Dendritic trafficking faces fundamental 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, toleary@brandeis.edu

ABSTRACT

Most neurons have complex morphologies with long processes and dynamic, nonuniform spatial expression patterns of subcellular organelles and macromolecules. It is thought that trafficking and delivery of this subcellular cargo depends on purely local signals, rather than a global addressing system. Yet it remains unclear how such a decentralized system performs in complex morphologies. We mathematically formalize a previously proposed "sushi belt" model of microtubule transport (Doyle and Kiebler, 2011) and show how arbitrarily complex spatial distributions of cargo can be achieved by local signaling mechanisms. We reveal that this model predicts an unavoidable and physiologically critical tradeoff between speed and precision of cargo delivery that can be tested experimentally. More sophisticated variants of the sushi belt model can provide both fast and precise transport; however, these require global tuning of trafficking kinetics, and their performance is fragile to changes in the required spatial distribution of cargo.

Keywords: Microtubules, Active transport, Tagging hypothesis, Dendrites, Transport, Morphology

INTRODUCTION

The axonal and dendritic trees of most neurons are stunningly complex. The logistical task of distributing biomolecular components within these morphologies is therefore considerable, especially for components that are synthesized in locations that are distal from their site of use. Although the molecular machinery responsible for active transport has been characterized in detail (Doyle and Kiebler, 2011; Buxbaum et al., 2014a; Hancock, 2014), there have been few attempts to understand how these mechanisms can be orchestrated to distribute cargo accurately and efficiently in realistic neuron morphologies.

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). Similarly, constitutive maintenance of cytoskeletal, membrane and signalling 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, many long-lasting forms of synaptic plasticity are known to depend on anterograde transport of mRNAs (Kandel, 2001; Puthanveettil et al., 2008) 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. More generally, any site or region of a dendrite that requires specific biochemical components presumably needs to signal demand for these components and sequester them as they are delivered.

Kinesin and dyenin motor proteins mediate transport along microtubules at rates that far exceed passive diffusion (Block et al., 1990; Smith and Simmons, 2001; Hirokawa et al., 2010; Gagnon and Mowry, 2011; Park et al., 2014). The movement of individual cargoes is often stochastic and bidirectional (Hancock, 2014), motivating the hypothesis that motor proteins search the dendritic tree for dropoff locations, instead of following a pre-determined path (Welte, 2004; Bressloff and Newby, 2009; Doyle and Kiebler, 2011). Thus, the local rate of transport and the decision to release cargo are thought to be predominantly controlled by localized signaling pathways, rather than a centralized addressing system. 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; Buxbaum et al., 2014b), and changes in microtubule-associate proteins (Soundararajan and Bullock, 2014). These pathways can be activated by the spatial pattern of synaptic activity, the spontaneous release of calcium from internal stores, or by released factors from other cells (Wong and Ghosh, 2002; Parrish et al., 2007; Park and Poo, 2013; Zagrebelsky and Korte, 2014).

Based on these observations, a number of reviews have advanced a conceptual model of molecular transport, 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 is typical in certain restaurants. Customers sitting alongside the belt (representing specific locations or synapses along a dendrite) each have specific and potentially time-varying demands for the amount and type of sushi they consume, but they can only choose from nearby plates as they pass.

While the sushi belt model is intuitively plausible, a number of open questions remain. How can a trafficking system based on localized signals be engineered to accurately generate spatial distributions of cargo? How can two spatially separated regions of the neuron avoid interfering with each other, even though they can’t directly communicate? 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?

Here, we formulate the sushi belt model in a minimal, biophysically plausible set of equations, and show that it can reliably produce complex spatial distributions of cargo. However, we also reveal several logistical constraints on this family of models. In particular, we show that fast cargo delivery is error-prone, in the sense that the amount of cargo delivered is disproportionate to the local demand for cargo. We outline experimental predictions to test whether this tradeoff is present in real neurons, or 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 (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 (Muller 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 Muller 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 cable with N compartments, the mass-action model is (Fig. 1E):

$$u_1 \underset{b_1}{\overset{a_1}{\rightleftharpoons}} u_2 \underset{b_2}{\overset{a_2}{\rightleftharpoons}} u_3 \underset{b_3}{\overset{a_3}{\rightleftharpoons}} \ \dots \ \underset{b_{N-1}}{\overset{a_{N-1}}{\rightleftharpoons}} u_N \qquad (1)$$

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 our 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) 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, measurements of these two quantities on a specific stretch of dendrite provide local estimates of 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; Muller 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.

