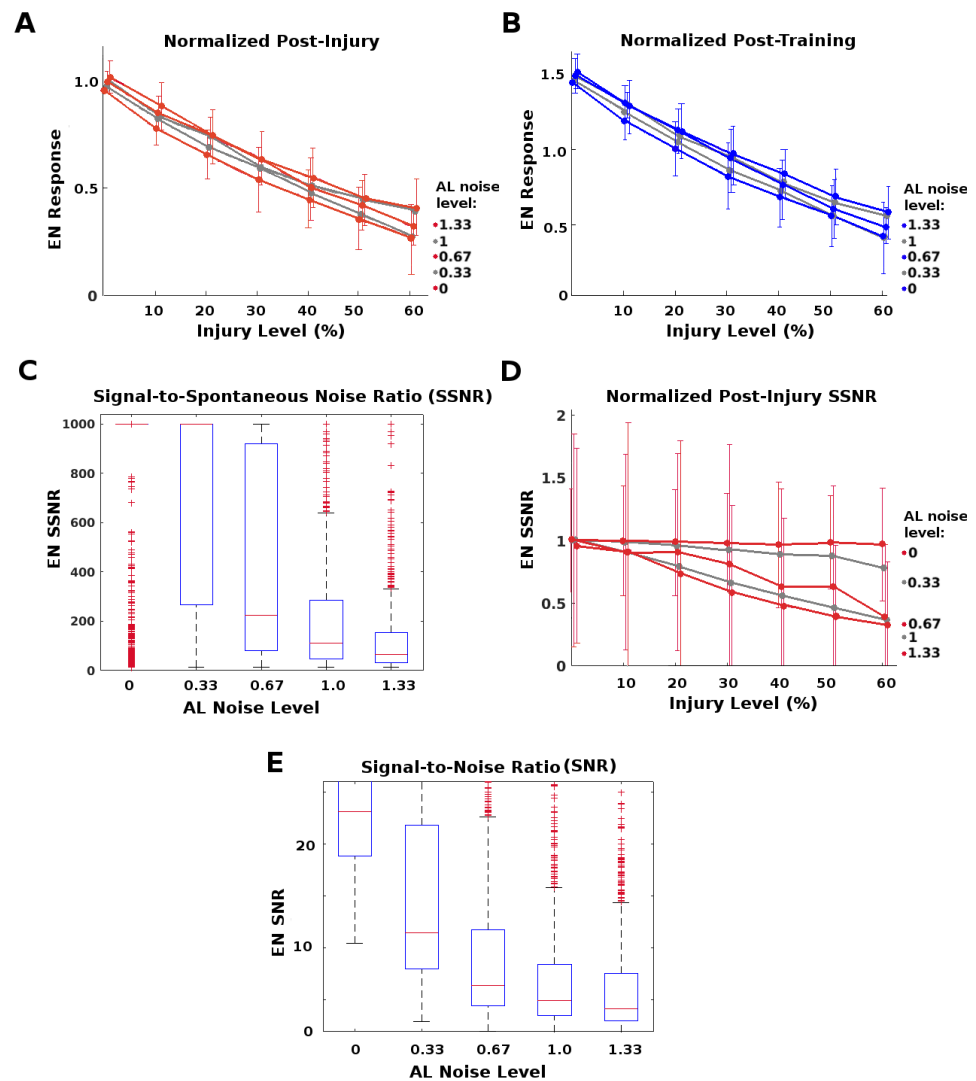
Figure 6. Protective effects of AL noise on the strongest (top-tranche) EN responses: Given RN injury, AL noise protected the strongest EN responses from loss (**A**,**B**), but exacted a cost in terms of signal-to-spontaneous noise ratio (SSNR) and signal-to-noise ratio (SNR) (**C–E**). (**A**): Post-injury EN responses of the top 15% tranche (i.e., the strongest odor responses), normalized by their pre-injury response, vs. injury level. Higher AL noise reduced attenuation from injury, at any level of injury. Each curve corresponds to a level of AL noise, from 0 to 1.33 where 1 = "natural" level. Pre-injury response = black line. (**B**): Post-injury EN responses of the top 15% tranche, normalized by their pre-injury response, vs. injury level. Higher AL noise allowed training to give full recovery of these top EN responses from larger injuries, $\approx 28\%$ injury given maximum noise vs. $\approx 20\%$ injury given no AL noise. (**C**): Naive healthy ratio of EN SSNR was much lower at high AL noise levels. (**D**): Post-injury SSNR, normalized by pre-injury ratios, vs. injury level. Injury lowered SSNR far more in moths with high AL noise. (**E**): Naive healthy SNR was much lower in moths with high AL noise.

However, high AL noise levels had a significant downside, namely, lower SNR (signal to noise ratio) and SSNR (signal to spontaneous noise ratio) values, as seen in Figure 6C–E. This suggests that the moth must trade-off between robustness to injury and signal quality.

In addition, we found that this protection did not apply to all EN responses: Top-scoring EN responses received more injury-mitigation benefit from higher AL noise levels than did average EN responses, i.e., $P(F) > \frac{\mu(F_j)}{\mu(F_h)}$. That is, the extra robustness to injury conferred by higher noise levels was greater for top-scoring responses than for average responses. This meshes with the notion that the system needs not protect all responses, just the ones most likely to exceed triggering threshold. Figure 7 shows this difference in protective effect, top-scoring *vs* average.
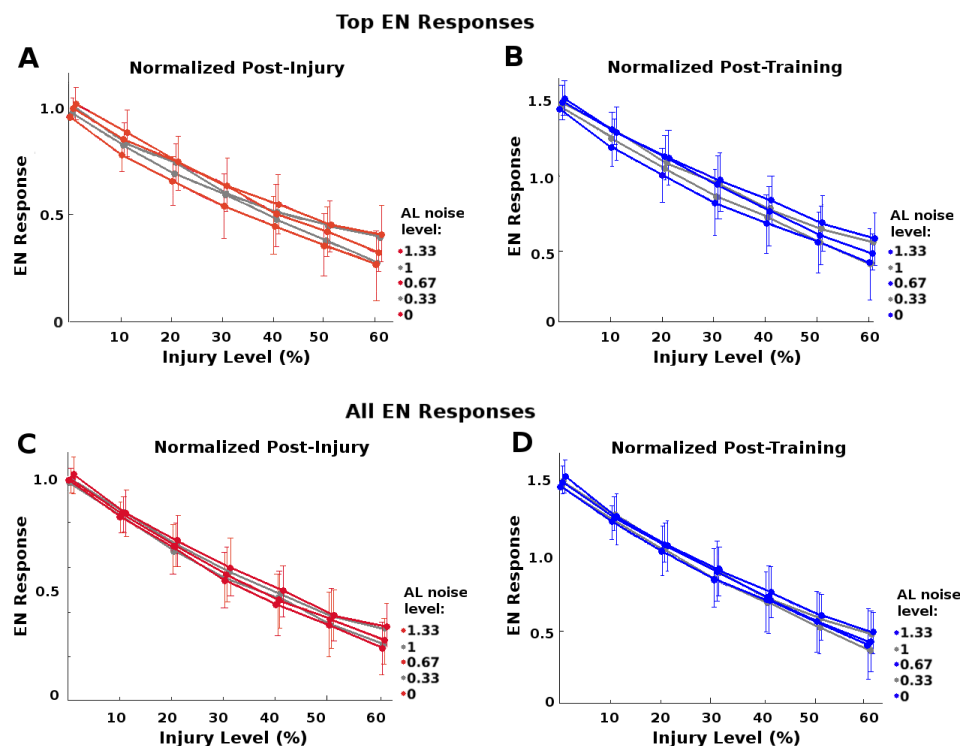

Figure 7. Protective effects of AL noise on strongest vs. average EN responses: Given RN injury, increased AL noise had a greater protective effect on the top 15% tranche of EN odor responses than on more average odor responses, both post-injury and post-training. Each curve corresponds to a noise level. A wider spread of curves indicates greater injury mitigation from higher noise. (**A**,**B**): Top 15% of EN responses, normalized to their pre-injury responses, post-injury (red, grey) and post-training (blue, grey). (These are the same subplots as in Figure 6A,B.) (**C**,**D**): All responses, normalized by their pre-injury responses, post-injury (red, grey) and post-training (blue, grey). Average EN responses had less injury mitigation benefit from high noise than top-scoring EN responses, seen in the narrower spreads.

