

1 **KymoButler: A deep learning software**
2 **for automated kymograph analysis**

3

4 Maximilian A. H. Jakobs*, Andrea Dimitracopoulos, Kristian Franze*

5

6 Department of Physiology, Development and Neuroscience, University of Cambridge,
7 Cambridge, UK

8

9 * To whom correspondence should be addressed: mi455@cam.ac.uk (M.A.H.J.) or
10 kf284@cam.ac.uk (K.F.)

11

12 **Abstract**

13 Knowledge about the dynamics of cytoskeletal proteins is key to understanding numerous
14 active cellular processes. However, quantifying cytoskeletal dynamics is challenging. Current
15 tracking algorithms often require human supervision and are less accurate than manual
16 analysis, which on the other hand is time-consuming and prone to unconscious bias. We
17 here developed and trained KymoButler, a deep neural network to trace dynamic processes
18 in kymographs, which are graphical representations of spatial position over time. We
19 demonstrate that KymoButler performs at least as well as manual tracking and outperforms
20 currently available automated tracking packages. Additionally, we successfully applied
21 KymoButler to a variety of different kymograph tracing problems. Finally, the network was
22 packaged in a web-based "one-click" software for use by the wider scientific community. Our
23 approach significantly speeds up data analysis, avoids unconscious bias, and represents a
24 step towards the widespread adaptation of Artificial Intelligence techniques in biological data
25 analysis.

26

27 Introduction

28 In eukaryotic cells, biopolymers such as microtubules and actin filaments (F-actin) provide
29 structural support and enable essential cellular functions including intracellular transport ^{1,2},
30 cell motility ³, and cell division ^{4,5}.

31 Both microtubules and F-actin are polar filaments with a +end and a –end which differ in
32 their chemical and dynamical properties. Microtubules, for example, exhibit a mostly stable -
33 end, while the +end undergoes rapid cycles of growth and shrinkage ⁶. Measurements of
34 microtubule dynamics are usually performed by genetically expressing fluorescent proteins
35 that preferentially bind to the filament ends, such as the +End-Binding protein 1 (EB1) ^{7,8}.
36 These fluorescent proteins (particles) are recorded using time-lapse fluorescence
37 microscopy and tracked with a variety of approaches.

38 Since actin and microtubules can only grow along their own axis, it is possible to visualise
39 and simplify filament end tracking by using kymographs ^{9,10} - 2D images generated by
40 stacking the intensity profile along a given line, e.g. the F-actin or microtubule axis, for each
41 time point of a movie. Thus, kymographs are length-time images showing labelled filament
42 ends as lines (**Fig. 1**). They are not limited to tracking cytoskeletal filaments but have been
43 widely employed to visualise biological processes across different length scales, ranging
44 from single molecule to cell tracking ^{11,12}.

45 Kymographs provide an elegant solution to the visualisation and quantification of particle
46 dynamics. In contrast to most currently available tracking software, which faces the difficult
47 computational problem of identifying corresponding particles in different frames, a
48 kymograph visualises this problem, and only requires the tracing of lines in an image, a
49 much simpler task for humans and machines alike. These lines then represent the track of a
50 filament, or any other process, so that measuring the lines' lengths and slopes allows to
51 calculate the average velocities and growth periods of a cytoskeletal filament, respectively.

52 Conventional kymograph tracing or particle tracking algorithms produce acceptable results
53 when applied to images with a high signal-to-noise ratio (SNR), but are exceedingly error-
54 prone at lower SNRs^{10,13}. While immunocytochemical stains may result in high quality
55 images with high SNR, live-cell imaging as required for the investigation of dynamic
56 processes usually suffers from autofluorescence, limited light exposure, and the low labelling
57 densities required to keep the cells undamaged. The resulting lower quality images often
58 require cumbersome manual error corrections, leading to similar time commitments as an
59 exclusively manual analysis. Thus, the problem of automatically, and reliably, tracking
60 dynamic processes in live cells is still largely unresolved, and any automation in kymograph
61 tracing that preserves the accuracy of manual annotation would represent a significant
62 improvement in the experimental workflow.

63 In recent years, Artificial Intelligence (AI), and particularly Deep Neural Networks, have been
64 very successfully introduced to data processing in biology^{14,15}. AI-based image analysis has
65 several advantages over other approaches: it is less biased than human users, takes a
66 shorter time to analyse immense datasets, and most importantly, comes closer to human
67 performance than conventional tracking algorithms¹⁴. Most AI approaches to image analysis
68 utilise Fully Convolutional Deep Neural Networks (FCNs) that were shown to excel at object
69 detection in images¹⁶⁻¹⁸. A convolutional neural network is able to use a multitude of hidden
70 layers to apply kernels of all shapes and sizes to images, filtering the information from the
71 background. This ability should, in theory, enable an FCN to trace biopolymer dynamics in
72 low SNR kymographs with unmatched precision.

73 Here we developed a stand-alone software, 'KymoButler', which is based on an FCN, to
74 automatically and reliably extract biopolymer dynamics from kymographs. Whilst we trained
75 the FCN on microtubule +end growth dynamics using manually traced kymographs of EB1-
76 GFP in neurons, the KymoButler software performs well on kymograph data of cytoskeletal
77 filaments in other cells, including EB3-GFP traces from mitotic HeLa cells and actin speckles

78 in *Aplysia* neuronal growth cones. Finally, the KymoButler outperforms conventional
79 automated tracking and, quite remarkably, several cases of manual tracing.

80 **Results**

81 Generation of training data, neural net training, and validation

82 Microtubules constitute a prevalent fraction of the filaments contained in growing neuronal
83 axons¹⁹. To generate kymographs capturing microtubule filament dynamics, we cultured
84 neurons dissociated from *Drosophila melanogaster* larvae, expressing EB1-GFP under the
85 endogenous *eb1* promoter, and tracked the dynamics of EB1 puncta in 520 axons (**Fig. 1A**).

86 In this model system, EB1-GFP puncta move in the axon either towards the cell body
87 (retrograde) or away from the cell body (anterograde). We generated kymographs by
88 manually tracing the axon and stacking the intensity profile along the axon for each frame
89 into one image (**Fig. 1B-C**). In these kymographs, individual EB1-GFP trajectories are
90 visually distinguishable as bright lines. We traced these trajectories by hand and colour-
91 coded them by directionality (anterograde or retrograde, **Fig. 1D**), creating a dataset of input
92 images (the raw kymographs) and labels (the traces).