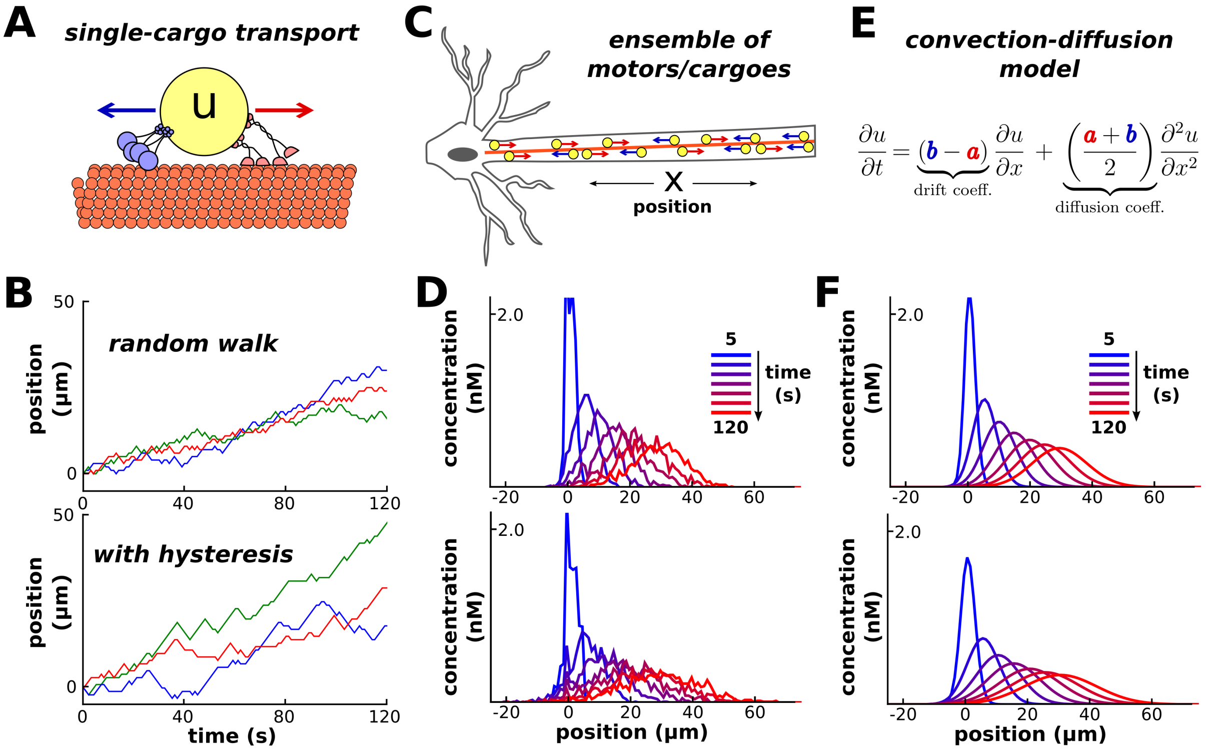

Figure 1. Model of microtubular transport. **(A)** Molecular cargo, *u*, undergoes stochastic back-and-forth movements driven by opposing motor proteins. **(B)** Three example random walks on a cable representing the movement of individual cargoes. In a simple random walk, each movement is independent 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 ensemble of synaptic cargoes transported along the length of a neurite. **(D)** The concentration profile of transported cargo along a cable over time, simulated as 1,000 independent random walks. Simulations with (bottom) and without (top) history-dependence. **(E)** In the limit of many individual 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 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 (Muller 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 our 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. This situation corresponds to the kinetics of motor proteins being dependent 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 $\tilde{u}$ (we will soon show how this quantity can be encoded by a localized biochemical signal). 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, with the slope of the linear gradient controlled by tuning the total amount of cargo in the cable (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 biologically plausible mechanism for controlling these rates. It is reasonable to assume that local trafficking rates can be modulated by a biochemical signal resulting from synaptic activation or local release of a growth factor. An experimentally characterized signal is provided by transient fluctuations in cytoplasmic calcium concentration, $[Ca]_i$, because increases in $[Ca]_i$ are known to simultaneously arrest anterograde and retrograde microtubular transport for certain cargoes (Wang and Schwarz, 2009). We therefore assume, for any pair of compartments, the anterograde rate constant is determined by 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