P-values are given in Appendix B as follows: *(i)* *P*-values for Figure 6A,B (and Figure 7A,B) are given in Tables A1 and A3; *(ii)* *P*-values for Figure 7C,D are given in Tables A2 and A4; *(iii)* *P*-values comparing injury-mitigating effects on top-scoring *vs* average responses, (i.e., Figure 7, A vs. C and B vs. D) are given in Tables A5 and A6. The *P*-values indicate that *(i)* increased AL noise correlated meaningfully with increased protective effect on the top-scoring responses; *(ii)* the protective effect was noticeably lower for average responses; and *(iii)* the protective effect was meaningfully greater for top-scoring than for average responses.

Results given for this set of experiments are for injury's effects on magnitude of EN response. Injury's effects on odor discrimination were only slightly affected by varying AL noise, since both trained odor and control were attenuated by injury (results not shown).

2.4. Ablation Is a Poor Proxy to Biological Injury

Neuronal pathologies are often modeled in a binary way, i.e., by treating a neuron and/or its connections as either fully functional or fully impaired, and ablation injuries are widely studied in theoretical and experimental settings. However, recent FAS studies show that most injured neurons maintain some residual firing rate activity. On large, homogeneous populations of neurons where outputs are pooled, such as the 30,000 RNs in the moth AL-MB, one can expect an approximate equivalence in ablation and FAS, modulo a conversion factor. This is because injuring or ablating any single RN has relatively small effect on the pooled total, and the overall effects of injury to the population can be approximated by average injury ratios. In this case, we estimate that ablation alone is roughly $1.75\times$ more harmful than FAS-like injury, i.e., ablating $n\%$ of neurons in a population causes the same relative drop in total summed FRs as FAS injury to $\approx 1.75n\%$ of the neuron population. For calculations, see section 5.3.3.

However, where neuron numbers are smaller and neural outputs are not pooled, so that individual neurons have relatively unique effects on the system, it is not clear that ablation effects can be reliably mapped to effects of more biologically-plausible FAS-like injuries. Injury to the AL→MB channel (i.e., PNs and QNs) resembles this situation, since there are only 5 PNs per AL glomerulus.

In ablation studies, injury levels are typically measured as percentage of neurons ablated. To assess whether ablation is a good proxy for naturalistic FAS injury, we ran experiments to test whether the impacts on EN response magnitude of ablation injury vs. FAS-like injury had a consistent $1.75\times$ relationship at these two locations, i.e., the RN channel (red stars in Figure 2) and the PN channel (orange stars in Figure 2). All parameters were generated from a MothNet template with AL noise at natural levels and number of QNs = 0 (QNs = 2 gave similar results). Half the moths were injured by ablation and half were injured by FAS, with injury levels from 0 to 60%, in order to compare the relative empirical effects on EN outputs. In each experiment, over 30 moth instances (31 to 40, mean = 35) were generated for each injury {type, level, location} datapoint.

The qualitative effects of ablation and FAS were similar, at each injury site. However, the relative quantitative effects (on EN outputs) of the two injury types varied greatly depending on the site of injury. For RN channel damage, ablation effects were roughly in line with that predicted by theory for large homogeneous populations, i.e., $1.75\times$ FAS damage. The match makes sense given the assumptions on number and distribution of RNs stated above, and is seen in Figure 8A.

In contrast, ablation injury to the PN channel was much less harmful relative to FAS-like injury than predicted by theory. For example, 10% ablation would theoretically induce the same EN loss as 17.5% FAS injury and 20% ablation would correspond to 35% FAS injury. However, in our experiments 10% ablation corresponded to only ~12% FAS injury (a ratio of 1.2), and 20% ablation corresponded to only ~25% (a ratio of 1.25). This effect is seen in Figure 8B by following horizontal lines, which correspond to equivalent EN loss, and comparing the percentage injury levels of (from right to left) experiment FAS, experiment ablation, and theoretical ablation. The experimental ablation levels required to induce a fixed EN loss were much closer to the FAS levels than theory predicted (ratio ~1.25 instead of 1.75).

We remark that this measured discrepancy between theoretical and actual effects is not at the site of injury, but at the downstream ENs, i.e., after the impact of the upstream injury has been nonlinearly modulated by moving through the cascaded system.

This variability in the ratio of ablation injury to equivalent FAS injury, dependent on which neurons are injured, suggests that ablation may be an unreliable proxy for naturalistic neuronal damage in some contexts.
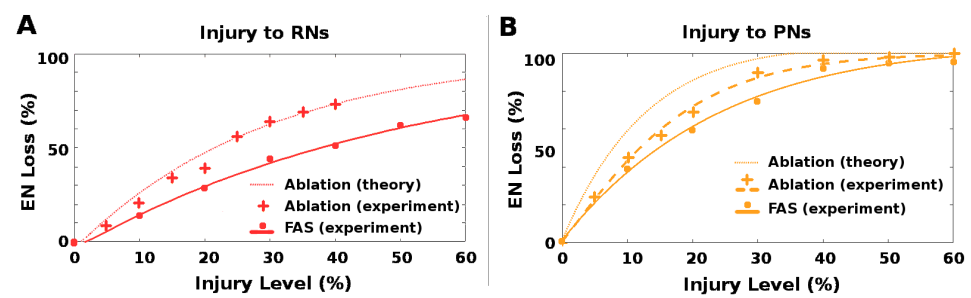
Figure 8. Level of Ablation and FAS injury needed to induce a given EN loss: The relationship of ablation and FAS injury effects varied with injury location, implying that ablation is an unreliable proxy for FAS injury. Key (both plots): dots and crosses are experimental results for FAS (dots) and ablation (crosses). Solid and dashed curves are fits to experimental results for FAS (solid) and ablation (dashed). The dotted curve is the theoretical result for ablation, given the FAS experimental results and a 1.75× ratio (see Section 5.3.3). (**A**) When RNs were injured, ablation induced loss to EN response consistent with theory (1 unit Ablation ≈1.75 units FAS injury). Curves do not meet at origin because they are fits to data points. (**B**) When PNs were injured, ablation induced a much smaller loss than theory would suggest (1 unit Ablation ≈1.25 units FAS injury). This can be seen by looking at horizontal lines (i.e., fixed EN loss), and comparing injury levels that induce this loss for (from right to left) experimental FAS injury, experimental ablation, and theoretically-expected ablation. PN injuries showed a large gap between experimental and theoretical ablation levels required to induce a given EN loss.

3. Discussion

Our simulations indicate that the neural mechanisms and motifs we tested have clear injury-mitigation properties. In this section, we suggest mechanisms by which these structures might protect readout neuronal activity from upstream injury. We note that from a functional point of view, overall resilience of a cascaded system is determined by whether downstream units can still transmit key readout signals to the rest of the body despite upstream damage. Cascaded networks are ubiquitous among biological neural systems, so the principles discussed in the moth's olfactory network may be applicable to other settings. We also discuss the discrepancy between axonal swelling injuries vs. ablation injuries. Finally, we argue that robustness to injury is a key principle of biological neural design.

3.1. Hebbian Plasticity-Induced Recovery from Injury

Learning in the moth olfactory network occurs via a combination of octopamine stimulation and Hebbian growth. Octopamine stimulation temporarily boosts neural firing rates during reinforcement by sugar reward, while Hebbian updates strengthen the synaptic weight w_{ab}, between two neurons a and b, proportionally to the product of their firing rates:

$$\Delta w_{ab} \propto f_a(t) f_b(t). \tag{6}$$

Injuries to the upstream regions of a network result in spike deletions and weaker encodings arriving at downstream neurons. If the damaged region cannot activate downstream neurons with the existing synaptic connection strengths, there is a functional loss of information. However, the combination of octopamine stimulation and a Hebbian update mechanism can evidently mitigate or reverse this effect. The original injured neurons are not themselves repaired, as plasticity only boosts downstream synaptic connections.