93 We then used these pairs of input-label images to train an FCN to separate pixels belonging
94 to an EB1-GFP trace from background pixels. We built a custom neural network based on
95 Google's "inception" architecture, the *Tracer FCN*¹⁷ (**Methods** and **Figure 1–figure**
96 **supplement 1**). Additionally, we designed a much faster, shallower FCN that only takes a
97 10% of the evaluation time of the Tracer FCN while maintaining similar levels of performance
98 in our system (**Figure 1–figure supplement 2**). Both FCNs take the input kymograph and
99 decompose it into several images, called feature-maps, through numerous convolution and
100 deconvolution steps. The final output is an image of the same size as the input image, in
101 which each pixel value corresponds to the probability p of this pixel being part of the
102 foreground (part of a trace). The nets were trained to recognise traces going from the left to

103 the right. Applying them to the original and the vertical mirror image allows to distinguish
104 between anterograde and retrograde traces, respectively.

105 We split our dataset into a validation set and a training set, by randomly selecting two
106 biological repeats with a total of 33 (~6%) kymographs as validation data. The training
107 dataset was used to iteratively change the FCN parameters to match the FCN output to the
108 manually traced lines (see **Methods**). This was done by minimising loss (a function that
109 quantifies the difference between the desired image and the FCN output) through stochastic
110 gradient descent and changing the network's parameters accordingly. The training of the
111 FCN stops when changing the parameters does not lead to any further decrease of the loss
112 (**Figure 1–figure supplement 2**). The validation data set was simultaneously used to
113 quantify how the FCN performed using a previously unseen dataset.

114 Trained FCNs assign the probability of being part of a trace to each pixel in the input image
115 (**Fig. 1F**). To convert these probability maps into tracks and compare them to the manual
116 data, we introduced a threshold value t : any pixel that had a larger value than t was
117 classified as being part of a track. The resulting binary image was then iteratively thinned so
118 that only traces with a width of one pixel remained, which was subsequently overlaid on the
119 manual data for comparison (**Fig. 1G**). The trained Tracer FCN showed a precise overlay
120 with the manual annotation in the validation data (see **Fig. 1H-I**). Often, the Tracer FCN
121 surpassed the accuracy of manual labelling, as it was able to recognise previously
122 unlabelled traces that were erroneously omitted.

123 Next, we quantified the effect that the threshold value t had on the output of the network by
124 introducing a similarity score that accounts for the differences between the output of the
125 Tracer FCN and the manual labels (**Fig. 1J**). A score of 1 would indicate a perfect overlay,
126 while a score of 0 would indicate no matches. For small t (0.01) we observed frequent
127 artefacts, for example the linking of parallel tracks. For high t (0.5) the predicted tracks were
128 too short. An optimum threshold was found around $t=0.2$ (**Fig. 1J**), which was therefore used

129 throughout this paper unless stated otherwise. The maximum similarity score we achieved
130 was ~0.7. As the KymoButler tends to outperform and detect more traces than identified by
131 the manual labelling (where faint or short traces are often missed), similarity was low (<1)
132 even when automated detection was close to an optimum. These results indicated that a
133 trained FCN is able to automate the kymograph tracing process, significantly reducing
134 research workload and avoiding biased data analysis.

135 The KymoButler software package

136 We packaged the trained FCNs into two easy-to-use interfaces for quick and fully automated
137 kymograph analysis: (1) a browser-based app with the shallow FCN (**Figure 1–figure**
138 **supplement 1**) to quickly drag & drop individual kymographs in order to analyse them
139 (<http://kymobutler.deepmirror.ai>) and (2) a simple command line python script to be used
140 offline with the full Tracer FCN (<https://github.com/MaxJakobs/KymoButler>). While the Tracer
141 FCN is preferable to precisely analyse large or more complex data sets, the web based
142 shallow version can be used to quickly assess the feasibility of the approach with a given
143 kymograph. In both cases, the user first has to generate a kymograph for their specific
144 problem, with any available kymograph generator (for example the Multi Kymograph Fiji
145 plugin (<https://imagej.net>), the KymographTracker package ⁹, or the KymographClear Fiji
146 plugin ¹⁰). The software then applies the FCN to the image twice (once to the original and
147 once to the vertical mirror image), thresholds the result, applies iterative thinning, generates
148 an overlay of predicted tracks onto the kymograph, and finally extracts and classifies each
149 connected line as a single trace (**Fig. 2**). In the software, the user can freely define the
150 threshold parameter t , the probability above which a pixel is considered to be part of a trace.
151 After the computation, which takes approximately 5-10 seconds on a conventional computer
152 (Tracer FCN on a CPU), the KymoButler generates several files including an overlay image
153 highlighting all the tracks found in different colours, and a CSV file per kymograph,
154 containing all track coordinates and track directionality for post-processing.

155 KymoButler outperforms conventional tracking software

156 To assess the performance of KymoButler, we compared it to manual kymograph tracing
157 and to the plusTipTracker package, which was explicitly written for tracking EB1-GFP puncta
158 ¹³. In conventional tracking algorithms such as the plusTipTracker, individual features are
159 first detected through local thresholding and then linked with each other between frames. We
160 compared the average track velocities (start-to-end velocity) and track lengths of EB1-GFP
161 puncta of our validation data set (33 previously 'unseen' kymographs, **Fig. 3**) for all the three
162 approaches. There was no significant difference between the average velocities
163 (KymoButler: $4.6 \pm 1.0 \mu\text{m}/\text{min}$, Manual: $4.3 \pm 0.9 \mu\text{m}/\text{min}$, plusTipTracker: $4.8 \pm$
164 $1.4 \mu\text{m}/\text{min}$, one way ANOVA, $p=0.16$, **Fig. 3A**). However, when plotting the velocities
165 calculated by the two algorithms against manually determined data in a 2D scatter plot, 97%
166 (32/33) of the velocities calculated by KymoButler fell within the standard deviation of the
167 manual data ($\pm 0.9 \mu\text{m}/\text{min}$), while this was only the case for 73% (24/33) of the velocities
168 calculated by plusTipTracker (**Fig. 3B**).