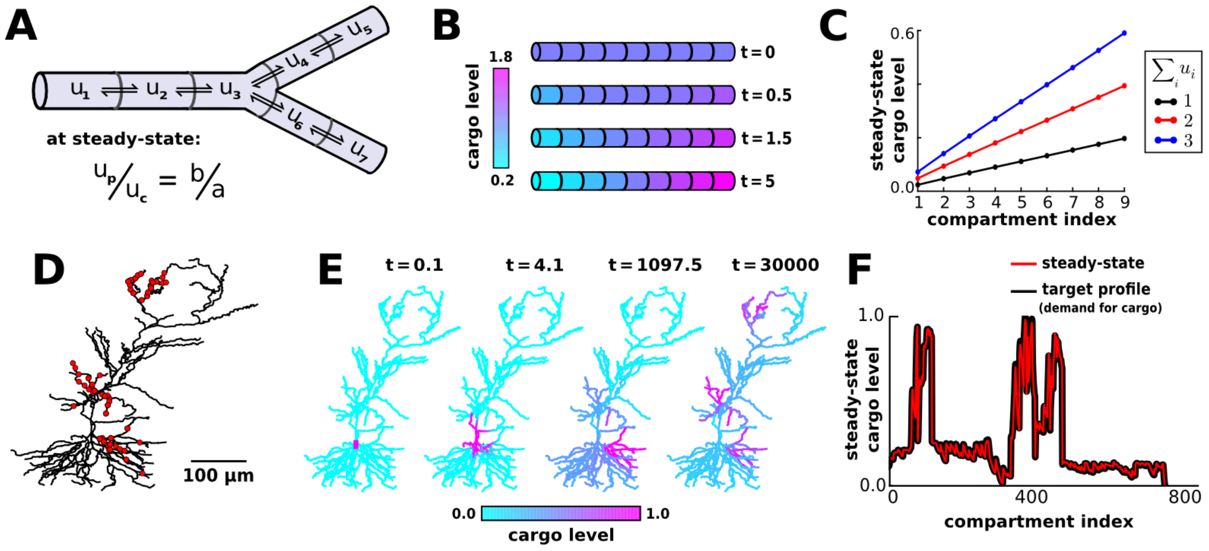

Figure 2. Local trafficking parameters determine the spatial distribution of biomolecules by a simple rule. **(A)** Diagram of the mass action transport model for a simple branched morphology. **(B)** A simulation of a nine compartment cable, with trafficking rate constants set to produce a linear gradient 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.

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 serve to arrest transport could provide a mechanism for distributing cargo to an arbitrary target profile. We note that other potential signalling pathways could achieve the same effect, so while there is direct evidence that $[Ca]_i$ as a key 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 extended an existing multi-compartment 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 approaches the desired steady-state exactly in the absence of noise.

Convergence rate

Biochemical processes are time-sensitive. For example, newly synthesized proteins must be delivered to synapses within ∼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. The convergence rate is determined by the smallest magnitude, non-zero eigenvalue of the system (see Methods), which can be thought of as a rate-limiting step or process for the system. Simulations confirmed this analysis and showed that the convergence rate 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; thus, 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 enough plasticity-related factors are delivered from the soma. 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). The presence of a bottleneck *shortens* the transport delay to proximal 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 (Fig. 3D-E). This might be experimentally tested by measuring the convergence time of a retrogradely transported molecule 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