The Hebbian mechanism alone is not sufficient to repair damage, because updates to synaptic strengths require sufficiently strong FRs in both the incoming and the receiving neurons. If injury reduces the FRs of upstream neurons, and these reduced inputs lead to lower FRs in downstream neurons, Hebbian updates are significantly degraded. Thus, octopamine-induced stimulation of the injured upstream neurons is crucial to post-injury

plasticity, because it temporarily boosts the FRs of injured upstream neurons back to levels that enable the Hebbian mechanism to strengthen the relevant connections.

We propose that degraded firing rates in downstream neurons are restored via the following mechanism (see schematic in Figure 9A):

1. Octopamine temporarily increases the firing rates of injured upstream neurons.
2. The transient boosted encodings are sufficient to trigger firing in the downstream neurons with the existing synaptic connection strengths.
3. Since neurons on both sides of the plastic connections are firing, Hebbian growth strengthens their connections.
4. Firing rates from the injured upstream region return to their reduced rate once octopamine is withdrawn. However, due to the stronger synaptic connections, these encodings are now sufficient to trigger the downstream neurons. This restores the transmission of key information to the rest of the system.

Because the {octopamine stimulation + Hebbian updates} learning mechanism is automated, i.e., hard-wired as a reward mechanism for adaptive stimuli such as sugar, it acts as a passive injury mitigation system (absent injury, it serves to boost network responses to adaptive stimuli). Since learning is activated repeatedly throughout life (by any rewarding stimulus), it can be expected to act post-injury as an automatic repair mechanism. Alternately, learning can be viewed as a built-in tuning mechanism that in event of injury serves to restore network responses towards their pre-injury states.

We hesitate to call this learning mechanism homeostatic, even though in the context of injury it automatically moves the system towards a prior state, because the restoration is one-way. Learning will not revert responses that have been previously strengthened (by learning itself). Rather, it is an automated mechanism for tuning a network towards stronger responses to adaptive stimuli, which in the event of injury has a homeostatic effect.

3.2. Inhibitory Neurons and Protective Canceling Out Effect

The moth olfactory network has both excitatory projection neurons and inhibitory projection neurons that feed-forward from the antennal lobe to the mushroom body. We propose a mechanism to explain how this can protect downstream neurons from the effects of upstream damage, assuming downstream dynamics depend on the summed input from upstream neurons:

$$(\mathbf{w} \cdot \mathbf{u}) = \mathbf{w}^+ \cdot \mathbf{u}^+ - \mathbf{w}^- \cdot \mathbf{u}^-, \text{ where} \tag{7}$$

$$
\begin{aligned}
\mathbf{w}^+ &= \text{connection weights from excitatory neurons} \\
\mathbf{u}^+ &= \text{FRs from upstream excitatory neurons} \\
\mathbf{w}^- &= \text{connection weights from inhibitory neurons} \\
\mathbf{u}^- &= \text{FRs from upstream inhibitory neurons.}
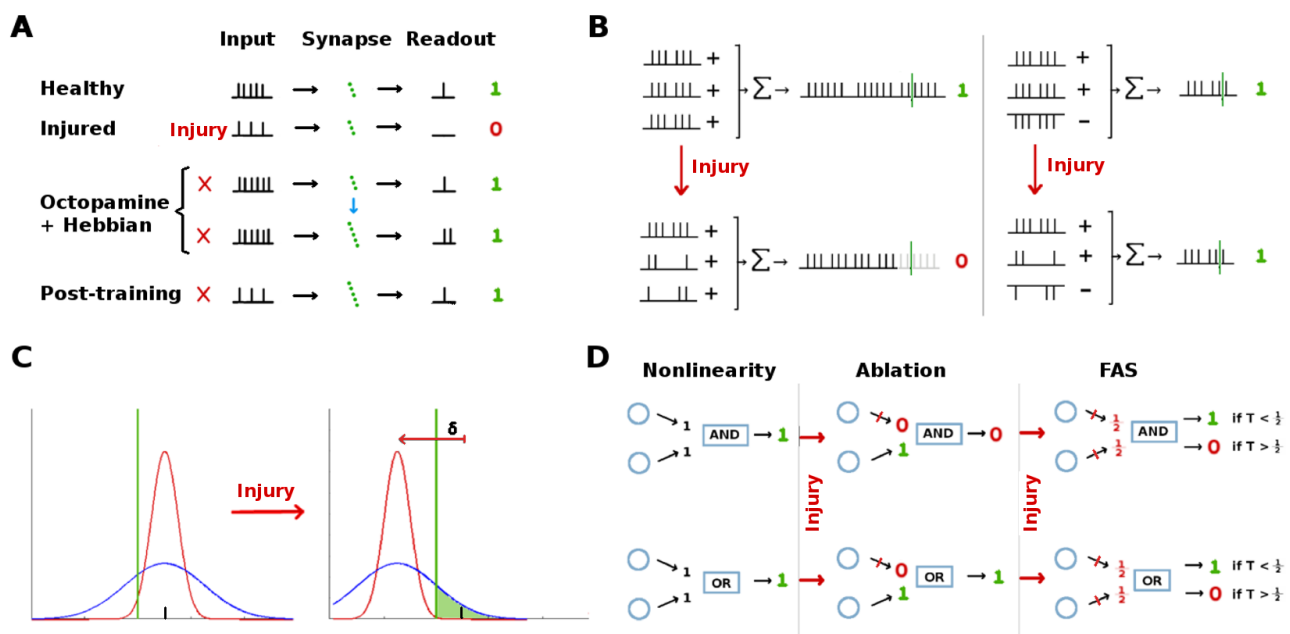\end{aligned}
$$

Figure 9. Injury mitigation hypotheses: In a cascaded network, various architectures can mitigate the effects of injury to upstream neurons by protecting or restoring functionality of downstream units. (**A**): Learning itself can compensate for injury: Octopamine temporarily stimulates the damaged neuron, allowing Hebbian growth to strengthen downstream synaptic connections.Though the injured neuron's signal is not restored, the downstream neurons receive an amplified input, cancelling out the injury. (**B**): Parallel inhibitory channels can reduce the effect of generalized injury because damage is split between excitatory and inhibitory signals, canceling out losses in terms of the inputs to downstream neurons. (**C**): Wide noise envelopes on upstream neuron outputs can protect the strongest stimulus responses from injury-induced attenuation δ, to the degree that their std dev $\sigma > \delta$. This allows the injured neuron's strongest responses to still exceed their activation threshold (green line) for downstream neurons, protecting downstream functionality. (**D**): Two simple examples of non-linearities that can result in qualititative change in the relative effects of ablation and FAS-like injury: In an AND gate, ablation can have a worse effect than FAS downstream, depending on the gate's input threshold T. In an OR gate, ablation can be harmless, while FAS can have a worse effect downstream, depending on T.

When FAS-like injury is applied to the PN+QN pipeline in our neural architecture, the net effect on the summed signal reaching downstream target neurons varies according to the proportion of QNs to PNs ($\mathbf{u}^- : \mathbf{u}^+$, assuming uniform weights $\mathbf{w}$). When all feedforward signals are excitatory (i.e., $\mathbf{u}^- = 0$), injury will always reduce the summed input reaching a downstream neuron. If QNs exist, however, and both PNs and QNs share the same exposure to injury, then the overall reduction to the summed input will be mitigated on average, since any injury to QNs will increase the summed input, offsetting decreases due to PN injury. A schematic of this "cancelling out" mechanism is shown in Figure 9B.

The injury resistance provided by high QN counts comes at a cost to other functionalities, e.g., higher spontaneous EN noise relative to odor response. Presumably, biological networks have QN counts which optimally balance the benefits of injury mitigation on one hand versus the need for high signal-noise-ratio, as well as other concerns such as the energy cost to the organism. If the QN counts are low (e.g., QN:PN $\leq$ 20%, as in the moth), this injury mitigation benefit is likely less important relative to other architectural or functional constraints. We note that learning and plasticity are not pre-requisites for this mechanism.

3.3. Upstream Noise Protects Downstream Behavior

Suppose that the behavioral response is preserved after injury if at least a subset of stimuli elicit downstream responses that exceed action-triggering thresholds. In this case, a large noise envelope on upstream neurons may help protect the network's functionality.

