Dendritic trafficking faces physiologically critical speed-precision tradeoffs

Alex H. Williams[1,2,3,*], **Cian O'Donnell**[2,4], **Terrence Sejnowski**[2,5], **and Timothy O'Leary**[6,7,*]

[1]Department of Neurosciences, University of California, San Diego, La Jolla, CA 92093, USA
[2]Howard Hughes Medical Institute, Salk Institute for Biological Studies, La Jolla, CA 92037, USA
[3]Department of Neurobiology, Stanford University, Stanford, CA 94305, USA
[4]Department of Computer Science, University of Bristol, Woodland Road, Bristol, BS8 1UB, United Kingdom
[5]Division of Biological Sciences, University of California at San Diego, La Jolla, CA 92093, USA
[6]Volen Center and Biology Department, Brandeis University, Waltham, MA 02454, USA
[7]Department of Engineering, University of Cambridge, Trumpington St, Cambridge, CB2 1PZ, United Kingdom
[*]Address correspondence to: ahwillia@stanford.edu, timothy.oleary@eng.cam.ac.uk

ABSTRACT

Neuronal function depends on the appropriate subcellular distribution of channels, receptors, mRNAs, and other components.Trafficking of subcellular cargo to specific locations is regulated by local signals such as synaptic input, yet it remains unclear how such a decentralized system performs in complex morphologies. We mathematically formalize a popular, conceptual model of microtubule transport (the "sushi-belt model", Doyle and Kiebler, 2011), and show that it can achieve arbitrarily complex spatial distributions of cargo in realistic morphologies. However, we reveal that this model predicts a critical tradeoff between the speed and precision of cargo delivery; given experimental estimates of trafficking kinetics the model predicts delays of many hours for modestly accurate cargo distribution. We explore the possibility that biological mechanisms might have more sophisticated, globally fine-tuned trafficking kinetics to enable both fast and precise transport; however, tuned mechanisms are less flexible, and are fragile to changes in spatial demand for cargo.

Keywords: Regulation, Active transport, Plasticity, Tagging hypothesis, Morphology, Motor proteins

INTRODUCTION

The axonal and dendritic trees of most neurons are remarkably complex. The logistical task of distributing biomolecular components within neuronal morphologies is therefore considerable, especially for components that are synthesized in locations that are distal from their site of use. Many neurophysiological processes, such as synaptic plasticity, neurite development and homeostatic regulation depend on dendritic trafficking, so characterizing the capabilities and limitations intracellular transport is crucial to our understanding of higher-order nervous system function. Although the molecular machinery responsible for active transport has been studied in detail (Doyle and Kiebler, 2011; Buxbaum et al., 2014a; Hancock, 2014), there have been few attempts to model how these mechanisms can be orchestrated to distribute cargo accurately and efficiently in realistic neuron morphologies, or the fundamental constraints that trafficking mechanisms may face.

Not all aspects of neurite metabolism depend on continual delivery of cargo from the soma. For example, some

forms of long-term potentiation (LTP) use local protein synthesis and can function in isolated dendrites (Kang and Schuman, 1996; Aakalu et al., 2001; Vickers et al., 2005; Sutton and Schuman, 2006; Holt and Schuman, 2013). Similarly, constitutive maintenance of cytoskeletal, membrane and signaling pathways is achieved in part by locally recycling or synthesizing components (Park et al., 2004, 2006; Grant and Donaldson, 2009; Zheng et al., 2015). Nonetheless, components originating in the soma need to find their way to synapses and dendritic branches in the first place, and in the long run, the mRNAs and the machinery that supports local biosynthesis need to be replenished. More crucially, many long-lasting forms of synaptic plasticity are known to depend on anterograde transport of mRNAs (Nguyen et al., 1994; Bading, 2000; Kandel, 2001) and specific mRNAs are known to be selectively transported to regions of heightened synaptic activity (Steward et al., 1998; Steward and Worley, 2001; Moga et al., 2004) or to developing synaptic contacts (Lyles et al., 2006). These observations fit with the well-known synaptic tagging hypothesis (Frey and Morris, 1997, 1998; Redondo and Morris, 2011), which proposes that synapses produce a biochemical tag that signals a requirement for synaptic building blocks as part of the plasticity process.

These neurophysiological functions rely on a number of second-messenger pathways that can also regulate the transport and recruitment of intracellular cargo (Doyle and Kiebler, 2011). Kinesin and dyenin motor proteins transport cargo along microtubules in a stochastic and bidirectional manner (Block et al., 1990; Smith and Simmons, 2001; Hirokawa et al., 2010; Gagnon and Mowry, 2011; Park et al., 2014). The stochastic nature of single-particle movements has led to the hypothesis that cargo are predominantly controlled by noisy, localized signaling pathways, rather than a centralized addressing system (Welte, 2004; Bressloff and Newby, 2009; Newby and Bressloff, 2010; Doyle and Kiebler, 2011; Buxbaum et al., 2015). These localized mechanisms are not fully understood, but are believed to involve transient elevations in second-messengers like $[Ca^{2+}]$ and ADP (Mironov, 2007; Wang and Schwarz, 2009), glutamate receptor activation (Kao et al., 2010; Buxbaum et al., 2014a), and changes in microtubule-associated proteins (Soundararajan and Bullock, 2014).

Based on these observations, several reviews have advanced a conceptual model in which cargo searches the dendritic arbor via a noisy walk before its eventual release or capture (Welte, 2004; Buxbaum et al., 2014a, 2015). Doyle and Kiebler (2011) refer to this as the "sushi belt model". In this analogy, molecular cargoes are represented by sushi plates that are distributed on a conveyor belt, as found in certain restaurants. Customers sitting alongside the belt (representing specific locations or synapses along a dendrite) each have specific and potentially time-critical demands for the amount and type of sushi they consume, but they can only choose from nearby plates as they pass.

Stated in words, the sushi belt model is intuitively plausible, yet a number of open questions remain. For example, how can a trafficking system based on local signals be tuned to accurately generate spatial distributions of cargo? Does the model predict cross-talk, or interference between spatially separated regions of the neuron that require the same kind of cargo? Within this family of models, are there multiple sets of trafficking parameters capable of producing the same distribution of cargo, and do they use qualitatively different strategies to produce it? Finally, how quickly and how accurately can cargo be delivered by this class of models, and do these measures of performance depend on morphology and the specific spatial pattern of demand?

Each of these high-level questions can be addressed in a parsimonious mathematical framework that is broadly applicable to many biological contexts and does not rely on stringent parameter estimates. Here, we present a simple, plausible formulation of the sushi belt model that reveals unanticipated features of intracellular trafficking. We show that this model can reliably produce complex spatial distributions of cargo across dendritic trees using local signals. However, we also reveal strong logistical constraints on this family of models. In particular, we show that

fast cargo delivery is error-prone, resulting in cargo distributions that do not match demand. Conversely, modestly precise cargo delivery requires many hours based on conservative estimates of trafficking kinetics. We present novel models that circumvent this speed-precision tradeoff. However, these possibilities pay a price of greater complexity and sensitivity to perturbations. Overall, these predictions suggest new experiments to probe the limitations and vulnerabilities of intracellular trafficking, and determine if more complex models beyond the sushi belt are required.

RESULTS

Model description

Transport along microtubules is mediated by kinesin and dynein motors, which are responsible for anterograde and retrograde transport, respectively (Hirokawa et al., 2010; Gagnon and Mowry, 2011). Cargo is often simultaneously bound to both forms of motor protein, resulting in stochastic back-and-forth movements with a net direction determined by the balance of opposing movements (Welte, 2004; Hancock, 2014; Buxbaum et al., 2014a, Fig. 1A).

To obtain a general model that can accommodate variations in the biophysical details, we interpreted microtubule-based transport as a biased random walk along a one-dimensional cable, corresponding to a section of neurite (Bressloff, 2006; Bressloff and Newby, 2009; Newby and Bressloff, 2010; Bressloff and Newby, 2013). For each time step (1 s), the cargo moves 1 μm forwards, 1 μm backwards, or remains in the same place, each with different probabilities. These parameters produce a qualitative fit to a more detailed biophysical model (Müller et al., 2008). In the simplest model, the probabilities associated with each movement are fixed and independent across each time step, with the forward jump more probable than a reverse jump, leading to a biased random walk (Fig. 1B, top panel). Existing biophysical models of single-cargo transport include mechanisms for extended unidirectional runs (Müller et al., 2008; Hancock, 2014). To account for these runs, we made a second version of the model in which the movement probabilities at each time step depend on the previous position of the particle (see *Methods*); the resulting trajectories change direction less frequently, resulting in a larger variance in cargo position over time (Fig. 1B, bottom panel).

While the position of individual cargoes is highly stochastic and dependent on the specific sequence of microscopic steps, the net movement of a population of cargoes (Fig. 1C) is predictable. This is seen in Figure 1D, which shows the distribution of 1000 molecules over time with (top panel) and without (bottom panel) unidirectional runs. Thus, bulk trafficking of cargo can be modeled as a deterministic process, which we refer to as the "mass-action model" of transport. The model discretizes the dendritic tree into small compartments, and describes the transfer of cargo between neighboring compartments as reactions with first order kinetics. In a neurite with N compartments, the mass-action model is (Fig. 1E):

$$u_1 \underset{b_1}{\overset{a_1}{\rightleftharpoons}} u_2 \underset{b_2}{\overset{a_2}{\rightleftharpoons}} u_3 \underset{b_3}{\overset{a_3}{\rightleftharpoons}} \ \ldots \ \underset{b_{N-1}}{\overset{a_{N-1}}{\rightleftharpoons}} u_N \tag{1}$$

where u_i is the amount of cargo in each compartment, and a_i and b_i respectively denote local rate constants for anterograde and retrograde transport. For now, we assume that these rate constants are spatially homogeneous, since the stochastic movements of the single-particle simulations do not depend on position. In this case, the mass-action model maps onto the well-known drift-diffusion equation (Fig. 1E Smith and Simmons, 2001) that can be used to estimate plausible parameter ranges from experimental data. The drift and diffusion coefficients are respectively proportional to the rate of change of the mean and variance of the ensemble distribution (see *Methods*); thus, an experimental measurement of these two quantities on a specific stretch of dendrite would provide a local estimate of