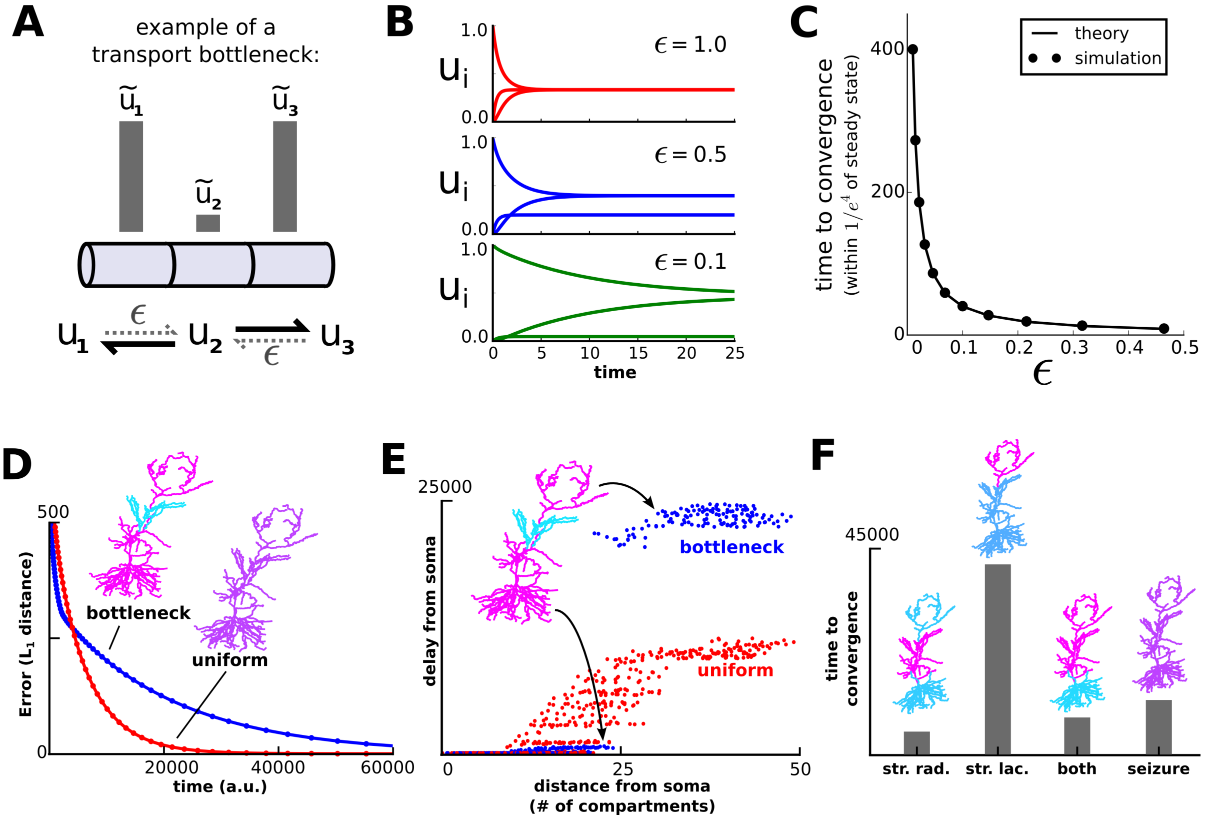

Figure 3. Convergence to steady-state is slow when molecules must be transported across bottlenecks —compartments with low target levels. **(A)** A three-compartment transport model, with the middle compartment acting as a bottleneck. The vertical bars represent the desired steady-state level of cargo in each compartment. The rate of transport into the middle compartment is small (ε, dashed harpoons) 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 distance) 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)** Predicted convergence times for various target distributions (*str. rad.*, stratum radiatum; *str. lac.*, stratum lacunosum/moleculare) in the CA1 model.

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. 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 disassembled or sequestered, it is reasonable to model detachment from the microtubule as

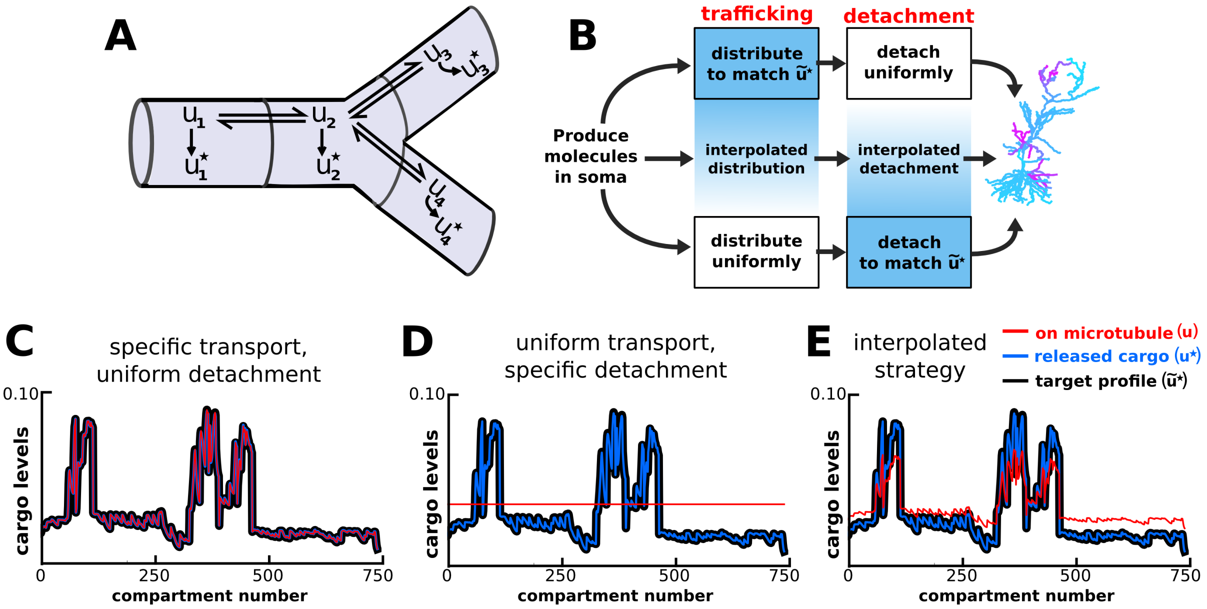

Figure 4. Multiple strategies for transport in a model including nonuniform microtubular detachment/activation. **(A)** Schematic of microtubular transport model with irreversible detachment in a branched morphology. The localized degradation reactions ($u_i^\star \xrightarrow{d_i} \varnothing$) are 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 (see circled red numbers).