Assume the firing rate of an upstream neuron FR responds to stimuli following a Gaussian distribution $N(\mu, \sigma)$, and that it needs to exceed a threshold T to activate downstream neurons. If the neural damage reduces this FR in average by δ, a large noise envelope (large σ) will ensure that some post-injury responses still exceed threshold, i.e., that $\mu - \delta + \sigma \geq T$. This idea is sketched in Figure 9C for two FRs characterized by $N(\mu, \sigma_1)$ and $N(\mu, \sigma_2)$ with $\sigma_1 > \sigma_2$.

Our experiments indicate that AL noise does enable the highest EN responses to exceed threshold after injury, even as the average EN response drops. However, the injury mitigating benefit of increased upstream noise comes at a cost to other system functionalities, e.g., it reduces signal-to-noise ratio. Noise levels in biological networks (such as in the antennal lobe) may represent an evolved/optimized trade-off between injury mitigation effects and negative side-effects such as reduced SNR. We note that the sparsity of the MB acts as a powerful noise filter [7]. Plasticity is not a pre-requisite to this mechanism.

3.4. Ablation is a Poor Proxy to Biological Injury

Neuronal injuries are often modeled in a binary way, i.e., by treating a neuron and/or its connections as either fully functional or fully impaired. Our results indicate, however, that in some situations ablations are a poor proxy for more naturalistic FAS-types of injuries regarding effects measured downstream from the injury site.

When the neuron population to be injured is large, and has pooled outputs to the next layers (in our model, the antennae/RNs), ablation maps to FAS injury in a predictable manner due to averaging effects over the population (see Section 5.3.3). However, when the neuron population is small (in our model, the PNs) the effects of ablation vs. FAS are not predictable. Ablation of PNs had much lower impact than large-population theory would predict

Our key finding is that ablation effects are inconsistent relative to FAS-like effects, depending on the location and characteristics of the injured neurons. This calls into question the value of ablation as a proxy for naturalistic neural injuries. We suggest that in systems with large numbers of somewhat interchangeable units (e.g., the 30,000 RNs) ablation is a suitable way to model injury; while in systems with smaller numbers of specialized neural units (e.g., the AL with 60 glomeruli) ablation is a poor injury model. This unsuitability is sharpened when the actionable effects of injury are measured downstream from the regions injured, because there is a complex interplay between the injuries and network nonlinearities, making the outcome somewhat unpredictable. Simple examples of possible effects of non-linearities (AND and OR gates) that might cause ablation vs. FAS-like injury models to diverge are shown in Figure 9D.

3.5. Limitations

Our computational model assumed only one readout neuron and one broadly-activating odor (two for the the discrimination experiments). A more detailed assessment of injury and mitigation might involve several readout neurons to allow for disparate effects on various readouts, as well as several time-varying odor mixtures, more narrowly-activating odors, and concentration gradients. In addition, we model learning-induced Hebbian plasticity in existing neurons, not other neurogenesis mechanisms. We model the effects of octopamine (since it is closely involved in learning), not other neuromodulators such as serotonin. We do not know if the PN+QN channel is a realistic target for injury: We chose it in order to investigate deficits caused by injury to innermost hubs. Our study certainly did not exhaust all potentially interesting combinations of structures-under-test and injuries. More elaborate models could incorporate extra-neural factors such as moth physiology and physiological status, and external environmental conditions such as temperature and humidity. These conditions are beyond the scope of MothNet and the in vivo experiments that inform it.

4. Conclusions

We investigated the moth olfactory network with the goal of understanding how its basic architectural components serve to make sensory processing robust to injury. Since most organisms are exposed to neuronal damage throughout their lives, it is important to understand how such neuronal circuits are structured to maintain functionality despite impairments. In this work, we showed explicitly how certain structural and functional network motifs act as injury mitigation mechanisms. Specifically, we examined how (i) Hebbian learning, (ii) high levels of noise, and (iii) presence of parallel inhibitory and excitatory connections, can support overall robustness to injury in the olfactory system in the *Manduca sexta* moth.

Our findings indicate that, in addition to accurate sensory processing, biological neural networks such as those found in the moth olfactory system hold robustness to injury as a central design principle. Our findings also suggest an additional hypothesis: Plasticity coupled with neuromodulatory stimulation, now central to learning, may have originally evolved as a repair mechanism for neural systems to offset injury and maintain function, and was only later ported to the task of developing responses to new information (exaptation). If this is the case, then the gift of learning is due originally to the exigencies of brain damage.

Our results also show that these architectures can in fact cause worse performance by some other performance metrics, e.g., SNR. Thus, trying to explain these architectures from the point-of-view of, for example, information theory risks running against the fact that they are actually suboptimal according to that particular lens. A more comprehensive and nuanced framing of the neural signal processing task, positing multiple design goals including injury mitigation, can enable better understanding of neurosensory processing. That is, a neural architecture can be understood only if its injury mitigation function, and the trade-offs between this and other desired functions, are considered. Indeed, it is possible that some neural structures and mechanisms, including the ability to learn, are best understood as evolutionary solutions to the challenge of maintaining function despite injury.

5. Materials and Methods

In this section, we first briefly outline the Moth Olfactory Network. We then describe focal axonal swelling (FAS), a characteristic form of neuronal injury utilized as a model of damage, and how it was applied to the network. Lastly, we provide details about the experimental setups involved in our key findings. A complete Matlab codebase for MothNet and the injury simulations is available at https://github.com/charlesDelahunt/BuiltToLast.

5.1. Brief Overview of the Moth Olfactory Network

For a schematic, see Figure 1. Receptor Neurons (RNs) in the Antennae detect relevant odors in the environment and transmit specific signals to the Antennal Lobe (AL) [31,32], which acts as a pre-amp, providing gain control and sharpening odor codes [33]. The AL neurons project odor codes forward to the Mushroom Body (MB) [34] by means of noisy [35] excitatory Projection Neurons (PNs), and to a smaller number of parallel inhibitory neurons (here called QNs). The Kenyon Cells in the MB fire sparsely (due to global inhibition from the Lateral Horn), and encode odor signatures as memories [36,37]. Finally, Extrinsic Neurons (ENs) are viewed as readout units that interpret the MB codes, delivering actionable output to the rest of the body [12,13]. In response to reward (sugar at the proboscis), a large neuron releases octopamine in the AL and MB. In the AL, this neuromodulator induces stronger responses in AL neurons (stimulation), though intra-AL connections are not plastic. Synaptic connections into and out of the MB (AL→MB and MB→ENs) are plastic [32,38] given octopamine, though octopamine does not stimulate KCs in the MothNet model (biological data is lacking). Learning fails in the absence of octopamine [39,40]. For accurate anatomy, see, e.g., [41]. Appendix A contains more detail,

and full details (biological and computational) can be found in [7] and in the codebase mentioned above.

5.2. Moth Template Parameters

In each experiment, over 30 moth instances (per data point) were randomly generated from the MothNet template defining the architecture, which included biologically-plausible choices for numbers of neurons, synaptic connection maps and weights, odor projection onto the glomeruli of the AL, learning rates, and SDE time and noise constants. To generate connection matrices, non-zero connections were randomly assigned according to architectural constraints and template parameters, then non-zero connection weights were then drawn from gaussian distributions with parameters dependent on types of neuron. For full details of how MothNet instances were generated from templates, please see [7]. The templates were realistic in the senses of having *(i)* PN firing rate behavior matching data from live moths and *(ii)* architecture parameters that match what is known from the literature [7]. Some templates were moved to the boundaries of, or out of, a known realistic regime by varying key parameters-under-test when required by the experiment.

1. The number of QNs per glomerulus varied from 0 to 7, in order to test the injury-mitigating effect of inhibitory QNs in parallel with PNs in the feed-forward AL→MB channel. PNs were fixed at 5 per glomerulus, as in moths. In moths, each QN has dendrites in several glomeruli. In our experiments, each MothNet QN has dendrites in exactly one glomerulus (like PNs), to make QN:PN ratios meaningful. Moths may have QN:PN ratio ≈ 0.2 (i.e., relatively few QNs), insofar as a ratio can be estimated (actual values are not known). We note that more QNs means that more individual MB neurons receive extra inhibition. The global inhibition mechanism modulated by the Lateral Horn is distinct, and is not affected by QN numbers.
2. The level of Gaussian noise affecting all AL neurons varied from 0 to 1.33, where 1.0 represents natural levels (i.e., fitted to in vivo data). This noise range was used for experiments on effects of AL noise levels. Other experiments used AL noise = 1.0.
3. For experiments examining effects of learning and of AL noise, we set the number of QNs per glomerulus to zero. Setting the number of QNs equal to 2 (equivalently, QN:PN ratio = 0.4) gave similar results. These QN values remain close to that of realistic models.