169 The average track lengths revealed by manual tracing, KymoButler, and plusTipTracker
170 differed significantly (**Fig. 3C**, $p < 10^{-23}$, one way ANOVA). A post-hoc analysis showed no
171 differences between KymoButler and manual analysis ($25 \pm 5 \text{ sec}$ and $23 \pm 4 \text{ sec}$, $p=0.16$,
172 Tukey-Kramer test). However, the plusTipTracker analysis significantly underestimated the
173 track times by about twofold ($12 \pm 2 \text{ sec}$, $p < 10^{-9}$, Tukey-Kramer test) (**Fig. 3C**). Additionally,
174 in 85% (28/33) of kymographs analysed with KymoButler, the average lengths of the traces
175 were within the standard deviation of the manual data ($\pm 5 \text{ sec}$), but only 1 out of the 33
176 axons analysed with plusTipTracker fell within the same region (**Fig. 3D**).

177 We noticed that for one kymograph the manual tracing resulted in much larger average EB1-
178 GFP track lengths than calculated by both KymoButler and plusTipTracker (dot 2 in **Fig. 3D**).
179 Revisiting the manual data revealed that several short tracks were unlabelled incorrectly
180 (black box in **Fig. 3F**). Additionally, some tracks were erroneously drawn too long, while

181 KymoButler broke them rightly into several shorter ones (red box in **Fig. 3F**), indicating that
182 KymoButler performs better than manual labelling on most kymographs.

183 KymoButler can be easily extended to other biological systems

184 We finally tested the capability of the KymoButler software to analyse kymographs
185 generated from different cell types and different cytoskeletal components. Note that we did
186 not retrain the Tracer FCN for these applications. First, we analysed time lapse movies of
187 EB3-GFP dynamics in interphase HeLa cells (**Fig. 4A**). After only changing the threshold
188 parameter to $t=0.1$, KymoButler predicted puncta trajectories as well as it did for *Drosophila*
189 *melanogaster* axon EB1-GFP. When comparing manually extracted traces with KymoButler
190 results of raw kymograph images, we did not find any significant differences between
191 average EB3-GFP microtubule growth velocities (Wilcoxon rank sum test, $p=0.98$) and
192 average growth times (Wilcoxon rank sum test, $p=0.61$) (**Fig. 4B**).

193 Remarkably, KymoButler was even able to quantify actin speckle velocities in *Aplysia* growth
194 cones. Average retrograde actin flow velocities showed no significant difference between
195 manual labelling and KymoButler analysis even though the network was only trained on
196 EB1-GFP puncta in axons (Wilcoxon rank sum test, $p=0.08$) (**Fig. 4D**).

197 Discussion

198 In this work, we used deep learning to optimise automated tracking of dynamic, fluorescently
199 labelled proteins in a noisy cellular environment. Fully convolutional neural networks (Tracer
200 FCNs) are nowadays widely applied for image recognition. Since tracking is *a priori* a visual
201 problem, we built an FCN for identifying traces in kymographs. We deployed our network in
202 two independent stand-alone software packages that take generic kymographs and output
203 all traces found in the image in a matter of seconds. Remarkably, the network not only
204 outperforms current particle tracking software and, in some cases, even manual tracking, but

205 it also performs just as well on kymographs of different dynamic processes, such as
206 fluorescence speckle microscopy.

207 Our KymoButler software has only one adjustable parameter: t , the threshold at which a
208 pixel is recognised as being part of a track. For our validation data, the best value for t was
209 0.2. This threshold generally depends on the SNR of the image. If the SNR is low, the FCN
210 is “less confident” about a given pixel, so that the threshold has to be smaller. More noisy
211 data, such as the HeLa cell EB3-GFP data or actin speckles shown in Figure 4, produced
212 good results with a smaller threshold value ($t=0.1$). Hence, the correct threshold has to be
213 chosen based on each biological application and imaging conditions.

214 Available automated kymograph analysis software was not suitable for tracing EB1-GFP
215 puncta in axons, mainly because these packages were susceptible to noise. The
216 KymographDirect package, for example, applies a global threshold to individual kymographs
217 to extract traces, thus being very prone to variations in background intensity and requiring
218 manual screening ¹⁰. Most other currently available packages require manual track tracing or
219 linking, defeating the purpose of a fully automated analysis ^{9,20}. An alternative approach
220 quantifies kymograph velocities through 2D autocorrelation, however, the analysis is limited
221 as trace lengths cannot be measured ²¹.

222 The current gold standard for automated tracking of microtubule dynamics is the
223 plusTipTracker package. When we compared KymoButler with manual and plusTipTracker
224 data, it performed at least as well as manual tracking, and much better than the
225 plusTipTracker. The mismatch between the plusTipTracker and manual traces is likely
226 because (1) “long” tracks have a tendency of being split into several shorter ones, since the
227 probability of linking errors increases with track length (**Supplementary Movie 1**), and (2)
228 “short” tracks are sometimes incorrectly linked due to background fluctuations
229 (**Supplementary Movie 2**). The first issue results in too short track lengths, and the second
230 causes inflated velocity measurements.

231 We propose that manual tracking is inferior to the KymoButler as it suffers from
232 inconsistency, bias, and is overall laborious. While the KymoButler analyses each
233 kymograph in the same way, manual tracing performance varies from one kymograph to the
234 next as well as between users. In our dataset, we frequently overestimated trace lengths, so
235 that the manual annotation yielded slightly larger track lengths than the KymoButler. In
236 future, KymoButler could be trained on a larger dataset traced by multiple researchers to
237 remove other inconsistencies that may be present in the dataset, thus further improving the
238 KymoButler's performance.

239 Additionally, KymoButler was able to analyse kymographs from different dynamic processes
240 such as retrograde actin flow in neuronal growth cones. This result highlights that particle
241 tracking does not depend on the precise nature of the particle, e.g. actin or EB1, but on the
242 task of tracing a line in an image, which should be the same for any dynamic process that
243 can be represented this way.