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:

$$\begin{tikzcd}
u_1 \arrow[r, shift left, "a_1"] \arrow[d, "c_1"'] & u_2 \arrow[l, shift left, "b_1"] \arrow[r, shift left, "a_2"] \arrow[d, "c_2"'] & u_3 \arrow[l, shift left, "b_2"] \arrow[r, shift left, "a_3"] \arrow[d, "c_3"'] & u_4 \arrow[l, shift left, "b_3"] \arrow[r, shift left, "a_4"] \arrow[d, "c_4"'] & \ldots \arrow[l, shift left, "b_4"] \\
u_1^\star \arrow[d, "d_1"'] & u_2^\star \arrow[d, "d_2"'] & u_3^\star \arrow[d, "d_3"'] & u_4^\star \arrow[d, "d_4"'] & \\
\varnothing & \varnothing & \varnothing & \varnothing &
\end{tikzcd} \tag{6}$$

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 our previous analysis (equation 3). We then choose detachment rate constants that transform the microtubular distribution into our 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.

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). 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 to speed up delivery globally. 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 (Fig. 5A). If the average 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

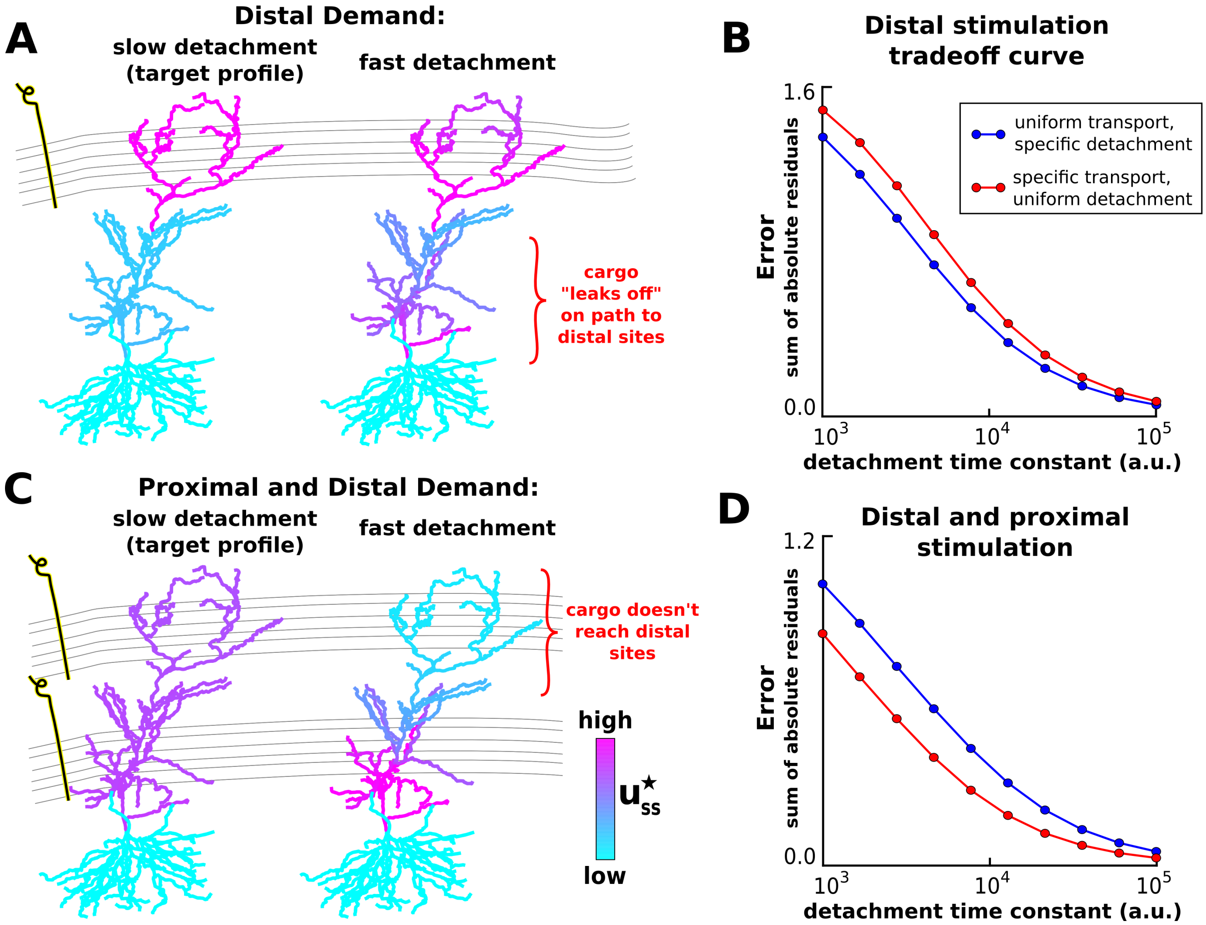

Figure 5. Proximal synapses capture more cargo at the expense of distal synapses when detachment rates are naïvely increased. **(A)** Delivery of cargo to the distal apical zone with slow *(left)* and fast detachment rates *(right)*. The achieved pattern does not match the target distribution when detachment 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.

