A control engineering perspective on the advantages of efference copies

Benjamin Campbell, Huai-Ti Lin, Holger Krapp

July 24, 2023

Imperial College London, Exhibition Rd, South Kensington, London, SW7 2BX
E-Mail: b.campbell21@imperial.ac.uk

Abstract

Biological systems have evolved to perform high-speed voluntary movements whilst maintaining robustness and stability. This paper examines a control architecture based on the principles of efference copies found in insect sensorimotor control which we call the inner-outer loop (IOL) controller. Within a control engineering framework, we benchmark the advantages of this control architecture against two common engineering control schemes: a pure feedback (PFB) and a two-degree-of-freedom (2DoF) controller. Our results show that the architecture performs particularly well for systems with noisy sensors, and systems with time-delays. We explain a hypothesis that the structure of the insect sensorimotor system allows the sensors to increase their operating range. Finally, we explore the robustness of the IOL controller and show that it achieves this increased performance with equal stability margins and stability robustness. Finally, we discuss the benefits of the IOL controller in relation to the properties of sensors employed in insect sensorimotor control.

1 Introduction

Biological systems use a different architecture to control attitude and movement compared to those frequently used in engineering. A typical feedback controller (Fig. 1a), works by sensing the current state, e.g. position, velocity, angle, and subtracting that from the desired state to give an error. The error is provided to a controller which generates an actuation command, e.g. force, voltage, torque. The actuation acts on the dynamics of the system to reduce the error, thereby ensuring that the state you are in, is your desired state [1].

In 2012 Krapp theorised that the sensorimotor control in insects differs by having two layers: an inner-loop and an outer-loop [2, 3, 4]. The inner-loop consists of a stabilising controller that compensates for any perceived change in state [5]. inner-loop control alone would result in a system that is unable to change state [4]. The outer-loop then has two roles which accommodate desired changes in state. Firstly, it directly provides an actuation command to the motor system to carry out the desired change. Secondly, it cancels the expected sensor response as a result of the actuation command. This means that the stabilising controller in the inner-loop does not undo the desired change (see [4]). In biology, this expected sensor response is called an efference copy [6, 7, 8, 9, 10].

Since 2012 several experimental studies provided physiological evidence for efference copies in insects. In a combined neurophysiological and behavioural study Kim et al. (2015) found that insects inhibit their sensor response during spontaneous yaw rotations (i.e. use

efference copies) [11]. Then in 2021 Fenk et al. showed that this was not the case for actuation due to external disturbances. These two results strongly support the inner-loop and outer-loop structure for insect sensorimotor control [5].

A control architecture that uses efference copies, using a predicted sensor response to desired changes in state, is not unique to insects; it is also found in primates [7]. However, compared to insects cellular evidence for its implementation is more challenging to find in primates.

Here we use the term forward model in the way it is often applied in the context of insect sensorimotor control [12]. It is a function, or filter, that maps a desired change in state into an expected sensor response. It is not an internal map of the world or an understanding of what will happen when certain decisions are taken, like what is described in [13]. It is also different to the concept of a forward model in much of the cerebellum literature [14, 15, 16, 17, 18], where typically a forward model refers to a model of the dynamics of the system. Whilst there are many hypotheses for different cerebellum control schemes, the most typical uses a model of the dynamics of the system to find an inverted model which can be used to generate a motor command [14]. Here, we refer to a forward model as a filter that predicts a sensor response from a desired state change. It is therefore not just a model of the dynamics, but a filter that models the insect controller, the dynamics, and the sensor properties.

The second aspect of the outer-loop is the generation of the actuation command to execute the desired state change or desired movement. Since the expected

sensor response is cancelled, this actuation must have been generated by a feedforward controller (otherwise there would be no sensor response to cancel). A feedforward controller is another function, or filter, that this time maps a desired state change into a motor command resulting in muscle contraction. Most of the time in engineering control theory feedforward controllers are used to supplement the actuation of the feedback controller when changing desired state, referred to as a two-degree-of-freedom (2DoF) controller (Fig. 1b) [19]. We differentiate this type of feedforward control from what is hypothesised in the insect sensorimotor system by name and by a simple principle. We refer to the feedforward controller in the insect sensorimotor system as a desired change controller; where the principle difference is that instead of supplementing the actuation of a feedback controller, it provides *all* of the actuation commands needed to change state.

The three controllers we compare in this paper are given in Figure 1. There is the feedback controller in Figure 1a, which we will refer to as the *pure* feedback controller (PFB), to differentiate it from the feedback component of the other two controllers. There is the two-degree-of-freedom controller (2DoF) in Figure 1b. Then there is the efference copy inspired inner-outer loop controller (IOL) in Figure 1c, which combines the desired change controller, forward model, and feedback control.