244 Future work will expand our approach to 2D or even 3D tracking problems. In this paper, we
245 drew 1D lines in 2D movies, extracted 2D (space and time) images (kymographs), and finally
246 traced 2D lines in those images. A similar, albeit computationally heavier, approach could
247 stack the frames of a 2D/3D movie on top of each other to generate a 3D/4D image (2D
248 space and time, or 3D space and time). The 2D/3D lines in those images can then be traced
249 by hand and a more complex FCN trained to recognise them. This approach could yield
250 human-like performance in higher dimensional automated tracking.

251

252 **online Methods**

253 Fly Stocks

254 The following stocks were used for expressing fluorescently tagged EB1: *eb1-gfp*²² and
255 *uas:eb1-gfp*²³. To include different genetic backgrounds in our training data we also co-
256 expressed two RNAi constructs: *uas:wh-RNAi* (Bloom# 35573) and *uas:dhcRNAi* (Bloom#
257 36698) of which the latter is known to cause a severe phenotype on EB1-GFP dynamics²⁴.
258 All *uas* constructs were driven by *elav-gal4* (Bloom# 458) and transgenic lines generated
259 through standard balancer crossing procedures.

260 *D. melanogaster* neuronal culture and EB1-GFP live imaging

261 Primary cell cultures were prepared similar as to²⁵. Third instar larvae were selected, and
262 their central nervous systems dissected. Subsequently, the tissue was dissociated in Hank's
263 Balanced Salt Solution (HBSS) supplemented with Dispase (Roche 049404942078001) and
264 Collagenase (Worthington Biochem. LS004214). The cells were plated in 30µl droplets of
265 Schneider's Medium (Thermo Fisher 21720024) supplemented with insulin (2 µg/ml Sigma
266 I0516) and fetal bovine serum (1:4 Thermo Fisher Scientific A3160801). We plated the drops
267 in ibidi glass-bottom µDishes (cat num 81158) and covered them with 25mm coverslips
268 (VWR) to create small culture chambers. The glass bottom dishes were previously coated
269 with Concanavalin A (5µg/ml, 1.5h at 37°C). The culture chambers were subsequently put at
270 26°C for 1.5h so that the cells settle on the coated surface of the dish. Then the chambers
271 were flipped to remove debris from the surface and left for 24 hours before imaging.

272 Live imaging movies were acquired on a Leica DMI8 inverted microscope at 63x
273 magnification and 26°C (oil immersion, NA=1.4). To reduce autofluorescence the culture
274 medium was replaced with Live Imaging Solution (Thermo Fisher A14291DJ). For EB1-GFP

275 imaging, an image was taken every 2 seconds for 70-150 frames depending on sample
276 bleaching rate. We imaged 520 axons from 26 different dishes.

277 We also treated the cells with Latrunculin B (10 μ M) and Ciliobrevin A (100 μ M). Both drugs
278 are known to perturb microtubule dynamics so that including movies acquired with these
279 treatments would again make our FCN more robust^{24,26}. In both cases the cells were first
280 allowed to attach to the coated glass for 1.5h post dissection before replacing the culture
281 medium with culture medium supplemented with Latrunculin B or Ciliobrevin A.

282 *Aplysia* neuronal culture and actin fluorescence speckle microscopy

283 *Aplysia* bag cell neurons were isolated and cultured as previously described in²⁷. Neurons
284 were then injected with alexa-568 labelled G-actin (Molecular Probes) at low levels,
285 appropriate for fluorescence speckle microscopy²⁸. The growth cone in Fig. 4B was imaged
286 on a spinning disk confocal microscope at 2 Hz sampling rate.

287 *HeLa* Cell culture and imaging

288 A HeLa stable cell line expressing LifeAct-GFP and EB3-mRFP²⁹, was maintained in
289 Dulbecco's Modified Eagles Medium (DMEM GlutaMAX; Gibco) supplemented with 10%
290 FBS and 50 U/ml penicillin and 50 μ g/ml streptomycin (Invitrogen) at 37 C under 5% CO₂.
291 Cells were imaged using an UltraView Vox (Perkin Elmer) spinning disc confocal microscope
292 with a 63X (NA 1.4) oil objective equipped with temperature and CO₂ controlling
293 environmental chambers and images were acquired using a Hamamatsu ImagEM camera
294 and Volocity software at a rate of 2 Hz (Perkin Elmer).

295 Kymograph generation and FCN training

296 The 520 neuronal axons were first traced by hand with the KymographTracker plugin for ICY
297 (<http://icy.bioimageanalysis.org>,⁹). We randomly choose two biological repeats (2x dishes,

298 33 axons, ~6%) as a validation data set, i.e. we did only use 489 axons as training data.
299 Subsequently we generated kymographs with the KymographTracker plugin and traced
300 EB1-GFP lines in those images by hand, using the same plugin. The traces were then
301 plotted in two images: one for retrograde tracks and one for anterograde tracks. We also
302 generated kymographs with a custom Mathematica script to obtain two slightly different
303 kymographs per axon. We then reflected each kymograph and the corresponding trace
304 images along the vertical (y) axis and stretched them along the x-axis to 0.5, 0.75, 1.25, and
305 1.5 their original length eventually resulting in a total number of 10,400 kymographs and their
306 respective manually traced images (two per kymograph). Hence our training/validation data
307 set comprises 9740/660 kymographs and their respective trace images.

308 We decided to design a Fully Convolutional Neural Network (FCN) to recognise the antero-
309 and retrograde lines in our noisy kymographs. An FCN does not exhibit any fully connected
310 layers, i.e. layers whose parameter number depends on the dimension of the input image,
311 but only calculates several parallel and consecutive image convolutions and/or
312 deconvolutions with trainable parameters. As the number of these parameters does not
313 depend on the size of the input image, kymographs do not have to be resized before
314 application of the FCN.

315 We used *Mathematica* (<http://wolfram.com>) to both generate and train our FCN. Even though
316 the network is fully convolutional, the *Mathematica* training algorithm needed all input
317 images to have the same dimensions. Thus, we divided each kymograph into tiles of 80x80
318 pixels so that one training “unit” comprised one input image and two output images, showing
319 anterograde and retrograde traces. To make training more efficient, we decided to only train
320 one network to recognise anterograde (left to right) tracks so that each of these sets was
321 again split into an input tile with the anterograde tracks and the vertically reflected input +
322 retrograde tile. The total number of tile pairs thus became 149,488 for the training data and
323 9740 for the validation data. In this way the final network would have to be called twice: once

324 on the original kymograph and once on the reflected one to detect both antero- and
325 retrograde traces.