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 loss of specificity and that the final destination of mis-trafficked cargo depends on the pattern of stimulation. We next asked whether we can tune the kinetic parameters of the transport mechanism to find an ideal compromise between precision and trafficking efficiency.

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. That is, synapses closer to the soma are more likely to capture cargo at the expense of distal synapses. 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).

Conservative 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 hard 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 μm^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

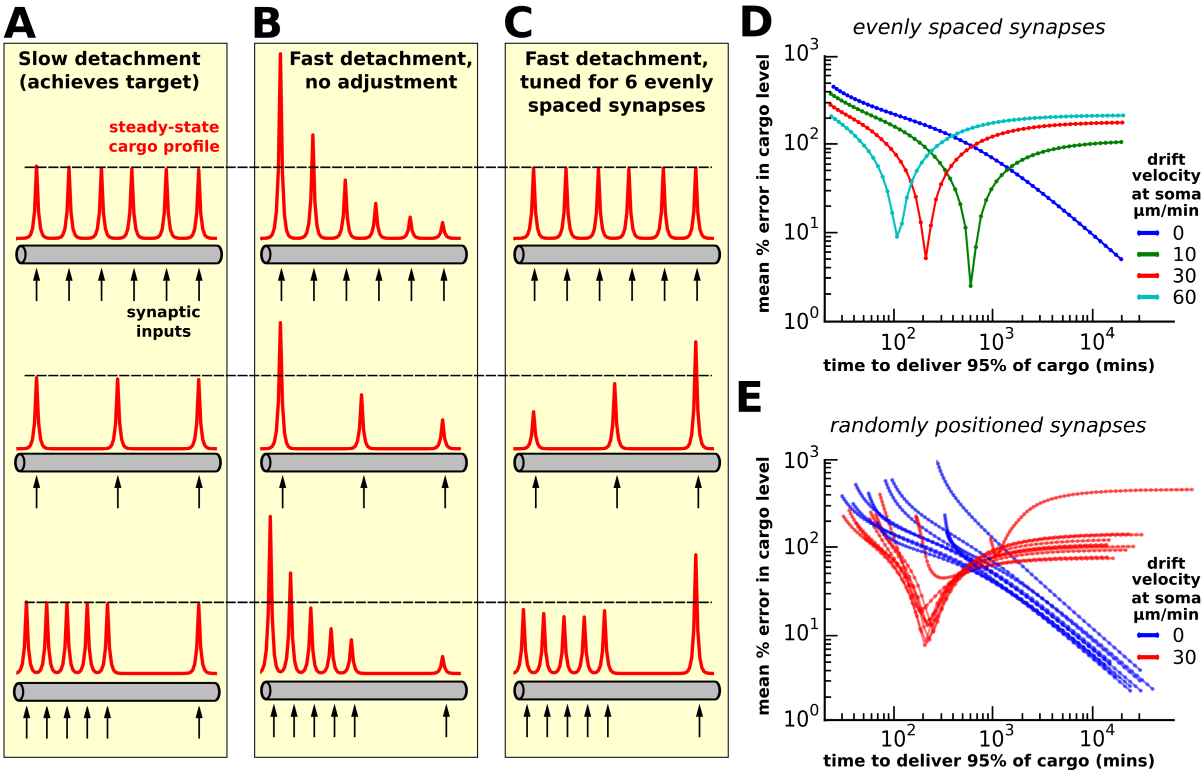

Figure 6. Tuning the model for speed and specificity is sensitive to the target distribution of cargo. **(A-C)** Cargo begins on the left end of an unbranched cable, and is ideally distributed equally amongst a number recently stimulated synaptic sites (black arrows). Steady-state cargo profiles (red) 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.