This paper is the first to analyse the inner-outer loop (IOL) insect sensorimotor control system and compare it with other typical controllers. Its primary structure has evolved over hundreds of millions of years specifically to control attitude and voluntary movements. By relating a control systems analysis of the structure, to the common sensor and dynamic properties of insects, we attempt to answer: what advantages does biology gain from using this architecture?

2 Formulation of the inner-outer insect sensorimotor control architecture

By far the most widespread and convenient way to examine control architectures is by using the complex frequency (Laplace) domain [1]. It is a mathematical tool that allows interactions between temporal filters to be treated algebraically, and convolution to be treated as multiplication.

The formulation in the Laplace domain for the inner-outer loop (IOL) controller is given in Figure 1c. The biological evidence it is based on primarily consists of the observation of efference copies [11, 5] and the knowledge that there must therefore be a "desired change controller" which actuates the desired changes in state. The order and structure of the outer-loop is still unknown [12].

Figure 1: Three control architectures, inside the blocks are the Laplace transforms of the dynamics, with block labels in black below. The red labels name the signals between blocks. The signal labelled "actuation" refers specifically to the actuation command, the actuator dynamics are then included in $P(s)$. (a) Shows a pure feedback controller (PFB) where the state is measured by sensors with some dynamics, that is subtracted from the desired state to give an error. The error is passed through a feedback controller which generates an actuation command. (b) Shows the two-degree-of-freedom (2DoF) controller where a feedback controller is combined with a feedforward component. (c) Shows a control systems formulation of the IOL control architecture. The desired state is sent to the feedforward desired change controller that generates the motor command to actuate the movement. The desired state is also sent to a forward model which predicts the sensory consequences of the desired state change and generates an efference copy. That efference copy signal cancels the sensor response. The feedback controller receives the error in sensor response and predicted sensor response.

There are two alternative possibilities for the structure shown in Figure 1c. In Figure 1c the forward model $F(s)$ and the desired change controller $C_2(s)$ receive their input from some internal decision to make a desired change in state. The two alternative possibilities are: (i) that the desired change controller $C_2(s)$ receives its input from the forward model $F(s)$, and (ii) that the forward model $F(s)$ receives its input from the desired change controller $C_2(s)$. The first option can be easily ruled out using a causality argument. If there was a time delay in the sensor, it would be needed in the forward model. If the desired change controller came after the forward model, it would need to contain a positive time delay which would be impossible. (ii) was ruled out because given unstable system dynamics an unstable forward model would be needed. Whereas, in the structure we opted to analyse, both the desired

change controller and the forward model take the desired state as their input. This means each block in the diagram is causal. It also means that given unstable dynamics, provided you have a stabilising desired change controller, the forward model can be stable. It is worth noting that there have not been physiological experiments showing the exact origin of the feedforward control command generated by the desired change controller $C_2(s)$, or the efference copy generated by the forward model $F(s)$.

There are three main metrics used to assess the suitability of a controller: (i) disturbance rejection, how quickly can the controller drive the actuators to remove external perturbations (e.g. gusts of wind), (ii) performance in reference tracking, how well the actual state tracks the desired state, and (iii) robustness, if the controller is built around an incorrect system dynamics model, will it still be stable?

By finding the transfer functions of the closed-loop controlled systems we can answer these questions. First, looking at disturbance rejection: the transfer function from a disturbance $U(s)$ to the actual state $X(s)$, is given by equation 1. Where $G(s)$ is the transfer function for how the disturbance affects the state.

$$\frac{X}{U}(s) = \frac{G}{1 + HC_1P} \tag{1}$$

The same equations are obtained for all three architectures in Figure 1. If all three architectures use the same feedback controller, they will all have the exact same performance in handling disturbances. This is intuitive, since in order to reject a disturbance it must be sensed, and therefore feedback control must be used to remove its impact.

For performance in tracking the desired state, all three controllers have different transfer functions. The transfer functions for the pure feedback controller (PFB, Fig. 1a) and 2DoF controller (Fig. 1b) are given in equations 2 and 3 respectively.

$$\frac{X}{X_0}(s) = \frac{C_1P}{1 + HC_1P} \tag{2}$$

$$\frac{X}{X_0}(s) = \frac{(C_1 + C_3)P}{1 + HC_1P} \tag{3}$$

The only difference between the two is that the 2DoF controller has an additional feedforward term. An advantage of the 2DoF controller over the PFB is that the feedforward controller $C_3(s)$ can be designed to improve reference tracking performance without impacting disturbance rejection. This means that the feedback component of the 2DoF controller can be designed with an emphasis on disturbance rejection, whereas in the PFB, C_1 must be designed with both in mind, as explained in [19].