326 Our approach to the precise architecture of the final Tracer FCN was purely empirical
327 comprising the following building blocks: (i) a convolutional layer with arbitrary kernel size
328 and number of output channels followed by a batch normalisation layer and a ‘leaky’ ramp
329 (leayReLU) activation function ($leayReLU(x) := \max(x, 0) - 0.1 \max(-x, 0)$), (ii) a dropout
330 layer that randomly sets 10% of all input values to zero during training to prevent ‘overfitting’
331 of the input data, (iii) a deconvolutional layer with arbitrary kernel size and number of output
332 channels to sharpen the input images again followed by a batch normalisation layer and a
333 leayReLU layer, (iv) a pooling layer with kernel size three to replace a given pixel with the
334 maximum value in its neighbourhood. The batch normalisation layer is useful to stabilize the
335 training procedure as it rescales inputs to the activation function (leayReLU) so that they
336 have zero mean and unit variance. The leayReLU prevents so-called dead ReLU’s by
337 applying a small gradient to values below 0. These building blocks were previously used for
338 image recognition tasks in *Google’s* inception architecture ¹⁷.

339 The architectures we settled on is shown in **Figure 1–figure supplement 1**. Six connected
340 “Trace Block” layers are used to denoise the image and highlight individual traces. The
341 precise architecture of these Blocks is again shown in **Figure 1–figure supplement 1**. This
342 block architecture allows a lot of flexibility with the choice of operation, for example the
343 convolving kernel size, throughout training and evaluation. A major feature of the Trace
344 Block architecture is the inclusion of deconvolutions. Without explicitly computing
345 deconvolutions in each block, as for example in the shallow FCN in **Figure 1–figure**
346 **supplement 1**, the final image is more blurred, and one is unable to segment individual
347 traces as efficiently. In the final step of both architectures all channels are projected on only
348 two and a softmax layer is applied so that the sum over those channels is one for each pixel.
349 The two channels can be interpreted as the probability of a given pixel to be part of the
350 background or a trace.

351 To train the FCN we quantified the difference between the FCN output o and the desired
352 target output t through a cross entropy loss layer ($CEloss(t, o) = -(t \cdot \ln(o) + (1 - t) \cdot$
353 $\ln(1 - o))$). Here t can be either 1 (background) or 2 (trace). For Example: The untrained
354 FCN will give 0.5 as the probability of each pixel to be part of the background as it has no
355 preference yet. The corresponding loss for a pixel that should be part of the background
356 (index=1) would be: $CEloss(0.5, 1) = 0.69$. During training this value might be updated to 0.9
357 decreasing the loss to $CEloss(0.9, 1) = 0.11$.

358 We trained the FCN through stochastic gradient descent. Here we first randomly subdivided
359 all training tile pairs into batches of 50. For each batch we then calculated the average cross
360 entropy loss and the gradient of this loss in all tuneable parameters, e.g. the kernel entries in
361 the convolutions. We then updated all the parameters σ in the network according to $\sigma' = \sigma -$
362 $\eta \partial_{\sigma} CEloss(t, o)$. Here ∂_{σ} denotes the partial derivative with respect to all parameters of the
363 FCN and η is the learning rate, i.e. the multiplier giving absolute value of the shift in σ at a
364 given step. Note that η is not fixed but is dynamically updated through the ADAM algorithm
365 ³⁰. This was repeated for all batches until the whole training dataset was visited by the
366 algorithm constituting one round. The FCN was trained until no decrease in the validation
367 data loss was observed anymore (5 Rounds). Every 10 minutes, the average loss was
368 calculated for the validation dataset to obtain a readout on how the FCN performs on
369 previously “unseen” data.

370 FCN performance evaluation

371 The direct output of both FCNs was an 80x80x2 tensor that assigns each pixel the
372 probability of being part of a trace (index=2) or the background (index=1). In order to
373 reconstitute traces from the FCN output we introduced a threshold value t for the second
374 index, above which we would consider a pixel being part of a trace. The training set
375 comprises many more background pixels than foreground pixels so that the FCN exhibits
376 small probabilities around traces, therefore the cut-off has to be chosen generally as an

377 unintuitively small value ($t < 0.5$). The thresholded output images were iteratively thinned until
378 they depicted lines of only one pixel wide.

379 To compare the FCN output with the manual annotation for the validation data we defined a
380 similarity score as a function of the threshold as follows: (i) Both the anterograde and the
381 retrograde trace probability map are calculated with the FCN and thresholded and dilated by
382 one pixel. (ii) Both dilated binary predictions (0=background, 1=trace) are multiplied with the
383 respective binary manual trace images and in the resulting image the total number of
384 pixels=1 counted (*ovlp*, a measure of the overlap between the prediction and the manual
385 annotation). (iii) We also calculated the total number of pixels=1 in the manual traced image
386 (N_m) and the prediction (N_p). (iv) The similarity score s was then given by:

387

$$s = \begin{cases} 1, & \text{if } N_m = 0 \ \& \ N_p = 0 \\ 1/N_m, & \text{elseif } N_p = 0 \\ 1/N_p, & \text{elseif } N_m = 0 \\ \frac{ovlp}{N_m(1 + \frac{|N_m - N_p|}{N_m})}, & \text{else} \end{cases}$$

388

389 In short: The similarity score measures the overlapping pixels in the prediction and the
390 manual annotation and divides them by the absolute number of pixels being part of a trace in
391 the manual annotation ($ovlp/N_m$). The result is divided by a factor measuring the difference
392 in pixels that are part of a trace between prediction and manual labelling to penalise large
393 discrepancies in total number of predicted pixels ($1 + |N_m - N_p|/N_m$). Since the prediction
394 rarely overlaps completely with the manual annotation and frequently finds more objects that
395 were previously labelled, a 'good' score lies at around 0.7.