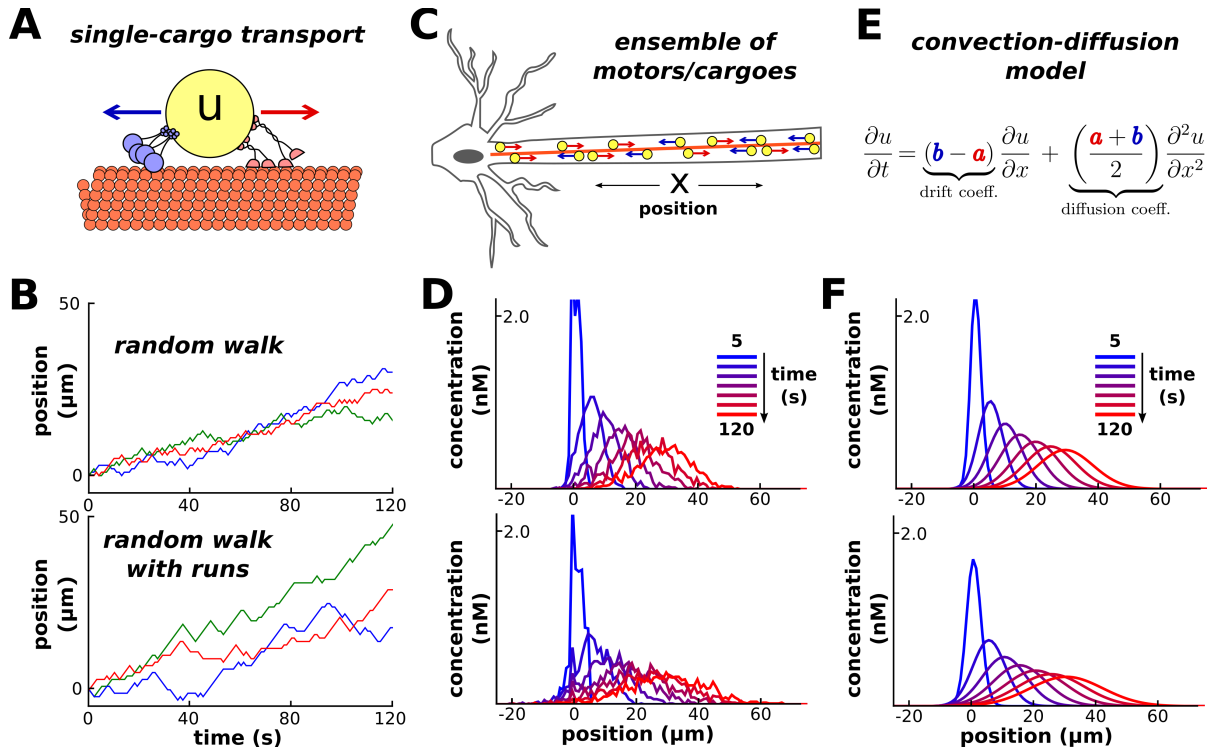

Figure 1. Model of intracellular transport. **(A)** Molecular cargo, u, on a microtubule undergoes stochastic back-and-forth movements driven by opposing motor proteins. **(B)** Three example random walks representing the movement of individual cargoes. In a simple random walk, each movement is independent of previous movements (top panel); longer unidirectional run result from adding history-dependence to the model, such that each movement is likely to continue in the same direction as the previous step (bottom panel). **(C)** An population of synaptic cargoes transported along the length of a neurite. **(D)** The concentration profile of a population of transported cargo along a neurite over time, simulated as 1,000 independent random walks. Simulations with (bottom) and without (top) history-dependence as in panel B. **(E)** In the limit of many individual cargo particles, the concentration of u is described by a mass-action model (equation 1). The parameters a and b respectively scale the anterograde and retrograde rate of transport. **(F)** The mass-action model provides a good fit for the simulations in panel D.

the trafficking parameters.

Figure 1F shows two mass-action models with 1 μm compartments. The trafficking rate constants were chosen to reproduce the mean and variance of the ensemble simulations in figure 1D. This resulted in $a \approx 0.42$ s^{-1} and $b \approx 0.17$ s^{-1} for the simulation without runs (top panel) and $a \approx 0.79$ s^{-1} and $b \approx 0.54$ s^{-1} with runs (bottom panel). These parameters produce mean particle velocities of 15 μm per minute for both simulations, which is within the range of experimental observations for microtubule transport (Rogers and Gelfand, 1998; Dynes and Steward, 2007; Müller et al., 2008). The variances of the particle distributions grow at a rate of 35 and 80 μm^2 per minute for the top and bottom panel, respectively.

The mass-action model is based on two important assumptions. First, we assumed that the net movement of molecules between neighboring compartments is a stochastic and memoryless process. This assumption is reasonable for cargoes that change direction often; specifically, on a length scale comparable to the size of the compartments. This appears to be the case under many (Müller et al., 2008; Verbrugge et al., 2009), though not necessarily all (Dynes and Steward, 2007; Soundararajan and Bullock, 2014) circumstances. However, we observed

that the mass-action model could still be well-fit to simulated data where this assumption was violated (Fig. 1F, lower panel). In fact, the mass-action model only breaks down in the limit of very large run lengths (i.e. unidirectional, non-stochastic transport; see *Methods*).

The second assumption of the mass-action model is that there are many transported particles distributed throughout the dendritic tree. Many types of dendritic cargo are present in high numbers (Cajigas et al., 2012), and a deterministic model can provide a good approximation in this regime (Fig. 1D). The model also provides insight into the stochastic dynamics of transport for cargoes with fewer copy numbers. Instead of interpreting u as the amount of cargo in each compartment, this variable (when appropriately normalized) can be interpreted as the probability that a single particle lies within a particular compartment. Thus, for a small number of transported cargoes, the mass-action model describes the average, or expected, distribution of the ensemble, with noise inversely proportional to the square root of the number of copies of cargo (see *Methods*).

A simple transport mechanism distributes cargo according to demand

We used the mass-action model for simulations because it is easy to extend it to branched morphologies (Fig. 2A), and to cases where the trafficking rate constants (a_i and b_i) are spatially non-uniform. These features are critical to model realistic biological settings where the kinetics of motor proteins depend on local biochemical signals (Mironov, 2007; Wang and Schwarz, 2009; Soundararajan and Bullock, 2014).

In the model, the exchange of cargo between compartments approaches a steady-state (*ss*), distribution over time (see *Methods*). The steady-state occurs when no net movement of cargo occurs between connected compartments. Mathematically, this occurs when:

$$\left.\frac{u_p}{u_c}\right|_{ss} = \frac{b}{a} \tag{2}$$

where u_p is the level in the "parent" compartment (closer to soma) and u_c is the level in the "child" compartment (closer to periphery); b and a refer to the trafficking rate constants between this pair of compartments.

Intuitively, in the mass-action model, the rate of cargo transfer is proportional to the concentration of cargo and the trafficking rate constant. Thus, the flow of cargo is equal and opposite when the ratio of cargo matches the reciprocal ratio of the rate constants between a pair of connected compartments.

The above result means that a specific spatial distribution of cargo can be achieved at steady state by modulating local trafficking rates. We denote the steady-state level of cargo in each compartment by $\tilde{u}$, which can be encoded by a localized biochemical signal as we will soon show. If we interpret the steady state as being a specific "target concentration" (determined by local demand), we see that the ratio of local trafficking rates must satisfy:

$$\frac{b}{a} = \frac{\tilde{u}_p}{\tilde{u}_c} \tag{3}$$

For example, we produced a linear expression gradient (Magee, 1998; Hoffman et al., 1997) by setting $\tilde{u}$ directly proportional to the distance from the soma. The trafficking rate constants in this case satisfy $b_i/a_i = i/i+1$ (where i indexes on increasing distance to the soma). Figure 2B-C shows that this rule produced the expected profile, revealing that the slope of the linear gradient is controlled by tuning the total amount of cargo in the neurite (Fig. 2C).

Together, this analysis shows that by manipulating local trafficking rates, arbitrary distributions of cargo can be achieved over time. We next incorporated a local signal to encode the demand for cargo at specific locations. The

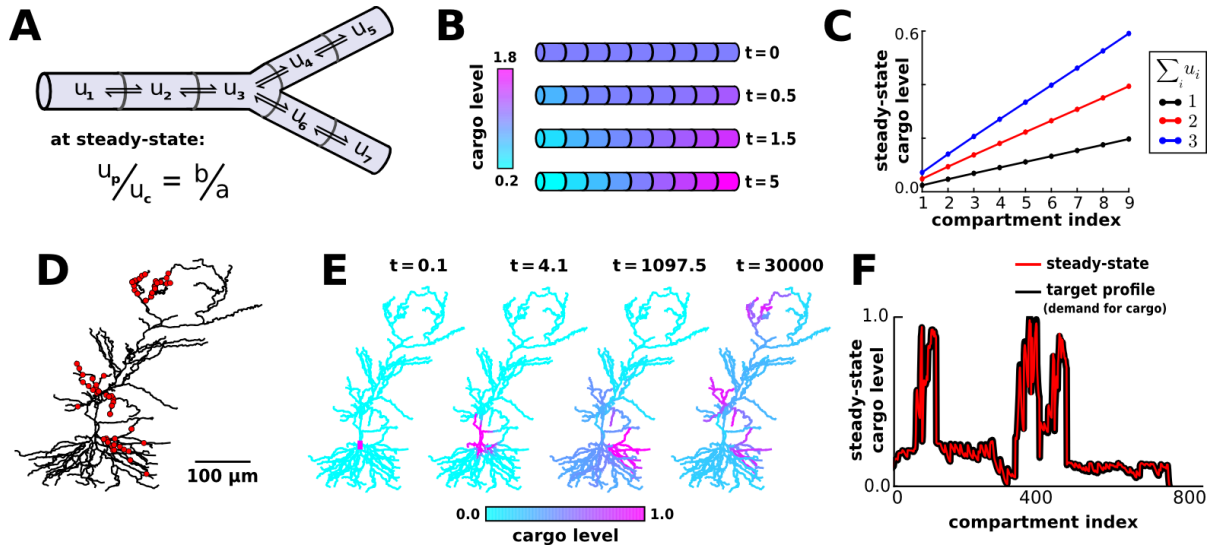

Figure 2. Local trafficking rates determine the spatial distribution of biomolecules by a simple kinetic relationship. **(A)** The mass action transport model for a simple branched morphology. **(B)** A simulation of a nine-compartment cable, with trafficking rate constants set to produce a linear gradient using the steady-state relation shown in panel A. **(C)** The slope of the linear gradient shown in panel B is directly proportional to the total amount of cargo in the model ($\sum_i u_i$); the slope increases with increasing cargo levels, but the linear profile is preserved. **(D)** A model of a CA1 pyramidal cell with 742 compartments adapted from Migliore and Migliore (2012); excitatory synapses were added at the locations marked by red dots. **(E)** The average membrane potential in each compartment of the CA1 model determined $\tilde{u}$, which was used to determine values for the trafficking rate constants by equation (3). Over time, the spatial distribution of cargo evolved to match local demand. **(F)** The steady-state profile of cargo for each compartment in the simulation shown in panel E is precisely matched to local demand.

molecular identity of this signal is not crucial to the behavior of this model, though we propose cytosolic calcium, $[Ca]_i$ since it is modulated by local synaptic activity and increases in $[Ca]_i$ are known to simultaneously arrest anterograde and retrograde microtubular transport for certain cargoes (Wang and Schwarz, 2009). We model this by assuming that the anterograde rate constant is determined by the calcium concentration in the parent compartment, $a = f([Ca]_p)$, and the retrograde rate constant is determined by calcium concentration in the child compartment, $b = f([Ca]_c)$, where f is a function that describes how calcium concentration alters the transport rates. This results in:

$$\frac{b}{a} = \frac{f([Ca]_c)}{f([Ca]_p)} = \frac{\tilde{u}_p}{\tilde{u}_c} \tag{4}$$

where $\tilde{u}_i = 1/f([Ca]_i)$.

Thus, in principle, local calcium transients that arrest transport could provide a mechanism for distributing cargo to an arbitrary target profile. We emphasize that other potential signaling pathways could achieve the same effect, so while there is direct evidence for $[Ca]_i$ as an important signal, the model can be interpreted broadly, with $[Ca]_i$ serving as a placeholder for any local signal identified experimentally. To test the global behavior of this model in a complex morphology, we implemented the trafficking model in an existing morphologically reconstructed model of a CA1 pyramidal neuron (Migliore and Migliore, 2012). Excitatory synaptic input was delivered to 120 locations within three dendritic regions (red dots, Fig. 2D), and the average membrane potential in each electrical

compartment determined the target level ($\tilde{u}_i$) in each compartment (Methods). This models how molecular cargo could be selectively trafficked to active synaptic sites (Fig. 2E, Supp. Video 1). Figure 2F confirms that the spatial distribution of u precisely approaches the desired steady-state exactly in the noise-free case.