The IOL control architecture (Fig. 1c) has the transfer function given by 4.

$$\frac{X}{X_0}(s) = \frac{C_2P + FC_1P}{1 + HC_1P} \tag{4}$$

However, the forward model $F(s)$ is not a random filter, it should turn the desired state changes into the expected sensor response. By looking at Figure 1c, following the loop from desired state to sensor response, we can see that the forward model should predict C_2PH. By substituting $F(s) = C_2PH$ into equation 4 we get a different-looking transfer function from desired state to actual state given in equation 5.

$$\frac{X}{X_0}(s) = C_2P \tag{5}$$

This effectively corresponds to open-loop control. Because there is a feedback loop, however, the disadvantages of *pure* open-loop control do not apply to the IOL architecture. The architecture can still compensate for external disturbances, and it can still deal with unstable dynamics.

For the first control metric (i) disturbance rejection, the equations show that all three controllers will have the exact same performance, given the same $C_1(s)$. This is not the case for the second control metric (ii) reference tracking, where the three architectures are formalised by different equations; the majority of the rest of the paper will be on simulations showing the differences in performance. Then finally (iii) the robustness of the controller to differences in the real system dynamics and expected system dynamics will be discussed in Section 4.4.

3   Methods

To explore the differences in the reference tracking performance between the inner-outer loop (IOL) controller and the two common engineering controllers we follow a simple procedure: select some system and sensor dynamics, design all three controllers around controlling those dynamics, and compare the response to a step change in the desired state. The responses of the closed-loop systems were simulated using a discrete approximation (ode45) of continuous time in Simulink MATLAB.

For a fair comparison, the limits of the actuators for each of the three systems must be the same. To ensure this, the maximum allowable actuation of the IOL control architecture was set to be within 1 % of the minimum of the other two controllers. The feedback controller was a proportional-integral-derivative (PID) controller, and the two-degree-of-freedom controller was the parallel 2DoF controller structure provided by MATLAB. The 2DoF controller has a PID in feedback and two setpoint weights in the feedforward component: one on the proportional term and one on the derivative term. Both controller architectures were configured using MATLABs optimal PID tuner with a preference put on fast response dynamics. The response time parameter in the tuner was set to be between 0.15 s and 0.175 s, and the robustness parameter was set to 0.45.

	Inner-outer controller
	Feedback controller
	Two-degree-of-freedom controller

Figure 2: Step response of all three controllers on system dynamics 1 in Table 1 with parameters $\tau = 0$ and $\omega = 10$. In both graphs, the response of MATLABs tuned 2DoF controller, and optimally tuned PID pure feedback controller (PFB) are compared with the response of the IOL control architecture. (a) shows the desired state and the actual state, the PFB overshoots 7.8 %, the 2DoF controller overshoots 5.7 %, whilst the IOL controller does not overshoot. The IOL controller has a settle time of 25 ms, compared with 39 ms for the PFB and 44 ms for the 2DoF controller. (b) shows the actuation used by the three controllers. The maximum actuation for the IOL controller is designed to be the same as the minimum of the other two, in this case the 2DoF controller (1.23 units), the PFB had a maximum actuation of 1.31 units.

Table 1: Three different transfer functions of system dynamics and sensor combinations, $P(s)H(s)$ in Figure 1.

1	2	3
$\dfrac{e^{-s\tau}}{\left(\dfrac{s}{\omega}+1\right)}$	$\dfrac{e^{-s\tau}}{s\left(\dfrac{s}{\omega}+1\right)}$	$\dfrac{e^{-s\tau}}{\left(\dfrac{s}{\omega}+1\right)^2}$

The controller architectures were tested on 18 different system dynamics. There were three structures given in Table. 1, each parameterised with different bandwidths and sensor delays, ω and τ respectively. This was to provide a range where either the sensor delay or the apparent delay due to the system dynamics dominated. We used three different sensor delays: 0 ms, 50 ms, and 100 ms. Each with two bandwidths: $1\,\mathrm{rad\,s^{-1}}$, and $10\,\mathrm{rad\,s^{-1}}$. To test each controller on the 18 dynamics, we had 54 different controllers.

An example step response controlling the first lot of system dynamics in Table 1 (with $\omega = 10$ and $\tau = 0$) is given in Figure 2. The MATLAB optimal gains for the feedback controller were $K_p = 0.86$ and $K_i = 19.03$, for the 2DoF controller the feedback component in this case was the same. The feedforward component is categorised by the two terms $b = 0$ and $c = 1$, where b and c are the setpoint weight on the proportional and derivative term respectively. $C_2(s)$ was chosen to have a zero at $10\,\mathrm{rad\,s^{-1}}$, and have poles at $27.75\,\mathrm{rad\,s^{-1}}$ and $20\,\mathrm{rad\,s^{-1}}$. Correspondingly, the forward model became a second-order low-pass filter with the same poles. The data on time delays was collected by repeating this process for all 18 different systems.