396 KymoButler software

397 The KymoButler software first applies either the deep Tracer FCN or the shallow FCN to a
398 given kymograph and its vertical reflection. The resulting foreground probability map is then

399 thresholded with the parameter t and thinned iteratively so that each trace is only one pixel
400 wide at any point. The thinned traces are then pruned by three pixels so that short branches
401 are deleted. Subsequently, each trace is segmented and selected only if it contains more
402 than 5 pixels and is at least 3 frames long. This step removes noise from the result. In the
403 final step, pixels that lie in the same row of the kymograph are averaged over so that the
404 resulting track has only one entry per frame.

405

406 Comparison between KymoButler and plusTipTracker

407 We used the plusTipTracker version 1.1.4 for *MATLAB* 2014a (mathworks.com) to analyse
408 the axons from our validation dataset (33 axons). In each movie we first selected a region of
409 interest comprising the axon and omitting very bright artefacts. To run the software, we first
410 varied the detection parameters to find those that result in similar total track numbers as the
411 manual kymograph tracing approach. We settled on the following detection parameters: $\sigma_1 =$
412 $1, \sigma_2 = 4, K = 8$. For tracking we chose: $minTrackLength = 3, maxGap = 2,$
413 $minSearchRad = 5, maxSearchRad = 15, maxFwAngle = 30, maxBwAngle = 10, shrinkV =$
414 $0,$ and $rFluc = 1.5$. Note that we set the shrinkage velocity to zero so that the plusTipTracker
415 does not try to calculate microtubule shrinkage events.

416 In order to compare the plusTipTracker to the KymoButler we wrote a short *Mathematica*
417 script that calculates the predicted tracks for the same 33 axons with the Tracer FCN and
418 exports them in a *MATLAB* friendly format. As with the plusTipTracker we ignored all traces
419 with track lengths below 3 frames. All subsequent data plotting and analysis was done in
420 *MATLAB*.

421

422 **Software**

423 Quick and easy cloud platform (Shallow FCN only): <http://www.kymobutler.deepmirror.ai>

424 GitHub with the command line interface (full Tracer FCN):

425 <https://github.com/MaxJakobs/KymoButler>

426 **Acknowledgements**

427 We would like to thank Paul Forscher for providing speckle microscopy time-lapse movies of

428 *Aplysia* growth cones, Eva Pillai for scientific input, proofreading, and logo design, Hannes

429 Harbrecht for fruitful discussions about FCN's; Hendrik Schuermann for help with kymograph

430 tracing; and the Mathematica stack exchange community

431 (<https://mathematica.stackexchange.com>) without whom this project would have taken

432 several decades longer. The authors acknowledge funding by the Wellcome Trust (Research

433 Grant 109145/Z/15/Z to M.A.H.J.), the Herchel Smith Foundation (Fellowship to A.D.), Isaac

434 Newton Trust (Research Grant 17.24(p) to K.F.), UK BBSRC (Research Project Grant

435 BB/N006402/1 to K.F.), and the ERC (Consolidator Award 772426 to K.F.).

436 **Author Contributions**

437 M.A.H.J. and K.F. conceived the project; M.A.H.J. and A.D. performed experiments;

438 M.A.H.J. and A.D. developed the software; M.A.H.J., A.D., and K.F. managed the project;

439 M.A.H.J., A.D., and K.F. wrote the manuscript.

440 **Competing Interests**

441 The authors declare no competing interests.

442

443 References

- 444 1. Franker, M. A. M. & Hoogenraad, C. C. Microtubule-based transport - basic
445 mechanisms, traffic rules and role in neurological pathogenesis. *J Cell Sci* **126**, 2319–
446 2329 (2013).
- 447 2. Mitchison, T. J. & Cramer, L. P. Actin-based cell motility and cell locomotion. *Cell* **84**,
448 371–379 (1996).
- 449 3. Gardel, M. L., Schneider, I. C., Yvonne Aratyn-Schaus & Waterman, C. M. Mechanical
450 Integration of Actin and Adhesion Dynamics in Cell Migration.
451 <http://dx.doi.org/10.1146/annurev.cellbio.011209.122036> **26**, 315–333 (2010).
- 452 4. Prosser, S. L. & Pelletier, L. Mitotic spindle assembly in animal cells: a fine balancing
453 act. *Nat. Rev. Mol. Cell Biol.* **18**, 187–201 (2017).
- 454 5. Lancaster, O. M. *et al.* Mitotic Rounding Alters Cell Geometry to Ensure Efficient
455 Bipolar Spindle Formation. *Dev. Cell* **25**, 270–283 (2013).
- 456 6. Brouhard, G. J. Dynamic instability 30 years later: complexities in microtubule growth
457 and catastrophe. *Mol. Biol. Cell* **26**, 1207–1210 (2015).
- 458 7. Piehl, M., Tulu, U. S., Wadsworth, P. & Cassimeris, L. Centrosome maturation:
459 measurement of microtubule nucleation throughout the cell cycle by using GFP-
460 tagged EB1. *PNAS* **101**, 1584–1588 (2004).
- 461 8. Ma, Y., Shakiryanova, D., Vardya, I. & Popov, S. V. Quantitative Analysis of
462 Microtubule Transport in Growing Nerve Processes. *Current Biology* **14**, 725–730
463 (2004).
- 464 9. Chenouard, N., Buisson, J., Bloch, I., Bastin, P. & Olivo-Marin, J.-C. Curvelet analysis
465 of kymograph for tracking bi-directional particles in fluorescence microscopy images.
466 in 3657–3660 (IEEE, 2010). doi:10.1109/ICIP.2010.5652479
- 467 10. Mangeol, P., Prevo, B. & Peterman, E. J. G. KymographClear and KymographDirect:
468 two tools for the automated quantitative analysis of molecular and cellular dynamics
469 using kymographs. *Mol. Biol. Cell* **27**, 1948–1957 (2016).