Convergence rate

Many neurophysiological processes are time-sensitive. For example, newly synthesized proteins must be delivered to synapses within $\sim$1 hour to support long-term potentiation in CA1 pyramidal cells (Frey and Morris, 1997, 1998; Redondo and Morris, 2011). More generally, biochemical components whose abundance and localization are regulated by cellular feedback signals need to be distributed within a time interval that keeps pace with demand, within reasonable bounds. We therefore examined how quickly the spatial pattern of cargo converged to its target distribution.

In equations (3) and (4), we implicitly required each $\tilde{u}_i > 0$ in order to avoid division by zero. Intuitively, if $\tilde{u}_i = 0$ for some compartment, then no cargo can flow through that compartment, cutting of more peripheral compartments from the transport system. Similarly, if certain $\tilde{u}_i$ are nearly equal to zero, then transport through these compartments will act as a bottleneck for transport, and convergence to the desired distribution will be slow.

Figure 3A-C illustrates and analyzes a bottleneck in a simple three compartment model. The two compartments on the end of the cable have the same desired level, $\tilde{u}_1 = \tilde{u}_3$; the middle compartment acts as a bottleneck when $\tilde{u}_2$ is very small (Fig. 3A). We can achieve this distribution with a simple mass-action model with only one free parameter, ε, which is the rate constant of trafficking into the middle compartment from either end; we fix the rate constant of trafficking out of the middle compartment to be 1 without loss of generality:

$$u_1 \underset{1}{\overset{\varepsilon}{\rightleftharpoons}} u_2 \underset{\varepsilon}{\overset{1}{\rightleftharpoons}} u_3 \tag{5}$$

We assume that u begins at one end of the cable, and examine the how taking ε to zero affects the convergence rate. Figure 3B shows that the convergence rate slows dramatically as ε decreases. Intuitively, ε can be thought of as the rate-limiting step for the system. More precisely, the rate of convergence is given by the smallest magnitude, non-zero eigenvalue of the system, which can be readily calculated (see *Methods*). In the model of Figure 3B, the smallest magnitude eigenvalue is $-\varepsilon$ leading to a convergence time constant of $1/\varepsilon$. Simulations confirmed this analysis and showed that the convergence time diverges to infinity as ε approaches zero (Fig. 3C).

We then asked whether the intuition from this three-compartment model extended to a cell with realistic morphology. We obtained qualitatively similar results. The CA1 model converged to a uniform target distribution more quickly than to a "bottleneck" target distribution, in which the middle third of the apical dendrite had low steady-state levels of cargo (Fig. 3D). Each pair of anterograde and retrograde rate constants was normalized to sum to one so that the differences in convergence were not due to the scale of the trafficking rate constants.

In addition to this global view of convergence (Fig. 3D), we considered how the transport bottleneck affected transport to individual dendritic compartments. Consider a scenario where transported cargo produces a local chemical reaction after a certain quantity of cargo accumulates; for example, a recently potentiated synapse might be stabilized after a threshold amount of plasticity-related factors is surpassed. Figure 3E plots the duration of time it took for u_i to reach a pre-defined threshold for each compartment as a measure of the local transport delay. As expected, introducing a bottleneck caused much longer delays to compartments distal to that bottleneck (Fig. 3E, upper right portion of plot). Interestingly, the presence of a bottleneck *shortens* the transport delay to proximal

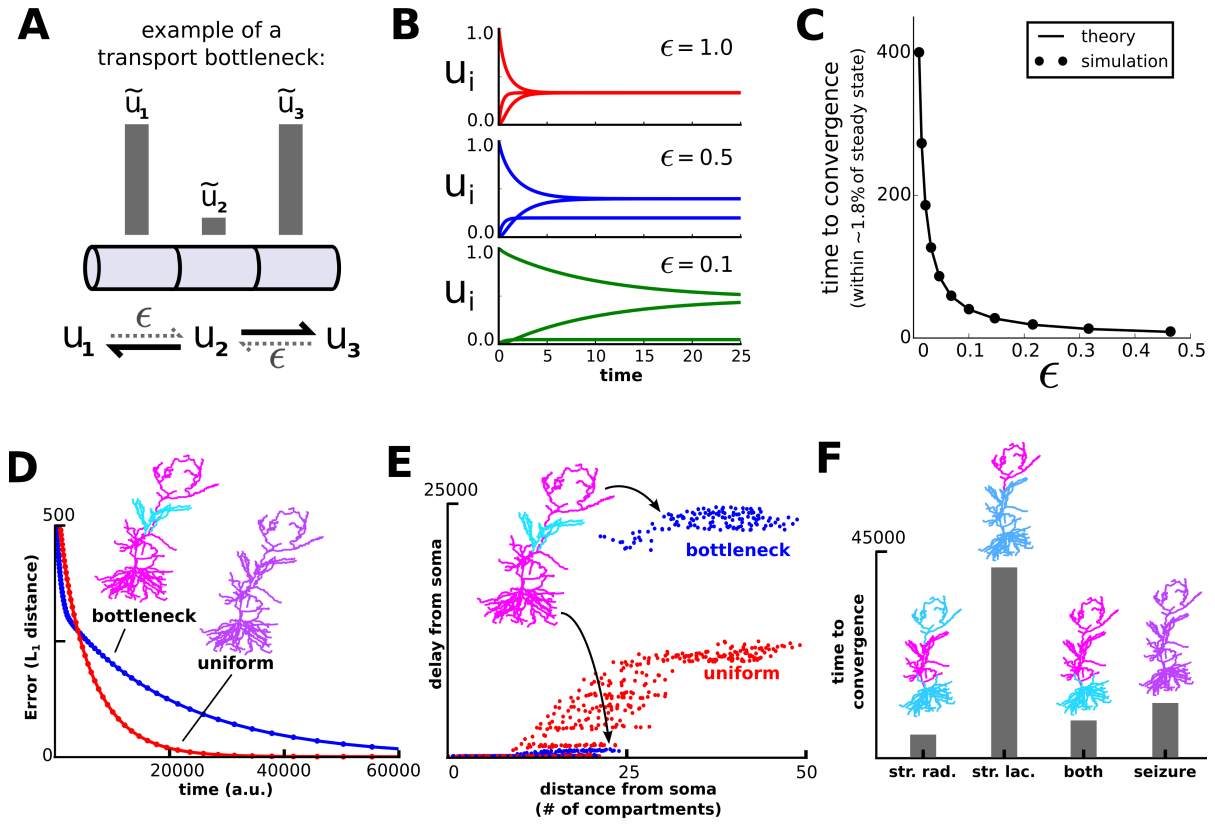

Figure 3. Transport bottlenecks caused by local demand profile. **(A)** A three-compartment transport model, with the middle compartment generating a bottleneck. The vertical bars represent the desired steady-state level of cargo in each compartment. The rate of transport into the middle compartment is small (ε, dashed arrows) relative to transport out of the middle compartment. **(B)** As ε decreases, the model converges more slowly and the steady-state level decreases in the middle compartment. **(C)** Simulations (black dots) confirm that the time to convergence is given by the smallest non-zero eigenvalue of the system (analytically calculated line). This eigenvalue can be thought of as the rate-limiting step of the system. **(D)** Convergence (L_1 norm or the sum of absolute value of the residuals) to a uniform target distribution (red line) is faster than a target distribution containing a bottleneck (blue line) in the CA1 model. **(E)** For all compartments that reach a threshold level ($u_i = 0.001$), the simulated time it takes to reach threshold is plotted against the distance of that compartment to the soma. **(F)** Convergence times for various target distributions (*str. rad.*, stratum radiatum; *str. lac.*, stratum lacunosum/moleculare) in the CA1 model. The timescale of all simulations in the CA1 model was set by imposing the constraint that $a_i + b_i = 1$ for each compartment.

compartments, compared to the uniform target distribution (Fig. 3E, lower left portion of plot). This occurs because cargo delayed by the bottleneck spreads throughout the proximal compartments, reaching higher levels earlier in the simulation. We observed qualitatively similar results for different local threshold values (data not shown).

The model predicts that transport to distal compartments will converge to steady-state at a faster rate when the steady-state level of cargo in the proximal compartments increases, since this removes a trafficking bottleneck between the soma and distal compartments. This can be experimentally tested by characterizing the convergence of a macromolecule that aggregates at recently activated synapses, such as *Arc* mRNA (Steward et al., 1998, see *Discussion*). To illustrate this in the model CA1 cell, we characterized the time course of transport to the distal

apical dendrite (stimulating stratum lacunosum/moleculare), proximal apical dendrite (stimulating stratum radiatum), entire apical dendrite (stimulating both layers), and entire cell (seizure). Notably, the model converges more slowly to distal input alone, than to paired distal and proximal input, or to an entirely uniform input distribution (Fig. 3F).

Microtubular trafficking, detachment and degradation on separated time scales

We have so far considered how a target distribution of cargo can be generated by modulating the local rates of molecular motors. However, while certain types of molecular cargo stay on the microtubule network (e.g. trafficked mitochondria), many kinds of cargo must detach from the microtubules in order to be used at their final location. For example, dendritic mRNAs are transported along microtubules within densely-packed granules, and are released following granule disassembly (Krichevsky and Kosik, 2001; Buxbaum et al., 2014b). In this case, and in other cases where the cargo is sequestered in transit, it is reasonable to model detachment from the microtubule as an irreversible process, followed by an eventual degradation.

This conceptual picture has been called the sushi belt model of transport (Doyle and Kiebler, 2011). This idea can be formalized as the following mass-action scheme:

$$\begin{array}{ccccccccc} u_1 & \underset{b_1}{\overset{a_1}{\rightleftharpoons}} & u_2 & \underset{b_2}{\overset{a_2}{\rightleftharpoons}} & u_3 & \underset{b_3}{\overset{a_3}{\rightleftharpoons}} & u_4 & \underset{b_4}{\overset{a_4}{\rightleftharpoons}} & \dots \\ \downarrow{\scriptstyle c_1} & & \downarrow{\scriptstyle c_2} & & \downarrow{\scriptstyle c_3} & & \downarrow{\scriptstyle c_4} & & \\ u_1^\star & & u_2^\star & & u_3^\star & & u_4^\star & & \\ \downarrow{\scriptstyle d_1} & & \downarrow{\scriptstyle d_2} & & \downarrow{\scriptstyle d_3} & & \downarrow{\scriptstyle d_4} & & \\ \varnothing & & \varnothing & & \varnothing & & \varnothing & & \end{array} \tag{6}$$

As before, a molecule u is transported along a network of microtubules (top row, in equation 6). In each compartment, molecules can irreversibly detach from the microtubules in a reaction $u_i \xrightarrow{c_i} u_i^\star$, where $u^\star$ denotes the biochemically active or released form of cargo. The final reaction, $u_i^\star \xrightarrow{d_i} \varnothing$, models degradation in each compartment. Note that only $u^\star$ is subject to degradation; the molecule is assumed to be protected from degradation during transport. This model can be extended to branched morphologies (Fig. 3A).

To analyze this system we first assume that these three processes — trafficking, detachment, and degradation — occur on separated timescales. If trafficking is sufficiently faster than detachment ($a, b \gg c$), then u approaches a quasi-steady state distribution defined by the previous analysis (equation 3). We then choose detachment rate constants that transform the microtubular distribution into the desired distribution for $u^\star$:

$$c_i \propto \frac{\tilde{u}_i^\star}{\tilde{u}_i} \tag{7}$$

Here, $\tilde{u}$ and $\tilde{u}^\star$ respectively denote the quasi-steady state distributions for u and $u^\star$, respectively. As long as degradation is sufficiently slow ($c \gg d$) the desired distribution is transiently achieved. For simplicity, we set $d_i = 0$ in simulations.

The addition of spatially varied detachment rates (c_i) produces a spectrum of strategies for achieving a desired distribution of cargo (Fig. 4B). We can reproduce the essential strategy used in figure 2 by choosing the transport rates (a_i,b_i) to match the target distribution of cargo during the trafficking phase of transport (using equation 3).