The demonstration of handling noisy systems was done by tuning two PID controllers to the system dynamics and only altering the response time parameter in the tuner, and therefore the bandwidth. It is not possible to perfectly simulate white noise due to its infinite frequency spectrum. Instead, a white noise approximation using a random sequence with a correlation time constant was injected at the sensor. The correlation time constant must be significantly smaller than the smallest time constant in the system being simulated, therefore, we used a value of 0.1 ms and a power of 0.001. The IOL controller was given the same inner-loop as the lower bandwidth feedback controller. The constraint of maximum actuation was relaxed for this demonstration, given the actuation used was largely determined by the response to noise.

4   Results

An example of how the three controllers perform when designed around the same system dynamics is given in Figure 2. To assess the performance of the reference tracking we measured the overshoot and the settle time. The overshoot is a measure of how far the actual state exceeds the desired state before settling; and the settle time is the time to reach and stay within 2 % of the desired state as shown in Figure 2a. We also measured the actuation costs in the form of the maximum actuation required, as well as the total actuation which was calculated by integrating over time until all three systems had settled. These values can be seen in Figure 2b.

The first observation in Figure 2a is that the inner-outer loop (IOL) control architecture does not overshoot, which is consistent with experiments where humans are tasked with reaching an object and the position of their hand is tracked [20, 18]. It can also be seen, as discussed above, that all three controllers respond to disturbances in the same way.

Figure 3: Performance of the three different controllers with different sensor delays. (a) Percentage overshoot of each controller at different sensor delays. (b) Settle time in seconds for each controller at different sensor delays.

4.1 Advantage 1: Mitigating the impact of sensor delays

The performance results for the overshoot and settle time at different sensor delays are shown in Figure 3. The inner-outer loop (IOL) controller never overshoots by design (Fig. 3a). For the other two controllers, the overshoot dramatically increases as the sensor delay is increased.

The average settling time (Fig. 3b) of the IOL controller with no sensor delay shows very little improvement compared with the two typical engineering controllers. However, as the sensor delay is increased, the difference in performance sees a large increase; the IOL controller outperforms the other controllers by an increasing margin. This is an inherent advantage of the IOL controller, as seen in equation 5 the transfer function from the desired state to the actual state is independent of the sensor properties. This means the reference tracking performance is independent of any delays in the sensors. Although, since a feedback component is required in all three control architectures for rejecting disturbances, the disturbance rejection performance will suffer from increased sensor delays. In Figure 3b there is a downward trend for the IOL control architecture when the sensor delay is increased. This is not because its performance improves with increasing sensor delays. It is that in comparison the other controllers use more actuation, therefore the IOL controller is "allowed" more actuation, and hence settles faster.

4.2 Advantage 2: Mitigation of sensor noise

In classical feedback control, there is a trade-off between noise rejection and reference tracking. If the sensor signal is heavily filtered, the feedback control will be slowed down. If the feedback controller has a high bandwidth, it will track the desired state faster. However, it will come at the cost of a stronger response

to noise from the sensors. A more formal mathematical evaluation of this principle is given in [21].

Figure 4: Noise anecdote

With the inner-outer loop (IOL) control architecture, this trade-off exists for the inner-loop only. The outer-loop is not affected by sensor noise. As a result, and because the outer-loop controls all desired changes in state, the performance of reference tracking is independent of sensor noise. This is demonstrated in Figure 4.

The behaviour of two PFB controllers highlights the trade-off. One of them rises to the desired state (unity) very quickly but is highly responsive to noise. The other one rejects the noise, however, it rises to the desired state much slower. The IOL control architecture is able to reject noise in the inner-loop, and change state quickly with the noise-independent outer-loop.

4.3 Advantage 3: Stopping sensor saturation

In their study on the use of efference copies in fruitflies, Kim et al. (2015) found the cancellation of the expected sensor response during spontaneous yaw turns

to take place in specific neurons of the animal's motion vision pathway [11]. Those neurons, called HS-cells, sense yaw rotations by processing optic flow [22] and send their output to various motor systems that control the fly's gaze and locomotor activities [23]. Like any other neuron in the nervous system, HS cells have a limited dynamic output range. Cancelling the predicted sensory response to desired state changes keeps the signal of HS cells within its dynamic operating range and therefore retains high sensitivity to external disturbances, effectively preventing output saturation. Besides keeping sensors in their operating range and therefore sensitive to state changes caused by external disturbances, responding only to unpredictable changes is a more energy-efficient strategy.

4.4 Controller robustness

Perfectly modelling a real system is not possible, it is, therefore, important to consider what happens when a controller is built around a simplified or inaccurate understanding of the sensors, actuators and system dynamics.