The extremes of the parameter regimes described above deviated significantly from calibrated models, and sometimes created moths with untenably noisy, dysfunctional EN responses to odor. Thus, we discarded moths with naive EN odor response-to-spontaneous FR (SSNR) outside an envelope defined by $\frac{\mu(F)}{\mu(s)} < 12$. These comprised about 12% of moths generated, with the percentage depending on the varied parameters: Templates with high numbers of QNs and/or very high AL noise had more rejected moths; templates with few QNs and normal/low AL noise had few rejected moths. Extra moths were generated as needed to match numbers across all experiments, so that each {parameter-under-test, injury level} pair (e.g., "4 QNs, 50% FAS injury") had ≥ 30 moth instances.

5.3. FAS-Like Injury

Focal Axonal Swellings (FAS) is a neural injury associated with traumatic brain injury (TBI), typically caused by physical shock. Examples in current events include blast injuries from recent wars, as well as impact injuries in contact sports. FAS presents as swollen neural axons (the signal delivery pipelines) with dramatic diameter changes, causing signals from the upstream source to be diminished or lost entirely before reaching downstream target neurons [17]. This degradation can be expressed as reduced FRs from upstream neurons, characterized in a computational model by [19], which found that signals traveling down an injured axon are attenuated to greater or lesser degree according to the amount of swelling and the firing rate of the signal. While ablation is a ready and oft-used means to

model neural injury, it imposes a binary "all-or-nothing" effect which is not present in FAS injuries. In these experiments we model neural injury according to [19].

5.3.1. FAS in the FR Model Context

While FAS models effects at the level of spike trains, it also has a meaningful representation in Firing Rate models such as MothNet. In particular, unlike ablation, FAS causes reduced but still non-zero FRs. In addition, the low-pass filtering effect of FAS, which impacts closely-bunched clusters of spikes more than sparse spikes, in analogous manner impacts high FRs more strongly than low FRs. Thus, FAS, applied in a FR model, results in neuron FRs being reduced but not ablated, with high-FR neurons affected more strongly than low-FR neurons. For a fixed amount of total damage, FAS results in relatively many partially-damaged neurons, while ablation results in relatively few fully-destroyed neurons. The structure of FAS injuries allows a principled conversion to the FR-based model context, as detailed below.

5.3.2. Injury Methods

FAS due to physical trauma does not affect all neurons in a targeted brain region equally, nor does it operate in an "all or nothing" way [42]. We used an injury regime derived in [19], which calculates the fractions of injured neurons falling into each of four injury types: FR unaffected (transmission); FR cut by half (reflection), FR destroyed (ablation), or FR filtered according to

$$f_{injured}(s) = L(f_{healthy}(s)), \tag{8}$$

where $f_* =$ firing rate, s is a stimulus, and L is a lowpass filter. FAS injury fractions were as follows: 15% transmission, 35% reflection, 35% ablation, and 15% low-pass filtering, where the low-pass filtering generally multiplied FRs by a factor of 0.7 to 0.9 depending on initial FR (taken from [17]). Neurons in the target group were randomly selected for injury according to the percentage specified in the particular experiment (0 to 60%). Each injured neuron in the target group was then randomly assigned one of the damage types.

Applying this injury regime to populations of PNs and QNs in the model is straightforward, since these neurons are modeled one-to-one (i.e., one neuron in the model represents one actual neuron). For example, given 300 PNs and a FAS injury level of 50%, 150 PNs would be randomly assigned (in fixed proportions) to have their FRs multiplied by either 1 (transmit), 0.5 (reflect), 0 (ablate), or variable (low-pass filter). At an ablation level of 50%, 150 PNs would be randomly assigned to have their FRs zeroed out.

Injury to RNs was treated differently because in MothNet each RN (inputting to one glomerulus in the AL) stands for ≈ 500 RNs in the moth. Injury of RNs was handled as follows, leveraging averages over large numbers of homogeneous neurons:

1. The FAS injury level m (% of neurons injured) was converted into a theoretically equivalent ablation injury level n (% of neurons ablated), where $m = 1.75n$ (see calculation below).
2. Each RN's FR was multiplied by $(1-n)$ (i.e., attenuation), since the glomerulus had fewer inputs.
3. The RN noise parameter was multiplied by $\frac{1}{\sqrt{1-n}}$, since the glomerulus was averaging fewer inputs, giving less noise reduction from averaging.
4. 15% of FAS injury was low-pass filtering, which depends on the injured neuron's firing rate. This part of the FAS injury to RNs varied according as the RN firing rate varied. So the actual $n(t)$ affecting a given RN fluctuated slightly with odor inputs.

5.3.3. Theoretical Correspondence of FAS to Ablation

We wish to compare levels of ablation and FAS injury, using as a measure the percentage of neurons injured. Since FAS does not usually destroy a neuron's FR while ablation

does, we expect that ablating n% of neurons will cause the same average drop in summed population FR as FAS injury to m% of neurons where $m > n$, as follows:

We assume a large homogeneous population of neurons. Given a fixed percentage of FAS injury, we can estimate the ablation level that gives the same average total reduction to summed FR:

$$FAS \text{ applied to } 100 \text{ units}$$
$$\implies 15 \text{ transmit} + 35 \text{ reflect} + 35 \text{ ablate} + 15 \text{ lowPass}$$
$$\implies (15 \times 1) + (35 \times 0.5) + (35 \times 0) + (15 \times 0.7) \approx 43$$
$$\implies \text{ablation applied to } 57 \text{ units}$$
$$\implies \text{conversion rate } = 100/57 \approx 1.75.$$

e.g., 20% ablation nominally corresponds to 35% FAS injury, in terms of total reduction in summed FRs over the injured population.

5.3.4. Location of Injured Regions

Two sites were targeted independently for FAS injury (RNs or PNs+QNs), with the choice determined by the experiment. See Figure 2.

- We posited the Antennae → AL channel, i.e., the RNs, as a likely site for FAS injury due to their exposure to external impacts. The red stars in Figure 2 show this injury location. There are about 30k RNs, which we divide as 500 RNs responding to each of 60 atomic odors. These atomic-focused groups send their inputs to a single glomerulus in the AL, where their inputs are averaged to reduce noise. The receptors for a given glomerulus are distributed across the antennae. Thus, we expect injury to an antenna to affect a roughly equal percentage of each glomerulus' 500 inputs.
- *PNs+QNs:* A channel carries PN (and QN) axons from the AL to the MB. We modeled damage to this channel by injuring n% of all neurons, with individual PNs and QNs given the same likelihood of injury. The orange stars in Figure 2 show this injury location. We assumed that the degree of global sparsity enforced on the MB by the Lateral Horn remained stable over modulations of input signal from the AL.

5.4. Simulation Protocols

Each experiment consisted of many moth instances, all generated from the same template with only the parameters-under-test varied. Over 30 moth instances (trials) were run for each parameter pair (e.g., "4 QNs, 50% FAS injury") with training on 5 odor puffs. Odors were randomly generated and projected onto the AL broadly ($\approx$25 glomeruli targeted). A single trial consisted of injuring and then training a single moth instance, in five or six stages:

0. (for discrimination experiments only) Pre-training to ensure initial discrimination: Trained odor plus octopamine were applied (3 odor puffs) with Hebbian plasticity activated.
1. Pre-injury baseline: Odor was applied without octopamine (15 odor puffs, each 0.2 mSec) to assess naive EN odor response.
2. Injury was applied.
3. Post-injury odor: Odor was applied without octopamine (15 odor puffs) to assess the effects of injury on EN odor response.
4. Training: Odor plus octopamine were applied (5 to 15 odor puffs), with Hebbian plasticity activated, to train the system (plasticity was coincident with octopamine, details in Appendix A).
5. Post-Training odor: Odor was applied without octopamine (15 odor puffs) to assess post-training EN odor response.