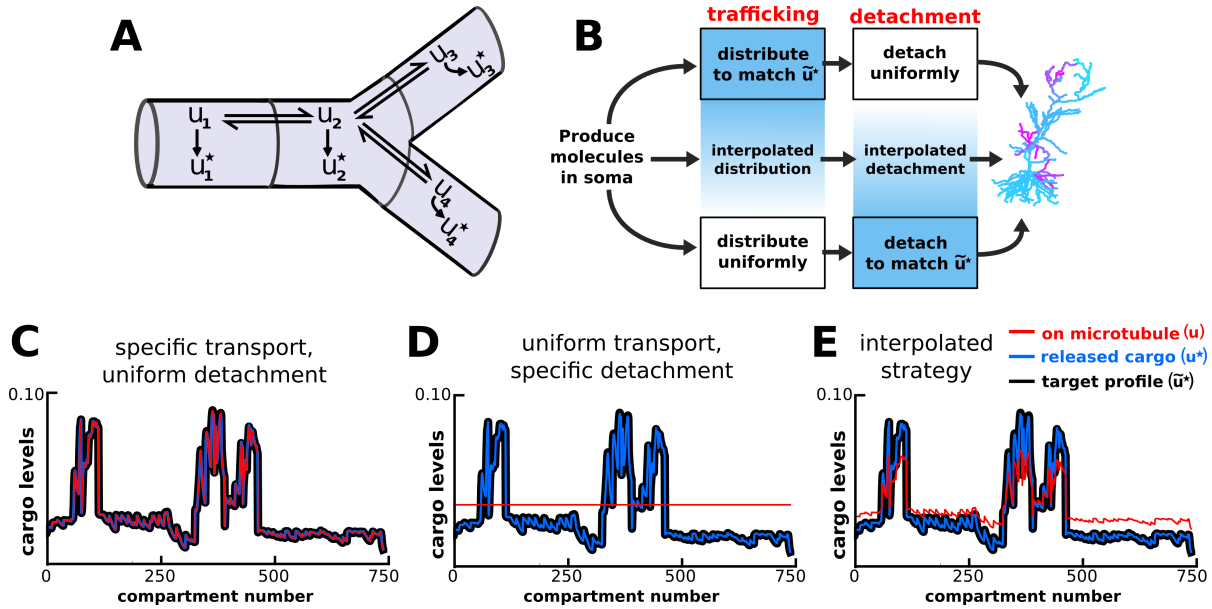

Figure 4. Multiple strategies for transport in a model including nonuniform microtubular detachment/activation. **(A)** Schematic of microtubular transport model with irreversible detachment in a branched morphology. The localized degradation reactions ($u_i^\star \xrightarrow{d_i} \varnothing$) are omitted for clarity. **(B)** Multiple strategies for producing a desired distribution of detached cargo ($\tilde{u}^\star$). When the timescale of detachment/delivery is sufficiently slow, the distribution of cargo on the microtubules approaches a quasi-steady-state (transport step). This known distribution can then be transformed into the desired distribution for $\tilde{u}$ (detachment step). As long as these two steps are appropriately matched (blends of blue), then the desired distribution will be achieved (CA1 cell, right). **(C-E)** Quasi-steady-state distributions of u (red), $u^\star$ (blue), and $\tilde{u}^\star$ (black) for the various strategies diagrammed in panel B.

Then, since the distribution is achieved at quasi-steady-state, the cargo can be detached uniformly ($c_i =$ constant, with $c \ll a, b$). Figure 4C shows simulated data confirming that the desired spatial distribution is first achieved along the microtubules (red line, Fig. 4C) and then maintained after cargo is uniformly detached (blue line, Fig. 4C).

A second strategy begins by choosing uniform transport rates ($a_i = b_i$), which evenly distributes cargo throughout the dendritic tree. The desired distribution is then achieved by locally delivering cargo at a rate proportional to the desired level ($c_i \propto \tilde{u}_i^\star$; Fig 4D). Unlike the first solution, this strategy avoids the transport bottlenecks examined in Figure 3, and can achieve target patterns with $\tilde{u}^\star$ equal to zero in certain compartments by setting $c_i = 0$.

We refer to the first strategy (Fig. 4C) as the *specific trafficking model*, because cargo is selectively transported to target sites. We refer to the second strategy (Fig. 4D) as the *uniform trafficking model*, because cargo is uniformly distributed throughout the dendrite. These two strategies represent extremes on a spectrum of possible models (Fig. 4B). Figure 4E shows the behavior of an intermediate model, whose parameters are a linear interpolation between the extreme strategies shown in Figure 4C and 4D. Thus, effective trafficking systems can be described as a spectrum of strategies that may be suited to different situations and purposes (see Discussion).

Non-specific cargo delivery occurs when trafficking and detachment occur on similar timescales

We have demonstrated possible strategies for neurons to achieve precise and flexible transport of cargo by assuming that cargo detachment is sufficiently slow relative to trafficking. However, biological neurons are unlikely to require perfect matching between demand and distribution of cargo and may therefore tolerate loss of precision in order

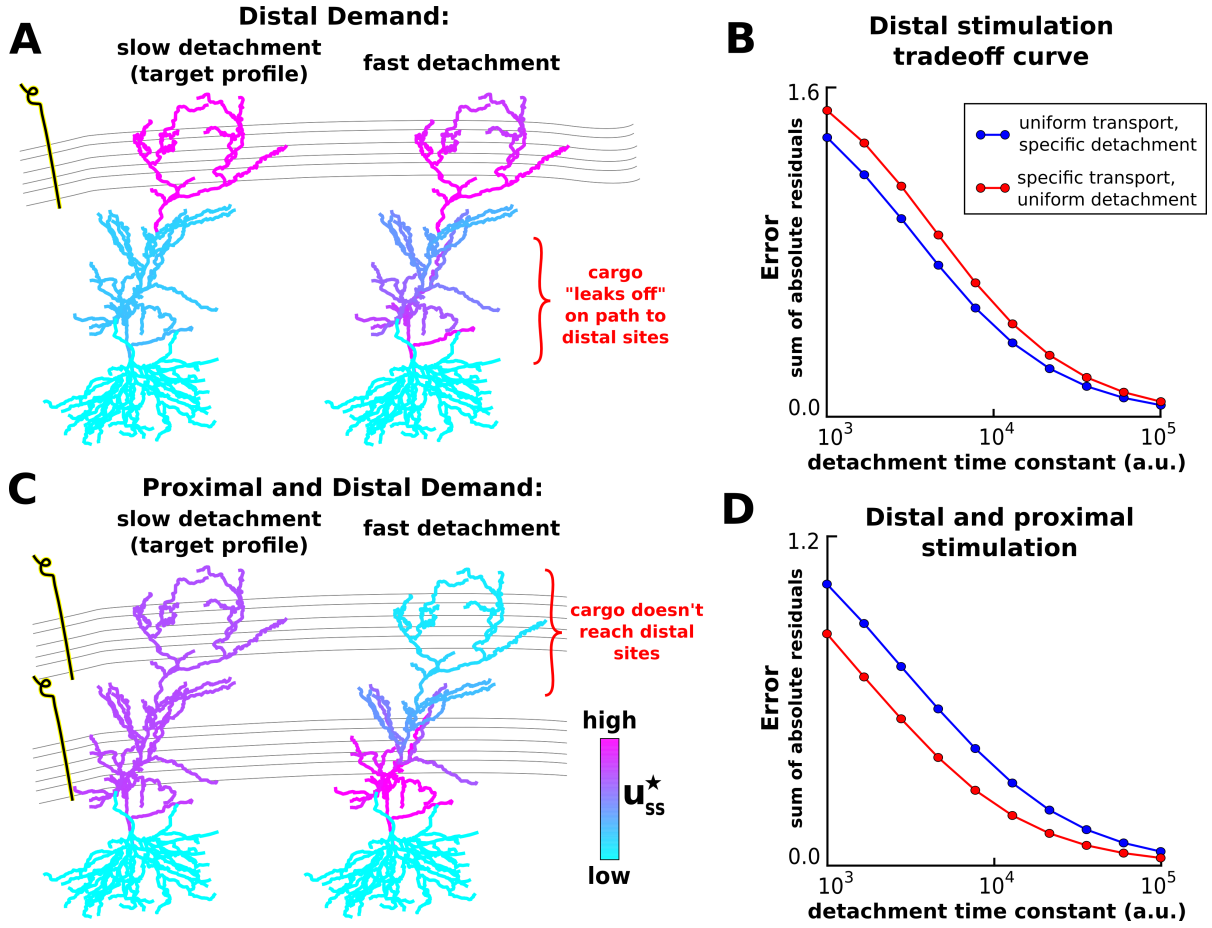

Figure 5. Proximal synapses capture more cargo at the expense of distal synapses when detachment rates are naïvely increased. **(A)** Delivery of cargo to the distal apical zone with slow *(left)* and fast detachment rates *(right)*. The achieved pattern does not match the target distribution when detachment is fast, since some cargo is erroneously delivered to proximal sites. **(B)** Tradeoff curves between non-specificity and convergence rate for two trafficking strategies (blue line, see Fig 4D; red line, see Fig 4C). **(C-D)** Same as **(A-B)** but for intended cargo delivery to the entire apical tuft. The timescale of all simulations was set by imposing the constraint that $a_i + b_i = 1$ for each compartment.

to speed up delivery. We therefore examined the consequences of relaxing the separation of time scales between transport and detachment.

Returning to the model of Figure 4, we consider a scenario where distal synaptic inputs on the apical tuft are stimulated, triggering local demand (Fig. 5A). If the detachment rate constants are sufficiently slow, then, as before, cargo is delivered selectively to the stimulated region (Fig. 5A, left). If we increase the detachment rates by a uniformly scaling, some cargo "leaks off" the microtubule path on its way to the distal synapses (Fig. 5A, right). We refer to this as a "non-specific" delivery of cargo, since cargo is not selectively delivered to the stimulated sites. Thus, speeding up detachment relative to transport improves the overall rate at which cargoes are delivered to synapses, but this comes at the cost of decreased accuracy of delivery.

Tradeoff curves between the average detachment rate constant and the non-specificity of transport for this stimulation pattern are shown in figure 5B. Importantly, this tradeoff exists for both trafficking strategies we

examined in figure 4 — the selective transport strategy (see Fig 4D) and uniform transport strategy (see Fig. 4C). For this stimulation pattern, the uniform trafficking strategy (Fig. 5B, blue line) outperforms the specific trafficking strategy (Fig. 5B, red line) since the latter suffers from bottleneck in the proximal apical zone.

The pattern of non-specific delivery is stimulation-dependent. When the entire apical tree is stimulated, fast detachment can result in a complete occlusion of cargo delivery to distal synaptic sites (Fig. 5C). As before, a tradeoff between specificity and delivery speed is present for both transport strategies (Fig. 5D). Interestingly, the specific trafficking strategy outperforms the uniform trafficking strategy in this case (in contrast to Fig. 5A-B). This is due to the lack of a bottleneck, and the fact that the uniform trafficking strategy initially sends cargo to the basal dendrites where it is not released.

Together, these results show that an increase in the efficiency of synaptic cargo delivery comes at the cost of specificity and that the final destination of mis-trafficked cargo depends on the pattern of stimulation.

Transport speed and precision can be optimized for specific spatial patterns of demand

In figure 5, we showed that scaling the detachment rates (c_i), while leaving the transport rates (a_i, b_i) fixed produces a proximal bias in cargo delivery. We reasoned that increasing the anterograde transport rate of cargo could correct for this bias, producing transport rules that are both fast and precise.