- 470 11. Twelvetrees, A. E. *et al.* The Dynamic Localization of Cytoplasmic Dynein in Neurons
471 Is Driven by Kinesin-1. *Neuron* **90**, 1000–1015 (2016).
- 472 12. Barry, D. J., Durkin, C. H., Abella, J. V. & Way, M. Open source software for
473 quantification of cell migration, protrusions, and fluorescence intensities. *J Cell Biol*
474 **209**, 163–180 (2015).
- 475 13. Applegate, K. T. *et al.* plusTipTracker: Quantitative image analysis software for the
476 measurement of microtubule dynamics. *Journal of Structural Biology* **176**, 168–184
477 (2011).
- 478 14. Mathis, A. *et al.* Markerless tracking of user-defined features with deep learning.
479 *arXiv.org cs.CV*, (2018).
- 480 15. Weigert, M. *et al.* Content-Aware Image Restoration: Pushing the Limits of
481 Fluorescence Microscopy. *bioRxiv* 236463 (2017). doi:10.1101/236463
- 482 16. Dai, J., Li, Y., He, K. & Sun, J. R-FCN: Object Detection via Region-based Fully
483 Convolutional Networks. 379–387 (2016).
- 484 17. Szegedy, C. *et al.* Going Deeper with Convolutions. *arXiv.org cs.CV*, (2014).
- 485 18. LeCun, Y. *et al.* Backpropagation Applied to Handwritten Zip Code Recognition.
486 <http://dx.doi.org/10.1162/neco.1989.1.4.541> **1**, 541–551 (2008).
- 487 19. Kapitein, L. C. & Hoogenraad, C. C. Building the Neuronal Microtubule Cytoskeleton.
488 *Neuron* **87**, 492–506 (2015).
- 489 20. Mukherjee, A. *et al.* Automated kymograph analysis for profiling axonal transport of
490 secretory granules. *Medical Image Analysis* **15**, 354–367 (2011).
- 491 21. Chan, C. E. & Odde, D. J. Traction Dynamics of Filopodia on Compliant Substrates.
492 *Science* **322**, 1687–1691 (2008).
- 493 22. Bulgakova, N. A., Grigoriev, I., Yap, A. S., Akhmanova, A. & Brown, N. H. Dynamic
494 microtubules produce an asymmetric E-cadherin-Bazooka complex to maintain
495 segment boundaries. *J Cell Biol* **201**, 887–901 (2013).

- 496 23. Jankovics, F. & Brunner, D. Transiently reorganized microtubules are essential for
497 zippering during dorsal closure in *Drosophila melanogaster*. *Dev. Cell* **11**, 375–385
498 (2006).
- 499 24. del Castillo, U., Winding, M., Lu, W., Gelfand, V. I. & Allan, V. Interplay between
500 kinesin-1 and cortical dynein during axonal outgrowth and microtubule organization in
501 *Drosophila* neurons. *eLife Sciences* **4**, e10140 (2015).
- 502 25. Sanchez-Soriano, N. *et al.* Are dendrites in *Drosophila* homologous to vertebrate
503 dendrites? *Dev. Biol.* **288**, 126–138 (2005).
- 504 26. Rao, A. N. *et al.* Cytoplasmic Dynein Transports Axonal Microtubules in a Polarity-
505 Sorting Manner. *Cell Rep* **19**, 2210–2219 (2017).
- 506 27. Forscher, P., Kaczmarek, L. K., Buchanan, J. A. & Smith, S. J. Cyclic AMP induces
507 changes in distribution and transport of organelles within growth cones of *Aplysia* bag
508 cell neurons. *Journal of Neuroscience* **7**, 3600–3611 (1987).
- 509 28. Danuser, G. & Waterman-Storer, C. M. Quantitative fluorescent speckle Microscopy of
510 cytoskeleton dynamics. *Annu Rev Biophys Biomol Struct* **35**, 361–387 (2006).
- 511 29. Fink, J. *et al.* External forces control mitotic spindle positioning. *Nat Cell Biol* **13**, 771–
512 778 (2011).
- 513 30. Kingma, D. P. & Ba, J. Adam: A Method for Stochastic Optimization. (2014).
- 514
- 515

516 **Figure legends**

517

518 **Figure 1: Generation of kymographs showing microtubule EB1-GFP dynamics and**

519 **subsequent training of the Tracer FCN. (A)** Fluorescence time-lapse images of a

520 drosophila neuron expressing EB1-GFP. A single EB1-GFP punctum is shown in four

521 consecutive frames (arrows). **(B)** Hand-drawn line along the axon building up each pixel row

522 of the kymograph. **(C)** Example kymograph obtained from the line shown in (B). Arrow: track

523 resulting from the EB1-GFP comet shown in (A). **(D)** Individual EB1-GFP traces were traced

524 by hand, distinguished by directionality (blue = anterograde, red = retrograde), and overlaid

525 on the kymograph. **(E)** Example output of the Tracer FCN applied to validation data (see

526 methods). An 80x80 pixel subimage from the kymograph shown in (D) (box) is fed to the

527 Tracer FCN. **(F)** The heat maps show the predicted probability p for each pixel being part of

528 a trace (top: anterograde traces, bottom: retrograde traces). **(G)** Tracer FCN prediction:

529 pixels were considered to be part of a track when $t > 0.2$, and iterative thinning was applied to

530 generate traces. **(H)** Hand-traced (manual) images for both directions. **(I)** the prediction

531 (orange) was overlaid with the manual annotation (blue); co-localised pixels appear pink.

532 The FCN fully recapitulated the hand-traced data and even recognised traces that were

533 omitted by mistake in hand tracings, even though it had never 'seen' this image during

534 training. **(J)** The performance of the Tracer FCN when applied to the whole validation data

535 set in terms of a manual to Tracer FCN similarity score (see methods) plotted as a function

536 of probability cut-offs t . The insets highlight the scores of the anterograde predictions of the

537 kymograph shown in (E). A maximum in similarity is achieved at $t = 0.2$. For larger p cut-off

538 values the network tends to return shorter traces than the manual labelling; for smaller t

539 tracks become incorrectly linked (left inset). Scale bars: 2 μm (horizontal), 25 sec (vertical).