Firing rate $f(t)$ from a single EN was recorded, to track the actionable effect of injury and training on the system's outputs. A timecourse of EN firing rates, from a typical experiment, is shown in Figure 3.

5.5. Additional Simulation Details

- *Learning compensates for injury:* Moth templates were biologically plausible, in the sense that their AL behavior matched data, and they demonstrated learning behavior. Moths in this experiment had 0 QNs, i.e., no feed-forward inhibitory signals from AL→MB (2 QNs per glomerulus gave very similar results). In separate experiments, either the RN channel (Antennae→AL) or the PN+QN channel (AL→MB) was injured with FAS (0% to 60%).
- *Parallel inhibitory neurons protect EN responses:* Moths were generated from biologically plausible template but with {0, 2, 4, 5, or 7} QNs per 5 PNs. The corresponding QN:PN ratios are {0, 0.4, 0.8, 1.0, or 1.4}. AL noise was set to a natural level (matching data). The PN+QN channel was injured with FAS (0% to 60%).
- *AL noise preserves the highest EN responses:* Moths templates had AL noise level from 0 to 1.33, where 1.0 corresponds to natural AL noise. Moth templates in this experiment had only excitatory PNs (i.e., #QNs = 0). Moth templates with 2 QNs per glomerulus gave similar results. The RN channel was injured with FAS from 0% to 60%.
- *Comparison of ablation vs FAS injuries:* To compare effects of ablation *vs* FAS injury, we ran parallel experiments using ablation and FAS. Injury levels were set between 0% and 60%, using either FAS type or ablation.

Author Contributions: Conceptualization, C.B.D., P.D.M., J.N.K.; methodology, P.D.M., C.D.; software, C.B.D.; validation, C.B.D.; formal analysis, C.B.D.; investigation, P.D.M., C.B.D.; writing—original draft preparation,C.B.D.; writing—review and editing, C.B.D., P.D.M., J.N.K.; visualization, C.B.D.; supervision, P.D.M., J.N.K.; project administration, J.N.K.; funding acquisition, J.N.K. All authors have read and agreed to the published version of the manuscript.

Conflicts of Interest: The authors declare no conflict of interest.

Appendix A. MothNet Architecture

We use close variants of the model of the moth olfactory network (MothNet) developed in [7], modifying the architecture features as needed for each experiment. We provide only selected details about the architecture here. For a fuller description and references, please see [7] and its associated codebase, as well as the codebase given for this paper.

MothNet uses firing rate dynamics for neural firing rates [43], evolved as stochastic differential equations [44], and Hebbian plasticity for synaptic weight updates.

For the experiments in this work, the relevant architectural structures of the MON were:

- *Antennae:* Network structures in which chemical receptors detect odor and send signals to the Antennal Lobe via receptor neurons (RNs). There are approximately 30k RNs (assumed here to be $\approx$ 500 per glomerulus) that MothNet combines into one averaged RN per glomerulus. All injury protocols applied to RNs accounted for this many-to-one abstraction (for details, see Section 5.3.2). While true RNs can respond to several atomic odors, MothNet makes the simplifying assumption that each RN responds to exactly one, and that all RNs responsive to a certain target innervate the same glomerulus in the AL. Thus atomic odors and AL glomeruli are assumed to correspond 1-to-1.
- *Antennal Lobe (AL):* The AL acts as a pre-amp, converting the weak electrical signals into an output signal strong enough to tolerate noise and allow further processing. It also modulates the odor's encoding via intra-AL lateral inhibition. The AL structure contains approximately 60 neural units (glomeruli) which process odors and send excitatory signals (via projection neurons PNs, $\approx$ 5 per glomerulus) and inhibitory

signals (via QNs) downstream to the mushroom body. In the version of MothNet used here, QNs have dendrites in one glomerulus, rather than in several (as in the actual MON), and were parametrized as inhibitory analogs to PNs, with the same connection distributions and FR behavior as PNs. This enabled us to vary the ratio of QNs to PNs from 0 to 1.4 for experiments.

- *Mushroom Body (MB):* The MB is a high-dimensional ($\approx$ 4000 neurons), sparsely-firing structure that encodes odor signatures and memories, and contains plastic synaptic connections. Odor responses feed-forward from the AL to the MB via PNs (excitatory) and (in moths) QNs (inhibitory). MB neurons feed-forward to the Readout Neurons. Sparsity in the insect MB is enforced by global inhibition either from the Lateral Horn (moth) or from the MB itself ([28], *drosophila* [29], locust [45,46]). In MothNet, MB sparsity is enforced by injecting a time-varying inhibitive input term into all MB neurons, such that only the most strongly-excited neurons fire and the percentage of MB neurons with positive FRs is fixed (at, e.g., 5 to 15%). This inhibitive term models input from the Lateral Horn, which is thus implicitly, not explicitly, modeled by global inhibition of the MB.

- *Readout Neurons (ENs):* Odor codes in the Mushroom Body (MB) feed-forward to Readout Neurons (Extrinsic Neurons, ENs), which are assumed to act as decision-making neurons. Strong EN responses trigger actionable messages (such as "fly upwind"). When assessing effects of injury, we focus on the ENs, since these represent the final, actionable output of the system. MothNet posits one EN, whose output FR serves to measure the functional effects of upstream injury.

- *Octopamine effects:* Octopamine is crucial to learning in the MON. In MothNet, effects of octopamine instantiate in three ways: (1) It stimulates AL neurons (RNs, LNs, PNs), making them more responsive to excitatory inputs and less responsive to inhibitory inputs; (2) It activates plasticity in AL→MB and MB→EN synaptic connections (in MothNet this simply means that plasticity in these regions is "turned on" when octopamine is non-zero; (3) It reduces the sparsity parameter of the KCs in the MB, to model that it makes KCs less responsive to Lateral Horn inhibition. Our experiments indicated that items *1* and *2* are necessary for learning (as in [39,40]), while item *3* is optional [7].

Appendix B. *p*-Values

We give *p*-values for Figures 5–7, as evidence of meaningful differences between injury mitigation associated with amount of AL noise, or number of parallel inhibitory connections (QN:PN ratio). We avoid the term "significant" in association with *p*-values, following the arguments in [30]. *p*-values here are calculated as the likelihood of the mean of one distribution occurring given another (implicitly null) distribution. *p*-values for post-injury responses had much higher *N*, i.e., number of moths per datapoint, because moth instances with any number of post-injury training puffs could be combined.

Appendix B.1. Effects of AL Noise

These tables refer to Figures 6A,B and Figure 7A–D. *p*-values are given in four tables for *top-scoring* responses and for *average* responses, both post-*injury* and post-*training*. Two more tables then give *p*-values for top-scoring vs. average responses, post-injury and post-training. "Post-injury" refers to the state after the system has been injured but before training. "Post-training" refers to the state after the system has been injured then trained. They are calculated from the same set of simulations as the figures.

The key points are that (i) the increased protective effect due to increased AL noise was meaningful on the top-scoring responses; (ii) the protective effect was noticeably lower for average responses; and (iii) the protective effect was meaningfully greater for top-scoring than for average responses, both post-injury and post-training.

Table A1. *p*-values for normalized post-injury effects of increased AL noise, on *top-scoring* responses (see Figure 6A). *p*-values are for a given AL noise vs. AL noise = 0.67 (i.e., the middle value). "0+" indicates *p*-value < 0.0005. *N* = 98 to 112 (mean = 107).

Injury Level	AL Noise Level				
(%)	0	0.33	0.67	1.0	1.33
10	0+	0.015	0.5	0.15	0.006
20	0+	0+	0.5	0.161	0.009
30	0+	0.001	0.5	0.027	0+
40	0+	0.363	0.5	0.001	0+
50	0+	0	0.5	0.036	0.013
60	0+	0.002	0.5	0+	0+

Table A2. *p*-values for normalized post-injury effects of increased AL noise on *average* responses (see Figure 7C). *p*-values are much larger than those corresponding to top-scoring responses (cf. Table A1), indicating weaker protective effects. *p*-values are for a given AL noise vs. AL noise = 0.67 (i.e., the middle value). "0+" indicates *p*-value < 0.0005. *N* = 98 to 112 (mean = 107).