We examined this possibility in a reduced model — an unbranched cable — so that we could develop simple heuristics that precisely achieve a desired distribution of cargo while minimizing the convergence time. In this model, cargo begins on the left end of the cable and is transported to a number of synaptic delivery sites, each of which is modeled as a double-exponential curve. We also restricted ourselves to investigating the uniform trafficking strategy (Fig. 4D); a similar analysis could in principle be done for the specific trafficking strategy.

As before, cargo can be precisely delivered to a variety of stimulation patterns when the detachment rate is sufficiently slow (Fig. 6A, Supp. Video 2); when the detachment rate is naïvely increased to speed up the rate of delivery, a proximal bias develops for all stimulation patterns (Fig. 6B, Supp. Video 3).

We then hand-tuned the transport rate constants to deliver equal cargo to six evenly spaced synaptic sites (top row of Fig. 6C, Supp. Video 4). Specifically, we increased the anterograde rate constants (a_i) and decreased the retrograde rate constants (b_i) by a decreasing linear function of position so that $a_{N-1} = b_{N-1}$ at the right side of the cable. On the left end of the cable, we set $a_1 = 0.5 + \beta$ and $b_1 = 0.5 - \beta$, where β is the parameter controlling anterograde bias. Intuitively, the profile of the proximal delivery bias is roughly exponential (Fig. 6B, top pattern), and therefore the anterograde rates need to be tuned more aggressively near the soma (where the bias is most pronounced), and more gently tuned as the distance to the soma increases.

However, this tuned model does not precisely deliver cargo for other stimulation patterns. For example, when the number of synapses on the cable is decreased, a distal delivery bias develops because too little cargo is released on the proximal portion of the cable (middle row, Fig. 6C; Supp. Video 5). Even when the number of synapses is held constant, changing the position of the synapses can disrupt equitable delivery of cargo to synapses. This is shown in the bottom row of figure 6C, where a distal bias again develops after the majority of activated synapses are positioned proximally. Thus, within the simple framework we've developed, the delivery of cargo can be tuned to achieve both precision and speed for a specified target distribution. However, non-specific cargo delivery occurs when different stimulation patterns are applied (assuming the transport parameters are not re-tuned).

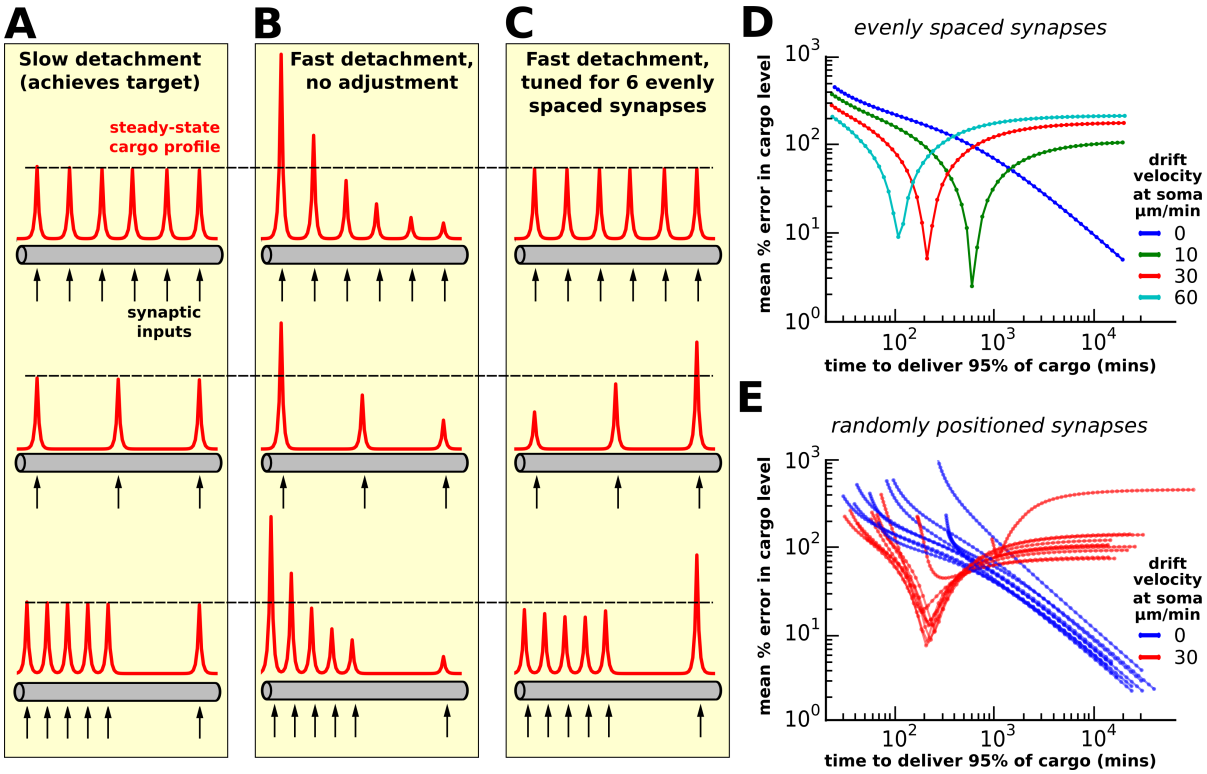

Figure 6. Tuning the model for speed and specificity is sensitive to the target distribution of cargo. **(A-C)** Cargo begins on the left end of an unbranched cable, and is ideally distributed equally amongst a number recently stimulated synaptic sites (black arrows). Steady-state cargo profiles (red) for three stimulation patterns (black arrows) along an unbranched cable. The dotted black line corresponds to the 'target' steady-state level at each delivery site. **(A)** When the timescale of detachment is sufficiently slow, cargo can be evenly distributed to the synapses regardless of their number and position. Transport parameters were set according to the procedure shown in figure 4D. **(B)** When detachment is näively increased (all rates multiplicatively scaled) a proximal bias in the steady-state distribution of cargo across all stimulation patterns. **(C)** Transport rate constants, a_i and b_i, were tuned to optimize the distribution of cargo to six equally spaced synapses (top row); detachment rate constants were the same as in panel B. Changing the number of synapses (middle row) or the position of the synapses (bottom row) causes the unequal distribution of cargo to synapses. **(D)** Tradeoff curves between non-specificity and convergence rate for six evenly spaced synapses (top row of A-C). Trafficking parameters were chosen so that the anterograde velocity decreased linearly over the length of the cable; the color of the lines shows the maximum velocity at the soma. The tradeoff curves shift to the left and becomes non-monotonic as the anterograde velocity increases. **(E)** Tradeoff curves for six randomly positioned synapses drawn uniformly across the cable. Ten simulations are shown for two levels of anterograde velocity (blue lines, 0 μm/min; red lines 30 μm/min); as before, the velocity linearly decreased across the length of the cable.

Conservative experimental estimates of trafficking parameters suggest that the tradeoff between speed and specificity is severe

To examine these observations over a larger range of stimulation patterns and transport parameters, we plotted tradeoff curves between delivery precision and speed. We first examined a cable with six evenly spaced delivery sites (same as top row of Fig. 6A-C). As before, a tradeoff between specifity and delivery speed exists for the rationally designed model, which assumes separated time scales of transport and detachment (blue line in Fig. 6D). To get

a rough estimate of how severe this tradeoff might be in real neurons, we set the length of the cable to 800 μm (roughly the length of an apical dendrite in a CA1 cell) and the diffusion coefficient to 4 $\mu m^2/s$ — an estimate on the upper end of what might be biologically achieved (see e.g., Caspi et al., 2000; Soundararajan and Bullock, 2014). Despite the optimistic estimate of the diffusion coefficient, the model predicts a severe tradeoff. It takes roughly 1 day to deliver cargo to match local demand with 10% average error, and roughly a week to deliver within 1% average error.

As the anterograde transport bias is introduced and increased, the optimal points along the tradeoff curve move to the left, representing faster transport times. The tradeoff curves also become nonmonotonic: the error (non-specificity) initially decreases as the detachment rate decreases, but begins to increase after a certain well-matched point. Points on the descending left branch of the curve represent cargo distributions with proximal bias (detachment is too fast); points on the ascending right branch correspond to distal bias (detachment is too slow).

As suggested by the simulations shown in Fig. 6A-C, changing the pattern of cargo demand changes the tradeoff curves. Thus, we performed simulations to calculate tradeoff curves for randomizing the number (between 3 and 9) and position of cargo demand hotspots (Fig. 6E). Notably, the untuned transport model (blue curves) always converge to zero error as the detachment rate decreases. In contrast, the model with anterograde bias (red curves) exhibit greater variability across demand patterns. Thus, for this model, it appears that the only way to achieve very precise, reliable and flexible transport is to have a slow detachment rate.

DISCUSSION

We formalized a well-known conceptual model of active transport — the sushi belt model (Doyle and Kiebler, 2011) — to examine how neurons distribute subcellular cargo subject to biologically realistic trafficking kinetics and dendritic morphologies. We found an inherent and surprisingly punishing trade-off between the specificity of cargo delivery and the time taken to transport it over realistic dendritic morphologies. In particular, conservative estimates based on experimental data predict delays of many hours or even days to match demand within 10%.

We formulated this model as a simple mass-action system that has a direct biological interpretation and permits analysis and simulation. Moreover, the model behaves macroscopically as a drift-diffusion process, providing a framework for interpreting and experimentally measuring model parameters (see *Methods*). From this we developed a family of models based on the assumption that sequestration of cargo from a microtubule occurs on a slower timescale than trafficking along it. Intuitively, this assumption implies that the cargo has sufficient time to sample the dendritic tree for potential delivery sites (Welte, 2004). We showed that the same distribution of cargo could be achieved by a family of model parameters ranging from location-dependent trafficking followed by uniform release (*specific trafficking model*) to uniform trafficking followed by location-dependent release (*uniform trafficking model*).

Experimental findings appear to span these possibilities, raising the intriguing question of whether transport systems might be tuned to suit the needs of specific neuron types, or physiological contexts. Kim and Martin (2015) identified three mRNAs that were uniformly distributed in cultured *Aplysia* sensory neurons, but were targeted to synapses at the level of protein expression by localized translation (supporting the uniform trafficking model). In contrast, the expression of *Arc* mRNA is closely matched to the pattern of Arc protein in granule cells of the dentate gyrus (Steward et al., 1998; Farris et al., 2014; Steward et al., 2015, supporting the specific trafficking model). Trafficking kinetics do not just differ across different cargoes — the same type of molecular cargo can exhibit diverse movement statistics in single-particle tracking experiments (Dynes and Steward, 2007). Our work places disparate

experimental findings into a common, mathematical framework and shows that the sushi belt model can produce the same distribution of cargo with various underlying parameters, supporting its biological plausibility.

Although the uniform and specific trafficking models can produce the same steady-state distribution of cargo, they converge to this distribution at different rates. The specific trafficking model can suffer from bottlenecks in the spatial pattern of cargo demand. Furthermore, these bottlenecks can be experimentally induced in cases of activity-dependent mRNA transport to test this prediction (see Fig. 3). Due to these bottlenecks, the uniform trafficking model often provides faster delivery than the specific trafficking model for a specified accuracy level (but see Fig. 5D for a counterexample). This appears to favor the uniform trafficking model from the point of view of performance and also due to the fact that uniform trafficking would be easier to implement biologically. However, neuronal activity has been observed to influencing trafficking directly (Mironov, 2007; Wang and Schwarz, 2009; Soundararajan and Bullock, 2014), and targeted trafficking may have other advantages or functions. For example, the non-uniform expression pattern of *Arc* mRNA could provide a physical substrate to store the spatial pattern of synaptic input over long time periods.