Firstly, we think about what happens if the controller has been designed around system dynamics that are not accurate, or have changed. To answer this we use a common metric in robust control, which we refer to as "stability robustness". Stability robustness is a mathematical formulation for *how wrong the expected system dynamics can be such that the whole system retains stability.* [24]

Mathematically, we can answer this question by applying a multiplicative error to the system dynamics and deriving conditions for stability for the different controllers using the small-gain theorem [24]. All three controllers result in having the exact same stability robustness (see Appendix A). So, despite the addition of the desired change controller and forward model, the inner-outer loop (IOL) controller provides the exact same stability robustness to "modelling" errors. This can be understood by considering the inner-loop controller as something that uses feedback to mitigate any unanticipated sensor response. This includes modelling errors: if the system dynamics aren't acting how the forward model expected, the feedback controller will act to obtain the expected sensor response.

In addition, we can find a condition at which the steady-state error in the feedback controller will converge to zero (App. A).

An interesting feature of the IOL architecture is that instead of stabilising state units, it stabilises sensory units. For example in the fly visual system, instead of stabilising roll angle, a potential internal representation of the state, the inner loop would be stabilising the response of a neuron that is tuned to sense a roll rotation. This has some interesting effects: if the sensors act as high-pass filters, or derivatives, for the state units this could result in steady-state errors. This can be seen by looking at the order of the equation 8 in Appendix A. But, it can also be explained intuitively; the stabilising inner loop may not see the

unexpected different sensor response for long enough to be able to mitigate it. This may appear to be a disadvantage, however, it works as an advantage for a different type of error. In biological systems sensors are less reliable, turning to the fly visual system again, the sensor response is dependent on the spatial properties of the environment [25]. Relating this back, if, instead of the system dynamics being different to the expectation, the sensor response is different to the expectation: the resulting error due to this could be reduced. In the two typical engineering controllers, if the sensors are slightly incorrect, the result is stabilising around the wrong point. With the IOL controller and transient filters, a trade-off can be made. By allowing errors resulting from incorrect system dynamics, resulting errors from incorrect sensors will be smaller. It is worth noting that this trade-off can also impact the response to external disturbances. This effect is specific and dependent on the system being controlled. Therefore, we have provided the six transfer functions - two types of error for the three different control architectures - in Appendix A.

To summarise, there is the exact same risk of becoming unstable due to unexpected system dynamics with the IOL control architecture as there is with the PFB and 2DoF controller - despite the more complicated structure. Secondly, in the engineering controllers, if there is a system dynamics modelling error there will generally not be any steady state error between the actual state and desired state. However, if there is an error in the sensor model in engineering controllers they will stabilise around the wrong point. The same applies to the IOL architecture when the sensors measure the absolute value. However, that is not always the case if the sensors in the IOL act like high-pass filters or derivatives, as commonly found in biology. Instead, small steady-state errors will be allowed when the system dynamics are different to expected. But the advantage is that when the sensor model is wrong, it will not simply stabilise around the wrong point. The steady-state error due to unexpected sensor dynamics will be reduced.

5 Conclusions

Motivated by recent studies in the insect sensorimotor system we developed a phenomenological control architecture and benchmarked its performance against common control engineering approaches. We find that the performance in tracking the desired state is independent of the sensor response for accurate controllers. This means that the control architecture is particularly good at handling sensors with time delays and significant levels of noise.

Sensory delays are common in biological systems [26, 27, 28], yet many animals are able to deliver high-performance control. This work shows that the inner-outer loop explains how that could be done. In particular for insects whose intrinsic system dynamics are fast, but still have significant sensor delays. This could also

suggest a reason why they have been first recorded in the fly visual system, one of the slower sensory systems in the fly [28, 25].

Noise is also prevalent in the nervous system [29]. We show that the inner-outer loop (IOL) architecture allows fast goal-directed behaviours despite sensor noise. This is due to the inherent property of the outer loop control signals being independent of the sensor response. Secondly, by subtracting the expected sensor response the sensors are kept in their operating range, maintaining sensitivity to external perturbations. This could cause a stronger signal power, and therefore, for a fixed level of noise, a better signal-to-noise ratio.

The visual system in flies, where there is physiological evidence for efference copies, does not respond linearly to velocity. Instead, it has a response that contains significant modulations that depend on the spatial properties of the environment [30, 28]. Despite this, flies remain incredibly agile and fast. We have shown the ability of the IOL architecture to suppress sensor noise and maintain fast desired state changes. A hypothesis is that this could provide an explanation for how the insect flight control system handles these large modulations from the visual sensors.