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 (other line colors), 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

A microscope image of a typical dendritic tree hints at the complex logistical task that the neuron must solve in order to distribute its biomolecular components. While this problem is relevant to all aspects of cellular physiology in dendrites, synapses in particular are known to signal demand for specific kinds of molecules dynamically and can initiate distal metabolic events including transcription (Kandel, 2001; Deisseroth et al., 2003; Greer and Greenberg, 2008; Ch'ng and Martin, 2011). This generates complex fluctuations in demand for specific kinds of molecular cargo that are involved in structural and physiological remodelling during plasticity and development. For example, excitatory synapses in the mammalian nervous system comprise hundreds of different proteins, including receptors, structural and anchoring proteins, and cytosolic signaling enzymes (Liu et al., 2014; Laßek et al., 2015). How do these components find their way to their final destination?

We formalized a well-known, parsimonious and plausible conceptual model of molecular transport, the sushi belt model (Doyle and Kiebler, 2011), to address this general question. We formulated this model as a simple mass-action system that has a direct biological interpretation and permits analysis and simulation. 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, a consequence of this assumption is 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 trafficking and release regimes: location-dependent trafficking followed by uniform release, uniform trafficking followed by location-dependent release, or a combination of these extreme cases (Fig. 4). Experimental findings appear to span these possibilities. Kim and Martin (2015) identified three mRNAs that were uniformly distributed throughout cultured *Aplysia* sensory neurons, but were targeted to synapses at the level of protein expression by localized translation. In contrast, the nonuniform 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). Even the same type of molecular cargo can show diverse movement statistics in single-particle tracking experiments (Dynes and Steward, 2007). Thus, our work places disparate experimental findings into a common conceptual framework and shows that a simple sushi belt model of transport can, in principle, maintain a specific distribution of cargo throughout a neuron and that this capacity is relatively independent of whether trafficking or cargo sequestration (or both) are controlled by local signals.

The model has a straightforward interpretation as a drift-diffusion process and the parameters describing drift rate and diffusion rate map directly to parameters in an equivalent mass-action scheme. This allowed us to estimate realistic rate parameters in the model from experimental data and then determine the consequences for distributing cargo throughout a typical dendritic arbor under a variety of assumptions and conditions. An inherent feature of the model is a trade-off between the rate at which cargo is distributed and the specificity, or 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. We found that 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 microtubule detachment rate increases the chances for cargo to be delivered prematurely, without sufficiently searching for an appropriate delivery site.

When transport and sequestration are solely controlled by local signals the model predicts that the speed-specificity tradeoff is severe. In this version of the model, explored in Figures 4, 5 and 6, 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 unfavourable scaling has a number of implications. First, it might explain in part why local protein synthesis is observed in neurites across many species and neuron types. Long delays caused by microtubule transport from the soma would make rapid plasticity and metabolic responses impossible, severely limiting the regulatory capacity of neurons (Frey and Morris, 1997, 1998; Redondo and Morris, 2011). This may place fundamental limits on what can be achieved with synapse-to-nucleus signaling, which is widely thought to be an essential mechanism of neuronal plasticity. However, evidently not all dendritically expressed proteins are locally synthesized, and transcriptional blockade is known to interfere with regulation of synaptic and dendritic function (Nguyen et al., 1994; Bading, 2000). This raises the question of whether delays for some kinds of proteins simply do not matter, or whether more sophisticated transport mechanisms might exist to mitigate delays.

We then asked whether the speed-specificity tradeoff could be circumvented by globally tuning basal anterograde and retrograde movement rates as a function of distance from the soma. Studies show that a directional bias in transport can be induced by changing the complement of motor proteins (Kanai et al., 2004; Amrute-Nayak and Bullock, 2012), and that such a change might be induced at the soma in response to synaptic activity (Puthanveettil et al., 2008). We were able to tune the transport rates to circumvent the speed-specificity tradeoff (Fig. 6C, top). Even when locations for cargo demand were uniformly placed across the cable, we found the trafficking rate constants needed to be tuned non-uniformly along a dendrite; in the cases we examined this could be achieved with a linearly decreasing profile (see *Methods*). However, biologically, it is unclear whether 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 very sensitive to changes in densities and spatial distributions of demand, as seen in Figure 6. This indicates that significant alterations in synaptic input or connectivity along a dendrite would preclude a tuned transport mechanism of this kind. In addition, the morphology of the neurites would affect the tuning, as introducing an anterograde bias can cause the accumulation of cargo at the tips of short branches, proximal to the soma (data not shown). However, other mechanisms, not modeled in this study, might prevent this accumulation; for example, upon reaching the end of a neurite, a unit of cargo could undergo retrograde transport for long distances (Soundararajan and Bullock, 2014).

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 potential explanation for this is that temporal delays in cargo distribution 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.