An inherent feature of all versions of the model is a trade-off between the time it takes to deliver cargo to its final destination and the accuracy with which its distribution matches the spatial profile of demand. We first asked whether the model could maintain precise delivery while relaxing the restriction of slow detachment relative to transport along a microtubule. Increasing the rate of detachment produces a proximal bias in the delivery of cargo, leading to a mismatch between supply and demand across a population of synapses. Intuitively, a fast detachment rate increases the chances for cargo to be delivered prematurely, before it has sampled the dendritic tree for an appropriate delivery site. If the transported cargo is involved in synaptic plasticity, this could produce heterosynaptic potentiation (or de-potentiation) of proximal synapses when distal synapses are strongly stimulated. Interestingly, this is indeed observed in some experimental contexts (Han and Heinemann, 2013). Although we refer to such non-specificity as "errors" in cargo delivery, it is possible that heterosynaptic plasticity relationships are leveraged to produce learning rules in circuits where synaptic inputs carrying different information sources target spatially distinct regions of the dendritic tree (Bittner et al., 2015). Thus, errors and delays might in fact be exploited biologically - especially in cases where the underlying mechanism make them inevitable. This cautions against normative assumptions about what is optimal biologically and instead emphasizes the importance of general relationships and constraints imposed by underlying mechanisms.

The sushi belt model predicts that the speed-specificity tradeoff is severe across physiologically relevant timescales. For accurate cargo delivery, the timescale of cargo dissociation needs to be roughly an order of magnitude slower than the timescale of cargo distribution/trafficking. Thus, if it takes roughly 100 minutes to distribute cargo throughout the dendrites, it will take roughly 1000 minutes ($\sim$16-17 hours) before the cargo dissociates and is delivered to the synapses. This unfavorable scaling may place fundamental limits on what can be achieved with nucleus-to-synapse signaling, which is widely thought to be an important mechanism of neuronal plasticity (Nguyen et al., 1994; Frey and Morris, 1997, 1998; Bading, 2000; Kandel, 2001; Redondo and Morris, 2011). Nonetheless, synaptic activity can initiate distal metabolic events including transcription (Kandel, 2001; Deisseroth et al., 2003; Greer and Greenberg, 2008; Ch'ng and Martin, 2011). This raises the question of whether delays for some kinds of proteins simply do not matter, or are somehow anticipated in downstream regulatory pathways. On the other hand, this might also suggest that other mechanisms, such as local translation of sequestered mRNA (Holt and Schuman, 2013), are in fact essential for processes that need to avoid nucleus-to-synapse signal

delays.

We then asked whether the speed-specificity tradeoff could be circumvented by globally tuning anterograde and retrograde trafficking kinetics as a function of distance from the soma. We were motivated by experimental studies which showed that a directional bias in transport can result from changing the proportions of expressed motor proteins (Kanai et al., 2004; Amrute-Nayak and Bullock, 2012); such a change might be induced in response to synaptic activity (Puthanveettil et al., 2008), or it may be that motor protein expression ratios are tuned during development to suit the needs of different neuron types. We were able to tune the transport rates to circumvent the speed-specificity tradeoff in an unbranched cable (Fig. 6C, top). However, even when locations for cargo demand were uniformly positioned, the trafficking rate constants needed to be tuned non-uniformly along the neurite. It is unclear whether real neurons are capable of globally tuning their trafficking rates in a non-uniform manner. Simply changing the composition of motor proteins (Amrute-Nayak and Bullock, 2012) is unlikely induce a spatial profile in transport bias. On the other hand, non-uniform modulation of the microtubule network or expression of microtubule-associated proteins provide potential mechanisms (Kwan et al., 2008; Soundararajan and Bullock, 2014), but this leaves open the question of how such non-uniform modifications arise to begin with. It would therefore be intriguing to experimentally test the existence of spatial gradients in anterograde movement bias, for example, using single-particle tracking in living neurons. Furthermore, although tuning transport bias can provide fast and rapid cargo delivery for a specific arrangement of synapses along a dendrite, tuned solutions are sensitive to changes in densities and spatial distributions of demand, as seen in Figure 6. In addition, the morphology of the neurites would affect the tuning, as introducing an anterograde bias naïvely can divert a surplus of cargo into proximal branches, resulting in an accumulation of cargo at branch tips (data not shown). Thus, although spatial tuning of transport rates can boost trafficking performance, the resulting tuned system may be extremely complex to implement and be fragile to perturbations.

Certain neuron types may nonetheless have morphology and synaptic connectivity that is sufficiently stereotyped to allow transport mechanisms to be globally tuned to some degree. In these cases the qualitative prediction that anterograde bias should decrease as a function of distance to the soma can be tested experimentally. Such tuning, while providing improved efficiency in trafficking intracellular cargo, would also constitute a site of vulnerability because alterations in the kinetics of transport or the spatial distribution of demand easily lead to mismatch between supply and demand of cargo. Molecular transport is disrupted in a number of neurodegenerative disorders (Tang et al., 2012). One obvious consequence of this is that temporal delays in cargo delivery may impair time-sensitive physiological processes. Our results highlight a less obvious kind of pathology: due to the broad requirement to appropriately match detachment and transport rates in the model, disruption of trafficking speed could additionally lead to spatially inaccurate cargo delivery. It would therefore be intriguing to examine whether pathologies caused by active transport defects are also associated with aberrant or ectopic localization of important cellular components.

Conceptual models form the basis of biological understanding, permitting interpretation of experimental data as well as motivating new experiments. It is therefore crucial to explore, quantitatively, the behaviour of a conceptual model by replacing words with equations. It is possible that active transport in biological neurons will be more robust and flexible than the predictions outlined here, which are based on a simple, but widely invoked, conceptual model. One noteworthy assumption of this model is that cargo dissociates from the microtubules by a single, irreversible reaction. In reality, the capture of cargo may be more complex and involve regulated re-attachment, which might prevent synapses from accumulating too much cargo. This possibility, and similar refinements to the sushi belt

framework, should be considered as dictated by future experimental results. The minimal models we examined provide an essential baseline for future modeling work, since it is difficult to interpret complex models without first understanding the features and limitations of their simpler counterparts.

What further experiments should be done to interrogate the model? A fundamental implication of the model is that the timescale of cargo dissociation from the microtubules is critical, as it not only determines the speed of cargo delivery, but is also directly related to the accuracy of a un-tuned transport system (Fig. 4) and the flexibility of a fine-tuned transport system (Fig. 6). Current experiments have used strong perturbations, such as chemically-induced LTP and proteolytic digestion, to induce the rapid release of cargo (Buxbaum et al., 2014a). Experimental tools that elicit more subtle effects on cargo dissociation and techniques that accurately measure these effects would enable a number of exciting investigations. For example, the sushi belt model predicts that fine-tuned transport systems are more sensitive to changes in the cargo dissociation rate. Finding evidence for fine-tuning would identify circumstances where neurons require both fast and accurate nucleus-to-synapse communication. In other contexts, neurons may tolerate sloppy cargo delivery or even take advantage of it (Han and Heinemann, 2013). Finally, for proteins that need to be rapidly, precisely, and flexibly delivered, neurons can sequester and constitutively replenish mRNA in the dendrites and rely on localized synthesis to control the spatial expression pattern (Holt and Schuman, 2013). In all cases, understanding the characteristics of intracellular transport provides clues about the underlying mechanisms of the trafficked cargo as well as it's consequences for higher level physiological functions.

METHODS

All simulation code is available online: https://github.com/ahwillia/Williams-etal-Synaptic-Transport

Model of single-particle transport

Let x_n denote the position of a particle along a 1-dimensional cable at timestep n. Let v_n denote the velocity of the particle at timestep n; for simplicity, we assume the velocity can take on three discrete values, $v_n = \{-1, 0, 1\}$, corresponding to a retrograde movement, pause, or anterograde movement. As a result, x_n is constrained to take on integer values. In the memoryless transport model (top plots in Fig. 1B, 1D, and 1F), we assume that v_n is drawn with fixed probabilities on each step. The update rule for position is:

$$x_{n+1} = x_n + v_n$$

$$v_{n+1} = \begin{cases} -1 & \text{with probability } p_- \\ 0 & \text{with probability } p_0 \\ 1 & \text{with probability } p_+ \end{cases}$$

We chose $p_- = 0.2$, $p_0 = 0.35$ and $p_+ = 0.45$ for the illustration shown in Figure 1. For the model with history-dependence (bottom plots in Fig. 1B, 1D, and 1F), the movement probabilities at each step depend on the previous movement. For example, if the motor was moving in an anterograde direction on the previous timestep, then it is more likely to continue to moving in that direction in the next time step. In this case the update rule is

written in terms of conditional probabilities:

$$v_{n+1} = \begin{cases} -1 & \text{with probability } p(-|v_n) \\ 0 & \text{with probability } p(0|v_n) \\ 1 & \text{with probability } p(+|v_n) \end{cases}$$

In the limiting (non-stochastic) case of history-dependence, the particle always steps in the same direction as the previous time step.

	$v_n = -1$	$v_n = 0$	$v_n = 1$
$p(v_{n+1} = -1)$	1	0	0
$p(v_{n+1} = 0)$	0	1	0
$p(v_{n+1} = 1)$	0	0	1

We introduce a parameter $k \in [0,1]$ to linearly interpolate between this extreme case and the memoryless model. The bottom plots of figure 1B, 1D were simulated with $k = 0.5$; in the bottom plot of figure 1F, a mass-action model was fit to these simulations.

Relationship of single-particle transport to the mass-action model

The mass-action model (equation 1, in the *Results*) simulates the bulk movement of cargo u across discrete compartments. The rate of cargo transfer is modeled as elementary chemical reactions (Keener and Sneyd, 1998). For an unbranched cable, the change in cargo in compartment i is given by:

$$\dot{u}_i = a u_{i-1} + b u_{i+1} - (a+b) u_i \tag{8}$$

For now, we assume that the anterograde and retrograde trafficking rate constants (a and b, respectively) are spatially uniform.

The mass-action model can be related to a drift-diffusion partial differential equation (Fig. 1E Smith and Simmons, 2001) by discretizing u into spatial compartments of size Δ and expanding around some position, x:

$$\dot{u}(x) \approx a\left[u(x) - \Delta\frac{\partial u}{\partial x} + \frac{\Delta^2}{2}\frac{\partial^2 u}{\partial x^2}\right] + b\left[u(x) + \Delta\frac{\partial u}{\partial x} + \frac{\Delta^2}{2}\frac{\partial^2 u}{\partial x^2}\right] - (a+b)\,u(x) \tag{9}$$

$$= a\left[-\Delta\frac{\partial u}{\partial x} + \frac{\Delta^2}{2}\frac{\partial^2 u}{\partial x^2}\right] + b\left[\Delta\frac{\partial u}{\partial x} + \frac{\Delta^2}{2}\frac{\partial^2 u}{\partial x^2}\right] \tag{10}$$

We keep terms to second order in Δ, as these are of order dt in the limit $\Delta \to 0$ (Gardiner, 2009). This leads to a drift-diffusion equation:

$$\dot{u}(x) = \frac{\partial u}{\partial t} = \underbrace{(b-a)}_{\text{drift coefficient}}\frac{\partial u}{\partial x} + \underbrace{\left(\frac{a+b}{2}\right)}_{\text{diffusion coefficient}}\frac{\partial^2 u}{\partial x^2} \tag{11}$$

Measurements of the mean and mean-squared positions of particles in tracking experiments, or estimates of the average drift rate and dispersion rate of a pulse of labeled particles can thus provide estimates of parameters a and b.