Another result is that the actuation costs were on average lower than for the other two controllers. The maximum actuation was not informative due to the design constraint ensuring a fair comparison. However, the total actuation over the course of settling for the IOL architecture averaged across the 18 systems tested was 28.4 % lower than the average of the other two controllers. Whilst a definite line between actuation and energy can not be made. It is most often the case in physical systems that more actuation results in more energy use. Energy efficiency is a general principle in biology and the most significant constraint for any successful evolutionary developments [31, 32]. The results on lower actuation are probably linked to the property of the IOL controller to not overshoot when assuming a new desired state.

A point that raises interesting questions is that the feedback controller does not control parameters in state units, but instead controls parameters in sensory units. This could potentially lend itself to a distributed control system. Instead of having many sensors integrated to provide the current state and then having a single, complicated, feedback controller. It allows each sensor system to sense something unanticipated and feedback directly to a corresponding motor system. This could potentially reduce computational overheads, inverting and combining information from multiple sensors, and provide a shorter neural pathway from sensor to muscle. Thereby, reducing feedback delays, and improving disturbance rejection performance. There is another potential advantage of this strategy. Different sensing modalities have different response delays. The IOL control architecture could provide a mechanism to utilise sensors operating at different speeds. Something similar is found in the fly gaze stabilisation system where the fast mechanical sensors provide muscle commands quickly, and the visual system provides muscle commands slower [33]. This may be a potential advantage of feedback controllers being in sensory units instead of state units.

Another quality of the IOL controller is that the desired changes in state, when given unlimited actuation capabilities, can often be arbitrarily fast. This is not the case for feedback control where high-gains can lead to instability, in particular for systems with sensor delays. All of the system dynamics evaluated here can be driven arbitrarily fast by the desired change controller. In the case of zeros in the numerator of the system dynamics, the settling time would not be able to be arbitrarily small. More biological system identification studies would shed more light on whether this potential benefit has been exploited in natural sensorimotor control designs.

This paper demonstrates that there are potential advantages to biological systems using control structures like the inner-outer loop controller described here. In particular given the agility and high-speed manoeuvres that are required in nature, with the sometimes delayed sensors available. Future work should include how these advantages can be retained through changing system dynamics. Whilst we show that it is equally robust, the performance is better when the system is fully adapted. The continuing work on adaptation in the insect sensorimotor system could provide valuable insights into how this is achieved, both in biological sensors [34, 35, 36] and during different locomotor states [37, 38, 39, 40]. It has been shown in electric fish that they can adapt their sensorimotor control to stabilise even intentionally destabilising dynamics [41]

To summarise, we explore a control architecture based on principles found in insect sensorimotor control. We show the architecture has advantages in systems with sensor delays, in particular when the dynamics are fast as is the case in flying insects. We have also shown that the inner-outer loop controller has very good properties when it comes to coping with noise from sensors, and provides a natural mechanism for keeping sensors in their operating range. These findings point towards the suitability of the architecture in biology and aim to provide some explanation for the advantages that efference copies can have in the sensorimotor control system.

A   Appendix 1

First, we look at the robustness of the control schemes to system dynamics that are different to the expectation. A multiplicative error can be modelled in the Laplace domain by making the plant $P(s)(1+\Delta_P(s))$; where $\Delta_P(s)$ is the Laplace transform of the error, and $P(s)$ is the model of the plant dynamics that were expected [24]. The conditions for closed-loop stability in both the inner-outer loop control system and the PFB and 2DoF controller can be derived as an extension to Nyquist using the small-gain theorem [24]. The result of a sufficient condition for closed-loop stability is given

in Equation 6

$$\|\Delta_P\|_\infty \left\|HC_1P(1 - HC_1P)^{-1}\right\|_\infty < 1 \qquad (6)$$

If this condition is satisfied, the system will be closed-loop stable [24]. It is immediate that the transfer functions $C_2(s)$, and $F(s)$ do not appear in Condition 6, and if they themselves are internally stable, they will not cause instability.

Even if the system with different to expected system dynamics remains stable, the performance can still suffer. To evaluate this we can find the transfer functions for the three control architectures including the multiplicative error. First, the inner-outer loop controller we substitute the plant with error into Equation 4, and leave the forward model not including the error. That gives the transfer function in Equation 7.

$$\frac{X}{X_0}(s) = C_2P\left(1 + \frac{\Delta_P(s)}{1 + HC_1P(1 + \Delta_P(s))}\right) \qquad (7)$$

From the final value theorem, the steady-state error in response to a unit step change in the desired state (e_{ss}) is given in Equation 8.

$$\lim_{s \to 0}\left\{\frac{C_2P\Delta_P(s)}{1 + HC_1P(1 + \Delta_P(s))}s\right\} = e_{ss} \qquad (8)$$

For the PFB and 2DoF the transfer functions are given in equations 9 and 10 respectively.

$$\frac{X}{X_0}(s) = \frac{C_1P(1 + \Delta_P(s))}{1 + HC_1P(1 + \Delta_P(s))} \qquad (9)$$

$$\frac{X}{X_0}(s) = \frac{(C_1 + C_3)P(1 + \Delta_P(s))}{1 + HC_1P(1 + \Delta_P(s))} \qquad (10)$$

Now we look at the impact of sensor dynamics that are different to expected. We take the same approach as above, but now there is a multiplicative error applied to the sensors. $H(1 + \Delta_H(s))$, where $H(s)$ is the expected sensor dynamics, and $\Delta_H(s)$ is the unexpected difference to the expectation. The stability robustness is the same as in Equation 6 for all three control architectures, but Δ_P is replaced by Δ_H.