What further experiments should be done to interrogate the model and test the predicted tradeoffs in speed, precision, and flexibility of transport? Our results suggest that the timescale of cargo dissociation from the microtubules is a critical parameter that merits careful examination. Local release of mRNA cargo can be induced by strong experimental manipulations, such as chemically-induced LTP and proteolytic digestion (Buxbaum et al., 2014a). Developing experimental tools that elicit more subtle effects on dissociation rate, as well as protocols to accurately observe and characterize dissociation in more naturalistic settings, would constrain the set of models we examined in this study. For example, our work predicts that decreasing the dissociation rate should produce minimal effects on the spatial distribution of cargo if dissociation occurs on a slower time scale than trafficking, whereas fine-tuned strategies (Fig. 6) are more sensitive to these manipulations.

It is possible that active transport in biological neurons will be more robust and flexible than our predictions, which are based on a simple, but widely invoked conceptual model. In particular, we modeled the dissociation of cargo from the microtubules as a single, irreversible chemical reaction. In reality, the capture of cargo may be more complex and involve regulated re-attachment. Similarly, the over-accumulation of cargo at a subset of synapses might be ameliorated if the dissociation process exhibited saturation or adaptation. Incorporating these details would change the behaviour of the model but would also need more detailed experiments to motivate them. Nonetheless, the minimal models we examined provide a basis for designing and interpreting future experiments, and serve as a rigorous exploration of an important and popular conceptual model.

METHODS

All simulation code is available on: 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 probabilites 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 utilizes *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 extreme 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 approximates the drift-diffusion partial differential equation (Fig. 1E) by discretizing u into spatial compartments. This can be seen by Taylor expanding around some position $u(x)$:

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

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

where Δ is the compartment size. We keep the first two terms of the Taylor expansion, due to the classic result that $dx^2 \sim dt$ for a diffusion process (Gardiner, 2009). In the limit $\Delta \to 0$, we arrive at the drift-diffusion equation.

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

When all cargo starts in one location (i.e. the initial condition is a delta function), then the distribution after some elapsed time will be a Gaussian function. The mean of this Gaussian function changes in direct proportion to the drift coefficient over time, while the variance is inversely proportional to the diffusion coefficient.

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 (i.e. assuming *ergodicity*), the ensemble of cargo will approach a Gaussian distribution by the Central Limit Theorem. This means that the mean and variance calculations for a single particle can be directly related to the ensemble statistics outlined in the previous paragraph. We find:

$$a = \frac{2p_+ - (p_+ - p_-)^2}{2}$$
$$b = \frac{2p_- - (p_+ - p_-)^2}{2} \qquad (460)$$

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.

Fitting local parameters of the mass-action model

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 simple, and well-known, approach to doing 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 coefficient. 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 (Caspi et al., 2000). Diffusion appropriately models particle motion 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.

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. A similar procedure (starting at the tips and moving in) can be applied to branched morphologies to prove the more general condition.

It is helpful to re-express the mass-action trafficking model as a matrix differential equation, $\dot{\mathbf{u}} = A\mathbf{u}$, where $\mathbf{u} = [u_1, u_2, ... u_N]^T$ 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:

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

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 our 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), which is available at: https://github.com/ahwillia/PyNeuron-Toolbox.

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 (e.g. mitochondria) 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.

The mass-action model nevertheless provides insight into these systems if we shift our perspective. Instead of interpreting u_i as the amount of cargo in compartment i, we can interpret u_i as the *probability* of a particle occupying compartment i at a particular time. This perspective evokes 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 mirrors our steady-state analysis of the deterministic model.

In the stochastic model, the number of molecules in each compartment follows 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 that there are 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 our 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 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.

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.

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.

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.

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.

Laßek, M., Weingarten, J., and Volknandt, W. (2015). The synaptic proteome. *Cell Tissue Res.*, 359(1):255–265.

Liu, X.-A., Kadakkuzha, B., Pascal, B., Steckler, C., Akhmedov, K., Yan, L., Chalmers, M., and Puthanveettil, S. V. (2014). New approach to capture and characterize synaptic proteome. *Proc. Natl. Acad. Sci. U.S.A.*, 111(45):16154–16159.

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.

Muller, 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. and Poo, M.-m. (2013). Neurotrophin regulation of neural circuit development and function. *Nat. Rev. Neurosci.*, 14(1):7–23.

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.

Parrish, J. Z., Emoto, K., Kim, M. D., and Jan, Y. N. (2007). Mechanisms that regulate establishment, maintenance, and remodeling of dendritic fields. *Annu. Rev. Neurosci.*, 30:399–423.

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.

Wong, R. O. and Ghosh, A. (2002). Activity-dependent regulation of dendritic growth and patterning. *Nat. Rev. Neurosci.*, 3(10):803–812.

Zagrebelsky, M. and Korte, M. (2014). Form follows function: BDNF and its involvement in sculpting the function and structure of synapses. *Neuropharmacol.*, 76:628–638.

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.