How does this equation relate to the model of single-particle transport (Fig. 1A-B)? For a memoryless biased random walk, the expected position of a particle after n time steps is $E[x_n] = n(p_+ - p_-)$ and the variance in position after n steps is $n\left(p_+ + p_- - (p_+ - p_-)^2\right)$. For large numbers of independently moving particles the mean and

variance calculations for a single particle can be directly related to the ensemble statistics outlined in the previous paragraph. We find:

$$a = \frac{2p_+ - (p_+ - p_-)^2}{2}$$

$$b = \frac{2p_- - (p_+ - p_-)^2}{2}$$

The above analysis changes slightly when the single-particle trajectories contain long, unidirectional runs. The expected position for any particle is the same $E[x_n] = n(p_+ - p_-)$; the variance, in contrast, increases as run lengths increase. However, the mass-action model can often provide a good fit in this regime with appropriately re-fit parameters (see Fig. 1F). As long as the single-particles have stochastic and identically distributed behavior, the ensemble will be well-described by a normal distribution by the central limit theorem. This only breaks down in the limit of very long unidirectional runs, as the system is no longer stochastic.

Estimating parameters of the mass-action model using experimental data

The parameters of the mass-action model we study can be experimentally fit by estimating the drift and diffusion coefficients of particles over the length of a neurite. A commonly used approach to this is to plot the mean displacement and mean squared displacement of particles as a function of time. The slopes of the best-fit lines in these cases respectively estimate the drift and diffusion coefficients in (11). Diffusion might not accurately model particle movements over short time scales because unidirectional cargo runs result in superdiffusive motion, evidenced by superlinear increases in mean squared-displacement with time (Caspi et al., 2000). However, drift-diffusion does appear to be a good model for particle transport over longer time scales with stochastic changes in direction (Soundararajan and Bullock, 2014).

The mass-action model can also be fit by tracking the positions of a population of particles with photoactivatable GFP (Roy et al., 2012). In this case, the distribution of fluorescence at each point in time could be fit by Gaussian distributions; the drift and diffusion coefficients are respectively proportional to the rate at which the mean and variance of this distribution changes.

These experimental measurements can vary substantially across neuron types, experimental conditions, and cargo identities. Therefore, in order to understand fundamental features and constraints of the sushi belt model across systems, it is more useful to explore relationships within the model across ranges of parameters. Unless otherwise stated, the trafficking kinetics were constrained so that $a_i + b_i = 1$ for each pair of connected compartments. This is equivalent to having a constant diffusion coefficient of one across all compartments. Given a target expression pattern along the microtubules, this is the only free parameter of the trafficking simulations; increasing the diffusion coefficient will always shorten convergence times, but not qualitatively change our results. In figures 1 and 6, we derived estimates of the trafficking parameters using drift and diffusion coefficients from single-particle tracking experiments (Caspi et al., 2000; Soundararajan and Bullock, 2014).

Steady-state analysis

The steady-state ratio of trafficked cargo in neighboring compartments equals the ratio of trafficking rate constants (equation 2). Consider a unbranched neurite with non-uniform anterograde and retrograde rate constants (equation

1). It is easy to verify the steady-state relationship in the first two compartments, by setting $\dot{u}_1 = 0$ and solving:

$$-a_1 u_1 + b_1 u_2 = 0 \Rightarrow \left.\frac{u_1}{u_2}\right|_{ss} = \frac{b_1}{a_1}$$

Successively applying the same logic down the cable confirms the condition in equation 2 holds globally. The more general condition for branched morphologies can be proven by a similar procedure (starting at the tips and moving in).

It is helpful to re-express the mass-action trafficking model as a matrix differential equation, $\dot{\mathbf{u}} = A\mathbf{u}$, where $\mathbf{u} = [u_1, u_2, \ldots u_N]^T$ is the state vector, and A is the state-transition matrix. For a general branched morphology, A will be nearly tridiagonal, with off-diagonal elements corresponding to branch points; matrices in this form are called Hines matrices (Hines, 1984). For the simpler case of an unbranched cable, A is tridiagonal:

$$A = \begin{bmatrix} -a_1 & b_1 & 0 & & \ldots & 0 \\ a_1 & -b_1 - a_2 & b_2 & 0 & & \\ 0 & a_2 & -b_2 - a_3 & b_3 & \ddots & \vdots \\ \vdots & 0 & a_3 & \ddots & & 0 \\ & & \ddots & & -b_{N-2} - a_{N-1} & b_{N-1} \\ 0 & & \ldots & 0 & a_{N-1} & -b_{N-1} \end{bmatrix}$$

For both branched and unbranched morphologies, each column of A sums to zero, which reflects conservation of mass within the system. The rank of A is exactly $N-1$ (this can be seen by taking the sum of the first $N-1$ rows, which results in -1 times the final row). Thus, the nullspace of A is one-dimensional (red lines in Supp. Fig. 1).

The desired steady-state distribution, $\tilde{\mathbf{u}}$, is an eigenvector that spans the nullspace of A. It is simple to show that all other eigenvalues A are negative using the Gershgorin circle theorem; thus, the fixed point described by equation 2 is stable. The convergence rate is determined by the non-zero eigenvalue with smallest magnitude of A.

CA1 pyramidal cell model

We obtained a model published by Migliore and Migliore (2012) from the online repository ModelDB (https://senselab.med.yale.edu/modeldb/), accession number 144541. We utilized this model to illustrate that the analysis holds for complex, branched morphologies. Similar results can be achieved for different morphologies and for different compartmental discretizations of the same morphology. We used the same spatial compartments used by Migliore and Migliore (2012) and set the trafficking and dissociation parameters of the mass-action transport model without reference to the geometry of the compartments. The mass-action model was simulated in Python; any simulations of the electrical activity of the model (see Fig. 2) were performed using the Python API to NEURON (Hines et al., 2009). We used a custom-written Python library to generate movies and figures for NEURON simulations (Williams, 2016).

Stochastic interpretation of the mass-action model

The mass-action approximation holds in the limit of having a large number of transported particles. Cargoes with small copy numbers will operate in a stochastic regime. Intuitively, when each compartment contains many particles, then small fluctuations in particle number don't appreciably change concentration. However, these fluctuations can

be functionally significant when the number of particles is small — for example, even a very energetically favorable reaction cannot occur with zero particles, but may occur with reasonable probability when a few particles are present.

We can thus interpret u_i as the *probability* of a particle occupying compartment i at a particular time. This invokes a standard technical assumption that the system is *ergodic*, meaning that position of one particle averaged over very long time intervals is the same as the ensemble steady-state distribution.

Thus, in addition to modeling how the spatial distribution of cargo changes over time, the mass-action model equivalently models a spatial probability distribution. That is, imagine we track a single cargo and ask its position after a long period of transport. The probability ratio between of finding this particle in any parent-child pair of compartments converges to:

$$\left.\frac{p_p}{p_c}\right|_{ss} = \frac{b}{a}$$

which matches the steady-state analysis of the deterministic model.

In the stochastic model, the number of molecules in each compartment converges to a binomial distribution at steady-state; the coefficient of variation in each compartment is given by:

$$\sqrt{\frac{1-p_i^{(ss)}}{n\,p_i^{(ss)}}}$$

This suggests two ways of decreasing noise in the system. First, increasing the total number of transported molecules, n, decreases the noise by a factor of $1/\sqrt{n}$. Additionally, transport is more reliable to compartments with a high steady-state occupation probability.

Incorporating detachment and degradation into the mass-action model

Introducing detachment and degradation reactions into the transport model is straightforward. For an arbitrary compartment in a cable, the differential equations become:

$$\begin{aligned} \dot{u}_i &= a_{i-1}u_{i-1} - (a_i + b_{i-1} + c_i)u_i + b_i u_{i+1} \\ \dot{u}_i^\star &= c_i u_i - d_i u_i^\star \end{aligned}$$

When $a_i, b_i \gg c_i \gg d_i$, then the variables u_i and $u_i^\star$ approach a quasi-steady-state, which we denote $\tilde{u}_i$ and $\tilde{u}_i^\star$. For simplicity we assume $d_i = 0$ in simulations. We present two strategies for achieving a desired distribution for $\tilde{u}_i^\star$ in figure 4C and 4D. To interpolate between these strategies, let F be a scalar between 0 and 1, and let $\tilde{u}^\star$ be normalized to sum to one. We choose a_i and b_i to achieve:

$$\tilde{u}_i = F\ \tilde{u}_i^\star + (1-F)/N$$

along the microtubular network and choose c_i to satisfy

$$c_i \propto \frac{\tilde{u}_i^\star}{F\ \tilde{u}_i^\star + (1-F)/N}$$

Here, N is the number of compartments in the model. Setting $F = 1$ results in the simulation in the "specific" trafficking model (Fig. 4C), while setting $F = 0$ results in the "uniform" trafficking model (Fig. 4D). An interpolated

strategy is shown in figure 4E ($F = 0.3$).

Globally tuning transport rates to circumvent the speed-specificity tradeoff

We investigated the *uniform trafficking model* with fast detachment rates in an unbranched cable with equally spaced synapses and $N = 100$ compartments. Multiplicatively increasing the detachment rates across the cable produced a proximal bias in cargo delivery which could be corrected by setting the anterograde and retrograde trafficking rates to be:

$$a_i = 0.5 + \beta \cdot \frac{N-1-i}{N-2}$$

$$b_i = 0.5 - \beta \cdot \frac{N-1-i}{N-2}$$

where $i = \{1, 2, ...N-1\}$ indexes the trafficking rates from the soma ($i = 1$) to the other end of the cable ($i = N-1$). Faster detachment rates require larger values for the parameter β; note that $\beta < 0.5$ is a constraint to prevent b_i from becoming negative. This heuristic qualitatively improved, but did not precisely correct for, fast detachment rates in the *specific trafficking model* (data not shown).

ACKNOWLEDGEMENTS

We thank Eve Marder, Subhaneil Lahiri, Friedemann Zenke, and Benjamin Regner for useful feedback on the manuscript, and thank Simon Bullock for useful discussion. This research was supported by the Department of Energy Computational Science Graduate Fellowship, NIH Grant 1P01NS079419, NIH Grant P41GM103712, the Howard Hughes Medical Institute, and the Charles A. King Trust.

COMPETING INTERESTS

The authors declare that there are no competing interests.

REFERENCES

Aakalu, G., Smith, W. B., Nguyen, N., Jiang, C., and Schuman, E. M. (2001). Dynamic visualization of local protein synthesis in hippocampal neurons. *Neuron*, 30(2):489–502.

Amrute-Nayak, M. and Bullock, S. L. (2012). Single-molecule assays reveal that RNA localization signals regulate dynein-dynactin copy number on individual transcript cargoes. *Nat. Cell Biol.*, 14(4):416–423.

Bading, H. (2000). Transcription-dependent neuronal plasticity. *Eur. J. Biochem.*, 267(17):5280–5283.

Bittner, K. C., Grienberger, C., Vaidya, S. P., Milstein, A. D., Macklin, J. J., Suh, J., Tonegawa, S., and Magee, J. C. (2015). Conjunctive input processing drives feature selectivity in hippocampal CA1 neurons. *Nat. Neurosci.*, 18(8):1133–1142.

Block, S. M., Goldstein, L. S. B., and Schnapp, B. J. (1990). Bead movement by single kinesin molecules studied with optical tweezers. *Nature*, 348(6299):348–352.

Bressloff, P. and Newby, J. (2009). Directed intermittent search for hidden targets. *New J. Phys.*, 11(2):023033.

Bressloff, P. C. (2006). Stochastic model of protein receptor trafficking prior to synaptogenesis. *Phys. Rev. E Stat. Nonlin. Soft. Matter. Phys.*, 74(3).

Bressloff, P. C. and Newby, J. M. (2013). Stochastic models of intracellular transport. *Rev. Mod. Phys.*, 85(1):135–196.