The transfer functions for the IOL, PFB, and 2DoF control architectures are given in Equations 11, 12, and 13.

$$\frac{X}{X_0}(s) = C_2P\left(\frac{1 + HC_1P}{1 + HC_1P(1 + \Delta_H(s))}\right) \qquad (11)$$

$$\frac{X}{X_0}(s) = \frac{C_1P}{1 + HC_1P(1 + \Delta_H(s))} \qquad (12)$$

$$\frac{X}{X_0}(s) = \frac{(C_1 + C_3)P}{1 + HC_1P(1 + \Delta_H(s))} \qquad (13)$$

Using the six transfer functions (three architectures and two types of error) the trade-offs for specific system dynamics and feedback controllers can be explored.

References

[1] K. Ogata *et al.*, *Modern control engineering*, vol. 5. Prentice hall Upper Saddle River, NJ, 2010.

[2] B. J. Hardcastle and H. G. Krapp, "Evolution of biological image stabilization," *Current Biology*, vol. 26, no. 20, pp. R1010–R1021, 2016.

[3] H. G. Krapp, "Multisensory control of stabilization reflexes," tech. rep., IMPERIAL COLL OF SCIENCE AND TECHNOLOGY LONDON (UNITED KINGDOM), 2012.

[4] H. G. Krapp, "How a fly escapes the reflex trap," *Nature neuroscience*, vol. 18, no. 9, pp. 1192–1194, 2015.

[5] L. M. Fenk, A. J. Kim, and G. Maimon, "Suppression of motion vision during course-changing, but not course-stabilizing, navigational turns," *Current Biology*, vol. 31, no. 20, pp. 4608–4619, 2021.

[6] B. Bridgeman, "A review of the role of efference copy in sensory and oculomotor control systems," *Annals of biomedical engineering*, vol. 23, no. 4, pp. 409–422, 1995.

[7] T. B. Crapse and M. A. Sommer, "Corollary discharge across the animal kingdom," *Nature Reviews Neuroscience*, vol. 9, no. 8, pp. 587–600, 2008.

[8] R. W. Sperry, "Neural basis of the spontaneous optokinetic response produced by visual inversion.," *Journal of comparative and physiological psychology*, vol. 43, no. 6, p. 482, 1950.

[9] E. von Holst, "The reafference principle," *The Behavioral Physiology of Animals and Man*, 1973.

[10] H. Straka, J. Simmers, and B. P. Chagnaud, "A new perspective on predictive motor signaling," *Current Biology*, vol. 28, no. 5, pp. R232–R243, 2018.

[11] A. J. Kim, J. K. Fitzgerald, and G. Maimon, "Cellular evidence for efference copy in drosophila visuomotor processing," *Nature neuroscience*, vol. 18, no. 9, pp. 1247–1255, 2015.

[12] B. Webb, "Neural mechanisms for prediction: do insects have forward models?," *Trends in neurosciences*, vol. 27, no. 5, pp. 278–282, 2004.

[13] M. Mischiati, H.-T. Lin, P. Herold, E. Imler, R. Olberg, and A. Leonardo, "Internal models direct dragonfly interception steering," *Nature*, vol. 517, no. 7534, pp. 333–338, 2015.

[14] D. M. Wolpert, R. C. Miall, and M. Kawato, "Internal models in the cerebellum," *Trends in cognitive sciences*, vol. 2, no. 9, pp. 338–347, 1998.

[15] D. M. Wolpert and M. Kawato, "Multiple paired forward and inverse models for motor control," *Neural networks*, vol. 11, no. 7-8, pp. 1317–1329, 1998.

[16] M. Kawato, "Feedback-error-learning neural network for supervised motor learning," in *Advanced neural computers*, pp. 365–372, Elsevier, 1990.

[17] P. Dean, J. Porrill, C.-F. Ekerot, and H. Jörntell, "The cerebellar microcircuit as an adaptive filter: experimental and computational evidence," *Nature Reviews Neuroscience*, vol. 11, no. 1, pp. 30–43, 2010.

[18] R. Osu, S. Hirai, T. Yoshioka, and M. Kawato, "Random presentation enables subjects to adapt to two opposing forces on the hand," *Nature neuroscience*, vol. 7, no. 2, pp. 111–112, 2004.