Injury Level	AL Noise Level				
(%)	0	0.33	0.67	1.0	1.33
10	0.032	0.057	0.5	0.355	0.267
20	0.026	0.001	0.5	0.495	0.061
30	0.028	0.022	0.5	0.225	0.008
40	0.056	0.703	0.5	0.023	0
50	0.033	0.018	0.5	0.132	0.087
60	0.003	0.045	0.5	0+	0+

Table A3. *p*-values for normalized post-training effects of increased AL noise on *top-scoring* responses (see Figure 6B). *p*-values are for a given AL noise vs. AL noise = 0.67 (i.e., the middle value). "0+" indicates *p*-value < 0.0005. *N* = 31 to 40 (mean = 35).

Injury Level	AL Noise Level				
(%)	0	0.33	0.67	1.0	1.33
10	0+	0.003	0.5	0.545	0.789
20	0+	0.001	0.5	0.866	0.525
30	0.001	0.004	0.5	0.29	0.183
40	0.008	0.116	0.5	0.404	0.006
50	0.1	0.086	0.5	0.072	0.007
60	0.034	0.011	0.5	0.011	0.001

Table A4. *p*-values for normalized post-training effects of increased AL noise on *average* responses (see Figure 7D). *p*-values are markedly larger than for those corresponding to top-scoring responses (cf. Table A3), indicating weaker protective effects. *p*-values are for a given AL noise vs. AL noise = 0.67 (i.e., the middle value). $N = 31$ to 40 (mean = 35).

Injury Level	AL Noise Level				
(%)	**0**	**0.33**	**0.67**	**1.0**	**1.33**
10	0.001	0.035	0.5	0.833	0.913
20	0.031	0.032	0.5	0.89	0.715
30	0.049	0.025	0.5	0.755	0.534
40	0.33	0.271	0.5	0.67	0.066
50	0.548	0.16	0.5	0.278	0.078
60	0.344	0.058	0.5	0.021	0.005

Table A5. *p*-values for comparison of AL noise's effects on *top-scoring* responses vs. *average* responses, post-injury (see Figure 7A vs. C). "0+" indicates *p*-value < 0.0005. $N = 98$ to 112 (mean = 106).

Injury Level	AL Noise Level				
(%)	**0**	**0.33**	**0.67**	**1.0**	**1.33**
10	0.995	0.667	0.224	0.064	0.157
20	0.903	0.447	0.078	0.182	0.06
30	0.787	0.395	0.211	0.015	0.022
40	0.747	0.243	0.12	0.004	0.009
50	0.676	0.242	0.154	0.024	0.002
60	0.613	0.256	0.079	0.01	0+

Table A6. *p*-values for comparison of AL noise's effects on *top-scoring* responses vs. *average* responses, post-training (see Figure 7B vs. D). "0+" indicates *p*-value < 0.0005. $N = 31$ to 40 (mean = 35).

Injury Level	AL Noise Level				
(%)	**0**	**0.33**	**0.67**	**1.0**	**1.33**
10	1	0.777	0.102	0.004	0.044
20	0.988	0.406	0.009	0.059	0.004
30	0.92	0.323	0.084	0+	0+
40	0.874	0.111	0.02	0+	0+
50	0.799	0.114	0.038	0+	0+
60	0.699	0.123	0.008	0+	0+

Appendix B.2. Effects of QN Ratio

Tables A7 and A8 give *p*-values for each {injury level, #QNs}, with respect to the next lower (adjacent) #QNs (same injury level). These tables refer to Figure 5A,B. "Post-injury" refers to the state after the system has been injured but before training. "Post-training" refers to the state after the system has been injured then trained. They are calculated from a different set of simulations, which however had identical parameters and gave the same overall results. In general the differences due to #QNs are meaningful, indicating that increasing the number of parallel inhibitory neurons improved injury mitigation properties.

Table A7. *p*-values for normalized post-injury responses (see Figure 5A). *p*-values are with respect to *adjacent* (next lower) QN numbers. "0+" indicates *p*-value < 0.0005. $N = 66$ to 97 (mean = 79).

Injury Level	Number of QNs			
(%)	2	4	5	7
0	0.556	0.306	0.813	0.955
10	0.007	0.01	0.286	0.06
20	0+	0.003	0.027	0+
30	0+	0+	0.105	0+
40	0+	0+	0.025	0+
50	0.015	0+	0.091	0+
60	0+	0+	0+	0+

Table A8. *p*-values for normalized post-training responses (see Figure 5B). *P*-values are with respect to *adjacent* (next-lower) QN numbers. "0+" indicates *p*-value < 0.0005. $N = 31$ to 62 (mean = 40).

Injury Level	Number of QNs			
(%)	2	4	5	7
0	0.316	0.365	0.867	0.162
10	0+	0.02	0.141	0.093
20	0+	0.001	0.035	0+
30	0+	0+	0.349	0+
40	0+	0+	0.055	0+
50	0.532	0+	0.004	0+
60	0+	0+	0+	0+

References

1. Higginson, A.D.; Barnard, C.J.; Tofilski, A.; Medina, L.; Ratnieks, F. Experimental Wing Damage Affects Foraging Effort and Foraging Distance in Honeybees Apis mellifera. *Psyche* **2011**, doi:10.1155/2011/419793.
2. Roberts, J.C.; Cartar, R.V. Shape of wing wear fails to affect load lifting in common eastern bumble bees (Bombus impatiens) with experimental wing wear. *Can. J. Zool.* **2015**, *93*, 531–537, doi:10.1139/cjz-2014-0317.
3. Pouget, A.; Narain, C. A Conversation with Alexandre Pouget. *Cold Spring Harb. Symp. Quant. Biol.* **2014**, *79*, 285–287, doi:10.1101/sqb.2014.79.14.
4. Ganguli, S.; Sompolinsky, H. Compressed Sensing, Sparsity, and Dimensionality in Neuronal Information Processing and Data Analysis. *Annu. Rev. Neurosci.* **2012**, *35*, 485–508, doi:10.1146/annurev-neuro-062111-150410.
5. Eisthen, H.L. Why Are Olfactory Systems of Different Animals So Similar? *Brain Behav. Evol.* **2002**, *59*, 273–293.
6. Klambt, C. Modes and regulation of glial migration in vertebrates and invertebrates. *Nat. Rev. Neurosci.* **2009**, *10*, 769–779.
7. Delahunt, C.B.; Riffell, J.A.; Kutz, J.N. Biological Mechanisms for Learning: A Computational Model of Olfactory Learning in the Manduca sexta Moth, With Applications to Neural Nets. *Front. Comput. Neurosci.* **2018**, *12*, 102, doi:10.3389/fncom.2018.00102.
8. Rouyar, A.; Deisig, N.; Dupuy, F.; Limousin, D.; Wycke, M.A.; Renou, M.; Anton, S. Unexpected plant odor responses in a moth pheromone system. *Front. Physiol.* **2015**, *6*, 148, doi:10.3389/fphys.2015.00148.
9. Cayre, M.; Strambi, C.; Strambi, A. Neurogenesis in an adult insect brain and its hormonal control. *Nature* **1994**, *368*, 57–59.
10. Martin, J.P.; Beyerlein, A.; Dacks, A.M.; Reisenman, C.E.; Riffell, J.A.; Lei, H.; Hildebrand, J.G. The neurobiology of insect olfaction: Sensory processing in a comparative context. *Prog. Neurobiol.* **2011**, *95*, 427–447.
11. Kvello, P.; Lofaldli, B.; Rybak, J.; Menzel, R.; Mustaparta, H. Digital, three-dimensional average shaped atlas of the heliothis virescens brain with integrated gustatory and olfactory neurons. *Front. Syst. Neurosci.* **2009**, *3*, 14, doi:10.3389/neuro.06.014.2009.
12. Campbell, R.; Honegger, K.; Qin, H.; Li, W.; Demir, E.; Turner, G. Imaging a Population Code for Odor Identity in the Drosophila Mushroom Body. *J. Neurosci.* **2013**, *33*, 10568–10581, doi:10.1523/JNEUROSCI.0682-12.2013.
13. Hige, T.; Aso, Y.; Rubin, G.M.; Turner, G.C. Plasticity-driven individualization of olfactory coding in mushroom body output neurons. *Nature* **2015**, *526*, 258–262.