Buxbaum, A. R., Haimovich, G., and Singer, R. H. (2014a). In the right place at the right time: visualizing and understanding mRNA localization. *Nat. Rev. Mol. Cell Biol.*, 16(2):95–109.

Buxbaum, A. R., Wu, B., and Singer, R. H. (2014b). Single-Actin mRNA Detection in Neurons Reveals a Mechanism for Regulating Its Translatability. *Science*, 343(6169):419–422.

Buxbaum, A. R., Yoon, Y. J., Singer, R. H., and Park, H. Y. (2015). Single-molecule insights into mRNA dynamics in neurons. *Trends Cell Biol.*, 25(8):468–475.

Cajigas, I. J., Tushev, G., Will, T. J., tom Dieck, S., Fuerst, N., and Schuman, E. M. (2012). The local transcriptome in the synaptic neuropil revealed by deep sequencing and high-resolution imaging. *Neuron*, 74(3):453–466.

Caspi, A., Granek, R., and Elbaum, M. (2000). Enhanced diffusion in active intracellular transport. *Phys. Rev. Lett.*, 85:5655–5658.

Ch'ng, T. and Martin, K. (2011). Synapse-to-nucleus signaling. *Curr. Opin. Neurobiol.*, 21:345–52.

Deisseroth, K., Mermelstein, P. G., Xia, H., and Tsien, R. W. (2003). Signaling from synapse to nucleus: the logic behind the mechanisms. *Curr. Opin. Neurobiol.*, 13(3):354–365.

Doyle, M. and Kiebler, M. A. (2011). Mechanisms of dendritic mRNA transport and its role in synaptic tagging. *EMBO J.*, 30(17):3540–3552.

Dynes, J. L. and Steward, O. (2007). Dynamics of bidirectional transport of *Arc* mRNA in neuronal dendrites. *J. Comp. Neurol.*, 500(3):433–447.

Farris, S., Lewandowski, G., Cox, C. D., and Steward, O. (2014). Selective Localization of *Arc* mRNA in Dendrites Involves Activity- and Translation-Dependent mRNA Degradation. *J. Neurosci.*, 34(13):4481–4493.

Frey, U. and Morris, R. (1998). Weak before strong: dissociating synaptic tagging and plasticity-factor accounts of late-LTP. *Neuropharmacol.*, 37(4-5):545–552.

Frey, U. and Morris, R. G. M. (1997). Synaptic tagging and long-term potentiation. *Nature*, 385(6616):533–536.

Gagnon, J. and Mowry, K. (2011). Molecular motors: directing traffic during RNA localization. *Crit Rev Biochem Mol Biol*, 46:229–39.

Gardiner, C. (2009). *Stochastic Methods: A Handbook for the Natural and Social Sciences (Springer Series in Synergetics)*. Springer, 4 edition.

Grant, B. D. and Donaldson, J. G. (2009). Pathways and mechanisms of endocytic recycling. *Nat. Rev. Mol. Cell Biol.*, 10(9):597–608.

Greer, P. L. and Greenberg, M. E. (2008). From synapse to nucleus: Calcium-dependent gene transcription in the control of synapse development and function. *Neuron*, 59(6):846 – 860.

Han, E. B. and Heinemann, S. F. (2013). Distal Dendritic Inputs Control Neuronal Activity by Heterosynaptic Potentiation of Proximal Inputs. *J. Neurosci.*, 33(4):1314–1325.

Hancock, W. O. (2014). Bidirectional cargo transport: moving beyond tug of war. *Nat. Rev. Mol. Cell Biol.*, 15(9):615–628.

Hines, M. (1984). Efficient computation of branched nerve equations. *Int. J. Biomed. Comput.*, 15(1):69–76.

Hines, M. L., Davison, A. P., and Muller, E. (2009). Neuron and python. *Front. Neuroinform.*, 3.

Hirokawa, N., Niwa, S., and Tanaka, Y. (2010). Molecular Motors in Neurons: Transport Mechanisms and Roles in Brain Function Development, and Disease. *Neuron*, 68(4):610–638.

Hoffman, D., Magee, J., Colbert, C., and Johnston, D. (1997). K^+ channel regulation of signal propagation in dendrites of hippocampal pyramidal neurons. *Nature*, 387:869–75.

Holt, C. E. and Schuman, E. M. (2013). The central dogma decentralized: New perspectives on RNA function and local translation in neurons. *Neuron*, 80(3):648–657.

Kanai, Y., Dohmae, N., and Hirokawa, N. (2004). Kinesin transports RNA: isolation and characterization of an RNA-transporting granule. *Neuron*, 43(4):513–525.

Kandel, E. R. (2001). The Molecular Biology of Memory Storage: A Dialogue Between Genes and Synapses. *Science*, 294(5544):1030–1038.

Kang, H. and Schuman, E. M. (1996). A requirement for local protein synthesis in neurotrophin-induced hippocampal synaptic plasticity. *Science*, 273(5280):1402–1406.

Kao, D.-I., Aldridge, G. M., Weiler, I. J., and Greenough, W. T. (2010). Altered mrna transport, docking, and protein translation in neurons lacking fragile x mental retardation protein. *Proc. Natl. Acad. Sci. USA*, 107(35):15601–15606.

Keener, J. P. and Sneyd, J. (1998). *Mathematical Physiology*. Springer, New York.

Kim, S. and Martin, K. C. (2015). Neuron-wide RNA transport combines with netrin-mediated local translation to spatially regulate the synaptic proteome. *eLife*, 4:e04158.

Krichevsky, A. M. and Kosik, K. S. (2001). Neuronal RNA GranulesA Link between RNA Localization and Stimulation-Dependent Translation. *Neuron*, 32(4):683–696.

Kwan, A. C., Dombeck, D. A., and Webb, W. W. (2008). Polarized microtubule arrays in apical dendrites and axons. *Proc. Natl. Acad. Sci. U.S.A.*, 105(32):11370–11375.

Lyles, V., Zhao, Y., and Martin, K. C. (2006). Synapse Formation and mRNA Localization in Cultured Aplysia Neurons. *Neuron*, 49(3):349–356.

Magee, J. (1998). Dendritic hyperpolarization-activated currents modify the integrative properties of hippocampal CA1 pyramidal neurons. *J Neurosci*, 18:7613–24.

Migliore, M. and Migliore, R. (2012). Know your current I_h: Interaction with a shunting current explains the puzzling effects of its pharmacological or pathological modulations. *PLoS ONE*, 7(5):e36867.

Mironov, S. L. (2007). ADP Regulates Movements of Mitochondria in Neurons. *Biophys. J.*, 92(8):2944–2952.

Moga, D., Calhoun, M., Chowdhury, A., Worley, P., Morrison, J., and Shapiro, M. (2004). Activity-regulated cytoskeletal-associated protein is localized to recently activated excitatory synapses. *Neurosci.*, 125(1):7–11.

Müller, M. J. I., Klumpp, S., and Lipowsky, R. (2008). Tug-of-war as a cooperative mechanism for bidirectional cargo transport by molecular motors. *Proc. Natl. Acad. Sci. U.S.A.*, 105(12):4609–4614.

Newby, J. and Bressloff, P. C. (2010). Local synaptic signaling enhances the stochastic transport of motor-driven cargo in neurons. *Phys. Biol.*, 7(3):036004.

Nguyen, P. V., Abel, T., and Kandel, E. R. (1994). Requirement of a critical period of transcription for induction of a late phase of LTP. *Science*, 265(5175):1104–1107.

Park, H. Y., Lim, H., Yoon, Y. J., Follenzi, A., Nwokafor, C., Lopez-Jones, M., Meng, X., and Singer, R. H. (2014). Visualization of dynamics of single endogenous mRNA labeled in live mouse. *Science*, 343(6169):422–424.

Park, M., Penick, E. C., Edwards, J. G., Kauer, J. A., and Ehlers, M. D. (2004). Recycling endosomes supply AMPA receptors for LTP. *Science*, 305(5692):1972–1975.

Park, M., Salgado, J. M., Ostroff, L., Helton, T. D., Robinson, C. G., Harris, K. M., and Ehlers, M. D. (2006). Plasticity-induced growth of dendritic spines by exocytic trafficking from recycling endosomes. *Neuron*, 52(5):817–830.

Puthanveettil, S. V., Monje, F. J., Miniaci, M. C., Choi, Y.-B., Karl, K. A., Khandros, E., Gawinowicz, M. A., Sheetz, M. P., and Kandel, E. R. (2008). A New Component in Synaptic Plasticity: Upregulation of Kinesin in the Neurons of the Gill-Withdrawal Reflex. *Cell*, 135(5):960–973.

Redondo, R. L. and Morris, R. G. (2011). Making memories last: the synaptic tagging and capture hypothesis. *Nat. Rev. Neurosci.*, 12(1):17–30.

Rogers, S. L. and Gelfand, V. I. (1998). Myosin cooperates with microtubule motors during organelle transport in melanophores. *Curr. Biol.*, 8(3):161–164.

Roy, S., Yang, G., Tang, Y., and Scott, D. A. (2012). A simple photoactivation and image analysis module for visualizing and analyzing axonal transport with high temporal resolution. *Nat. Protoc.*, 7(1):62–68.

Smith, D. and Simmons, R. (2001). Models of Motor-Assisted Transport of Intracellular Particles. *Biophys. J.*, 80(1):45–68.

Soundararajan, H. C. and Bullock, S. L. (2014). The influence of dynein processivity control, maps, and microtubule ends on directional movement of a localising mRNA. *eLife*, 3:e01596.

Steward, O., Farris, S., Pirbhoy, P. S., Darnell, J., and Driesche, S. J. V. (2015). Localization and local translation of *Arc/Arg3.1* mRNA at synapses: some observations and paradoxes. *Front. Mol. Neurosci.*, 7.

Steward, O., Wallace, C., Lyford, G., and Worley, P. (1998). Synaptic Activation Causes the mRNA for the IEG *Arc* to Localize Selectively near Activated Postsynaptic Sites on Dendrites. *Neuron*, 21(4):741–751.

Steward, O. and Worley, P. F. (2001). Selective Targeting of Newly Synthesized *Arc* mRNA to Active Synapses Requires NMDA Receptor Activation. *Neuron*, 30(1):227–240.

Sutton, M. A. and Schuman, E. M. (2006). Dendritic protein synthesis, synaptic plasticity, and memory. *Cell*, 127(1):49–58

Tang, Y., Scott, D. A., Das, U., Edland, S. D., Radomski, K., Koo, E. H., and Roy, S. (2012). Early and selective impairments in axonal transport kinetics of synaptic cargoes induced by soluble amyloid β-protein oligomers. *Traffic*, 13(5):681–693.

Verbrugge, S., van den Wildenberg, S. M., and Peterman, E. J. (2009). Novel Ways to Determine Kinesin-1's Run Length and Randomness Using Fluorescence Microscopy. *Biophys. J.*, 97(8):2287–2294.

Vickers, C. A., Dickson, K. S., and Wyllie, D. J. A. (2005). Induction and maintenance of late-phase long-term potentiation in isolated dendrites of rat hippocampal CA1 pyramidal neurones. *J. Physiol.*, 568(3):803–813.

Wang, X. and Schwarz, T. L. (2009). The Mechanism of Ca^{2+}-Dependent Regulation of Kinesin-Mediated Mitochondrial Motility. *Cell*, 136(1):163–174.

Welte, M. A. (2004). Bidirectional transport along microtubules. *Curr. Biol.*, 14(13):R525–R537.

Williams, A. H. (2016). PyNeuron Toolbox. https://github.com/ahwillia/PyNeuron-Toolbox.

Zheng, N., Jeyifous, O., Munro, C., Montgomery, J. M., and Green, W. N. (2015). Synaptic activity regulates AMPA receptor trafficking through different recycling pathways. *eLife*, 4.