[19] M. Araki and H. Taguchi, "Two-degree-of-freedom pid controllers," *International Journal of Control, Automation, and Systems*, vol. 1, no. 4, pp. 401–411, 2003.

[20] L. Botzer and A. Karniel, "Feedback and feedforward adaptation to visuomotor delay during reaching and slicing movements," *European Journal of Neuroscience*, vol. 38, no. 1, pp. 2108–2123, 2013.

[21] K. Astrom and R. Murray, "Frequency domain design," *Feedback Systems: An Introduction for Scientists and Engineers*, pp. 315–346, 2008.

[22] K. Hausen, "Decoding of retinal image flow in insects," *Visual Motion and its Role in the Stabilization of Gaze*, pp. 203–235, 1993.

[23] R. Hengstenberg, "Gaze control in the blowfly calliphora: a multisensory, two-stage integration process," in *Seminars in Neuroscience*, vol. 3, pp. 19–29, Elsevier, 1991.

[24] M. Green and D. J. Limebeer, *Linear robust control*. Courier Corporation, 2012.

[25] G. K. Taylor and H. G. Krapp, "Sensory systems and flight stability: what do insects measure and why?," *Advances in insect physiology*, vol. 34, pp. 231–316, 2007.

[26] N. Saijo, I. Murakami, S. Nishida, and H. Gomi, "Large-field visual motion directly induces an involuntary rapid manual following response," *Journal of Neuroscience*, vol. 25, no. 20, pp. 4941–4951, 2005.

[27] P. B. Matthews, "The human stretch reflex and the motor cortex," *Trends in neurosciences*, vol. 14, no. 3, pp. 87–91, 1991.

[28] A. Borst and M. Egelhaaf, "Principles of visual motion detection," *Trends in neurosciences*, vol. 12, no. 8, pp. 297–306, 1989.

[29] A. A. Faisal, L. P. Selen, and D. M. Wolpert, "Noise in the nervous system," *Nature reviews neuroscience*, vol. 9, no. 4, pp. 292–303, 2008.

[30] H. G. Meyer, J. P. Lindemann, and M. Egelhaaf, "Pattern-dependent response modulations in motion-sensitive visual interneurons—a model study," *PLoS One*, vol. 6, no. 7, p. e21488, 2011.

[31] S. B. Laughlin, "Insects as a guide to neural circuit design," *Handbook of Brain Microcircuits*, p. 404, 2010.

[32] J. E. Niven and S. B. Laughlin, "Energy limitation as a selective pressure on the evolution of sensory systems," *Journal of Experimental Biology*, vol. 211, no. 11, pp. 1792–1804, 2008.

[33] D. A. Schwyn, F. J. Heras, G. Bolliger, M. M. Parsons, H. G. Krapp, and R. J. Tanaka, "Interplay between feedback and feedforward control in fly gaze stabilization," *IFAC Proceedings Volumes*, vol. 44, no. 1, pp. 9674–9679, 2011.

[34] S. B. Laughlin and R. C. Hardie, "Common strategies for light adaptation in the peripheral visual systems of fly and dragonfly," *Journal of comparative physiology*, vol. 128, pp. 319–340, 1978.

[35] J. Van Hateren, "Processing of natural time series of intensities by the visual system of the blowfly," *Vision research*, vol. 37, no. 23, pp. 3407–3416, 1997.

[36] R. A. Harris, D. C. O'Carroll, and S. B. Laughlin, "Contrast gain reduction in fly motion adaptation," *Neuron*, vol. 28, no. 2, pp. 595–606, 2000.

[37] R. Wolf, A. Voss, S. Hein, and M. Heisenberg, "Can a fly ride a bicycle?," *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences*, vol. 337, no. 1281, pp. 261–269, 1992.

[38] R. Rosner, M. Egelhaaf, and A.-K. Warzecha, "Behavioural state affects motion-sensitive neurones in the fly visual system," *Journal of Experimental Biology*, vol. 213, no. 2, pp. 331–338, 2010.

[39] M. E. Chiappe, J. D. Seelig, M. B. Reiser, and V. Jayaraman, "Walking modulates speed sensitivity in drosophila motion vision," *Current Biology*, vol. 20, no. 16, pp. 1470–1475, 2010.

[40] C. M. Niell and M. P. Stryker, "Modulation of visual responses by behavioral state in mouse visual cortex," *Neuron*, vol. 65, no. 4, pp. 472–479, 2010.

[41] Y. Yang, D. G. Yared, and N. J. Cowan, "Can a fish learn to ride a bicycle? sensorimotor adaptation to destabilizing dynamics in weakly electric fish," *bioRxiv*, pp. 2023–01, 2023.