14. Maia, P.D.; Raj, A.; Kutz, J.N. Slow-gamma frequencies are optimally guarded against effects of neurodegenerative diseases and traumatic brain injuries. *J. Comput. Neurosci.* **2019**, *47*, 1–16.

15. Smith, K.J. Conduction properties of central demyelinated and remyelinated axons, and their relation to symptom production in demyelinating disorders. *Eye* **1994**, *8*, 224–237.

16. Maia, P.D.; Kutz, J.N. Identifying critical regions for spike propagation in axon segments. *J. Comput. Neurosci.* **2014**, *36*, 141–155.

17. Maia, P.D.; Kutz, J.N. Compromised axonal functionality after neurodegeneration, concussion and/or traumatic brain injury. *J. Comput. Neurosci.* **2014**, *27*, 317–332.

18. Maia, P.D.; Hemphill, M.A.; Zehnder, B.; Zhang, C.; Parker, K.K.; Kutz, J.N. Diagnostic tools for evaluating the impact of Focal Axonal Swellings arising in neurodegenerative diseases and/or traumatic brain injury. *J. Neurosci. Methods* **2015**, *253*, 233–243.

19. Maia, P.D.; Kutz, J.N. Reaction time impairments in decision-making networks as a diagnostic marker for traumatic brain injuries and neurological diseases. *J. Comput. Neurosci.* **2017**, *42*, 323–347.

20. Lusch, B.; Weholt, J.; Maia, P.D.; Kutz, J.N. Modeling cognitive deficits following neurodegenerative diseases and traumatic brain injuries with deep convolutional neural networks. *Front. Neurosci.* **2018**, *123*, 154–164.

21. Rudy, S.; Maia, P.D.; Kutz, J.N. Cognitive and behavioral deficits arising from neurodegeneration and traumatic brain injury: A model for the underlying role of focal axonal swellings in neuronal networks with plasticity. *J. Syst. Integr. Neurosci.* **2016**, doi:10.15761/JSIN.1000120.

22. Weber, M.; Maia, P.D.; Kutz, J.N. Estimating memory deterioration rates following neurodegeneration and traumatic brain injuries in a Hopfield network model. *Front. Neurosci.* **2017**, *11*, 623.

23. Kunert, J.M.; Maia, P.D.; Kutz, J.N. Functionality and Robustness of Injured Connectomic Dynamics in C. elegans: Linking Behavioral Deficits to Neural Circuit Damage. *PLoS Comput. Biol.* **2017**, *13*, 1–21, doi:10.1371/journal.pcbi.1005261.

24. Hebb, D.O. *The Organization of Behavior: A Neuropsychological Theory*; Wiley: New York, NY, USA, 1949.

25. Roelfsema, P.R. and Holtmaat, A. Control of synaptic plasticity in deep cortical networks. *Nature Reviews Neuroscience* **2018**, doi:10.1038/nrn.2018.6

26. Riffell, J.A.; Lei, H.; Abrell, L.; Hildebrand, J.G. Neural Basis of a Pollinator's Buffet: Olfactory Specialization and Learning in *Manduca sexta*. *Science* **2012**, doi:10.1126/science.1225483

27. Wilson, R.I.; Laurent, G. Role of GABAergic Inhibition in Shaping Odor-Evoked Spatiotemporal Patterns in the Drosophila Antennal Lobe. *J. Neurosci.* **2005**, *25*, 9069–9079, doi:10.1523/JNEUROSCI.2070-05.2005.

28. Bazhenov, M.; Stopfer, M. Forward and Back: Motifs of Inhibition in Olfactory Processing. *Neuron* **2010**, *67*, 357–358, doi:10.1016/j.neuron.2010.07.023.

29. Lin, A.C.; Bygrave, A.M.; de Calignon, A.; Lee, T.; Miesenböck, G. Sparse, decorrelated odor coding in the mushroom body enhances learned odor discrimination. *Nat. Neurosci.* **2014**, *17*, 559–568.

30. Wasserstein, R.L.; Lazar, N.A. The ASA Statement on p-Values: Context, Process, and Purpose. *Am. Stat.* **2016**, *70*, 129–133, doi:10.1080/00031305.2016.1154108.

31. Wilson, R.I. Neural and behavioral mechanisms of olfactory perception. *Curr. Opin. Neurobiol.* **2008**, *18*, 408–412, doi:10.1016/j.conb.2008.08.015.

32. Masse, N.Y.; Turner, G.C.; Jefferis, G.S. Olfactory Information Processing in Drosophila. *Curr. Biol.* **2009**, *19*, R700–R713, doi:10.1016/j.cub.2009.06.026.

33. Bhandawat, V.; Olsen, S.R.; Gouwens, N.W.; Schlief, M.L.; Wilson, R.I. Sensory processing in the Drosophila antennal lobe increases reliability and separability of ensemble odor representations. *Nat. Neurosci.* **2007**, *10*, 1474–1482.

34. Campbell, R.A.; Turner, G.C. The mushroom body. *Curr. Biol.* **2010**, *20*, R11–R12, doi:10.1016/j.cub.2009.10.031.

35. Galizia, C.G. Olfactory coding in the insect brain: data and conjectures. *Eur. J. Neurosci.* **2014**, *39*, 1784–1795, doi:10.1111/ejn.12558.

36. Perisse, E.; Burke, C.; Huetteroth, W.; Waddell, S. Shocking Revelations and Saccharin Sweetness in the Study of Drosophila Olfactory Memory. *Curr. Biol.* **2013**, *23*, R752–R763, doi:10.1016/j.cub.2013.07.060.

37. Honegger, K.S.; Campbell, R.A.A.; Turner, G.C. Cellular-Resolution Population Imaging Reveals Robust Sparse Coding in the Drosophila Mushroom Body. *J. Neurosci.* **2011**, *31*, 11772–11785, doi:10.1523/JNEUROSCI.1099-11.2011.

38. Cassenaer, S.; Laurent, G. Hebbian STDP in mushroom bodies facilitates the synchronous flow of olfactory information in locusts. *Nature* **2007**, *448*, 709–713.

39. Hammer, M.; Menzel, R. Learning and memory in the honeybee. *J. Neurosci.* **1995**, *15*, 1617–1630.

40. Hammer, M.; Menzel, R. Multiple Sites of Associative Odor Learning as Revealed by Local Brain Microinjections of Octopamine in Honeybees. *Learn. Mem.* **1998**, *5*, 146–156.

41. Løfaldli, B.; Kvello, P.; Mustaparta, H. Integration of the antennal lobe glomeruli and three projection neurons in the standard brain atlas of the moth Heliothis virescens. *Front. Syst. Neurosci.* **2010**, *4*, 5, doi:10.3389/neuro.06.005.2010.

42. Wang, J.; Hamm, R.J.; Povlishock, J.T. Traumatic Axonal Injury in the Optic Nerve: Evidence for Axonal Swelling, Disconnection, Dieback, and Reorganization. *J. Neurotrauma* **2011**, *28*, 1185–1198, doi:10.1089/neu.2011.1756[PII].

43. Dayan, P.; Abbott, L.F. *Theoretical Neuroscience: Computational and Mathematical Modeling of Neural Systems*; The MIT Press: Cambridge, MA, USA, 2005.

44. Higham, D.J. An Algorithmic Introduction to Numerical Simulation of Stochastic Differential Equations. *SIAM Rev.* **2001**, *43*, 525–546, doi:10.1137/S0036144500378302.

45. Papadopoulou, M.; Cassenaer, S.; Nowotny, T.; Laurent, G. Normalization for Sparse Encoding of Odors by a Wide-Field Interneuron. *Science* **2011**, *332*, 721–725, doi:10.1126/science.1201835.

46. Gupta, N.; Stopfer, M. Functional Analysis of a Higher Olfactory Center, the Lateral Horn. *J. Neurosci.* **2012**, *32*, 8138–8148, doi:10.1523/JNEUROSCI.1066-12.2012.