540

541 **Figure 1–figure supplement 1: FCN architecture.** Left: An input 80x80 pixel image is first
542 fed into 2 consecutive Tracer Blocks that each output 110 80x80 images (feature maps). Then
543 a Dropout Layer deletes (randomly) 10% of all pixels in all feature maps (only during training).
544 The result is again computed through four Tracer Blocks. Subsequently, the resulting 110
545 feature maps are projected on two with a 1x1 convolution, the result transposed and a softmax
546 operation applied so that the two entries in each pixel of the 80x80 matrix sum up to 1. The
547 result then comprises two 80x80 images: one whose pixel values give the probability of being
548 part of the foreground (prob fg) and one whose pixel values give the probability of being part
549 of the background (prob bg). Only convolution and deconvolution operations are used, hence
550 the network does not depend on the input image size and can be applied to images that are
551 not 80x80 pixels large. Right Top: One Tracer Block comprises six parallel net chains. (1) the
552 identity convolution 1x1 with 10 output maps. (2) a 1x1 convolution followed by a 3x3
553 convolution with 20 output maps. (3) a 1x1 convolution followed by a 5x5 convolution with 20
554 output maps. (4) a 1x1 convolution followed by a 9x9 convolution with 20 output maps. (5) a
555 1x1 convolution followed by a 3x3 deconvolution with 20 output maps. (6) a 3x3 max pooling
556 operation followed by a 1x1 convolution with 20 output maps. The resulting feature maps are
557 concatenated along the first dimension to generate 110 feature maps as an output of the block.
558 Right Bottom: As this net can be computationally demanding for web form applications and
559 hence expensive to maintain we also designed a shallower FCN: This net does not comprise
560 any parallel blocks and only evaluates one 3x3 convolution followed by a 5x5 convolution and
561 a 3x3 deconvolution.

562

563 **Figure 1–figure supplement 2: Loss Curves for training and validation data.** Top:
564 Validation and batch Loss curves for the Tracer FCN. The FCN was trained for 5 Rounds, i.e.
565 full dataset visitations. 50 input tiles were summed to one batch and the loss calculated on
566 each batch (orange). Additionally, the loss on the validation data set was calculated every 10
567 minutes (blue dots and curve). The loss reaches a plateau after ~4 Rounds. Bottom: The Batch

568 loss of the Tracer FCN (blue, same data as in the orange curve above) and the batch loss for
569 the shallow FCN from Fig. S2.

570

571 **Figure 2: The KymoButler package - an automated software for kymograph analysis.**

572 **(A)** The software is first presented with a kymograph image input of any format. **(B)**

573 Subsequently, the Tracer FCN is applied to the image twice (once to the original and once to
574 the vertical reflection) resulting in two heat maps that assign each pixel the probability of

575 being part of an antero- or retrograde trace (top two panels). Then the software binarizes the

576 resulting images with a user-given threshold t (here $t=0.2$). The binary images are then

577 thinned iteratively, and each line gets segmented as one track (blue and red lines, bottom

578 two panels). **(C)** The software then generates multiple output files: an overlay of the

579 segmented tracks with the original image (shown, each colour represents a distinct track)

580 and a CSV file per kymograph, with every trace's coordinates. Scale bars: 2 μm (horizontal),

581 25 sec (vertical).

582

583 **Figure 3: KymoButler microtubule dynamics analysis outperforms conventional**

584 **tracking algorithms. (A)** Average EB1-GFP velocities per axon were similar for manual

585 tracing, the KymoButler, and plusTipTracker package ($p=0.17$ ANOVA). Each dot represents

586 one axon and the boxplots show the median and the upper and lower quantiles. **(B)** 2D

587 scatterplot of the average velocities calculated with KymoButler (green dots) and

588 plusTipTracker (magenta dots) against the average velocities calculated via manual tracing.

589 Black lines indicate a deviation of $\pm 0.9 \mu\text{m}/\text{sec}$ from the identity line, corresponding to the

590 standard deviation of the manually traced velocities. **(C)** Boxplots of the average track

591 lengths, i.e. the time during which EB1-GFP puncta were visible, calculated with manual

592 tracing, KymoButler, and the plusTipTracker. The average track length was approximately

593 half as long when the plusTipTracker package is used, compared to the manual tracing and

594 KymoButler ($p < 10^{-9}$, Tukey-Kramer test), which yielded similar results. **(D)** 2D scatter plot of
595 the average track lengths calculated with the KymoButler (green dots) and plusTipTracker
596 (magenta dots) against the average track lengths calculated via manual tracing. Black lines
597 again indicate the standard deviation of the manual data. **(E)** Kymograph of data point 1
598 labelled in (D) with overlaid manually labelled traces and the predicted traces of KymoButler
599 (each colour represents one segmented track). There is an excellent correspondence
600 between the tracks obtained by both approaches. **(F)** Kymograph of data point 2 labelled in
601 (D) with overlaid traces. KymoButler breaks up several tracks more accurately than the
602 manual tracking (red box, long trace in the centre, red arrow) and adds several shorter
603 tracks that were incorrectly omitted in the manual approach (black box, black arrow). Only
604 tracks longer than 2 frames were included in the analysis. **(G)** Zoom into the red box shown
605 in (F). Scale bars: $2\mu\text{m}$ (horizontal), 25 sec (vertical).

606

607 **Figure 4: KymoButler efficiently analyses particle tracks in other biological systems.**

608 **(A-B)** Analysis of EB3-GFP in HeLa cells. **(A)** A kymograph was extracted from an
609 interphase HeLa cell expressing EB3-GFP and subsequently analysed by hand and with
610 KymoButler. The heatmap represents the probability map generated by KymoButler, the blue
611 lines correspond to the hand traced EB3-GFP lines, and the coloured lines represent the
612 traces recognised by KymoButler. The threshold t was set to 0.1. Scale bars: $5\mu\text{m}$
613 (horizontal), 10 sec (vertical) **(B)** Average EB3-GFP velocities and growth times obtained by
614 manual tracing and KymoButler analysis. No significant differences were found (Wilcoxon
615 rank sum test, $p=0.98$ velocities, growth times $p=0.61$). **(C-D)** Analysis of actin speckle
616 dynamics in *Aplysia* growth cones. **(C)** Kymograph of fluorescently labelled G-actin, and
617 analysed traces with $t=0.1$. Scale bars: $5\mu\text{m}$ (horizontal), 20 sec (vertical). **(D)** Average actin
618 speckle velocities are similar for manual and KymoButler analysis (test, $p=0.08$). Tracks less
619 than 6 frames long were omitted from the analysis.

620 **Movie legends**

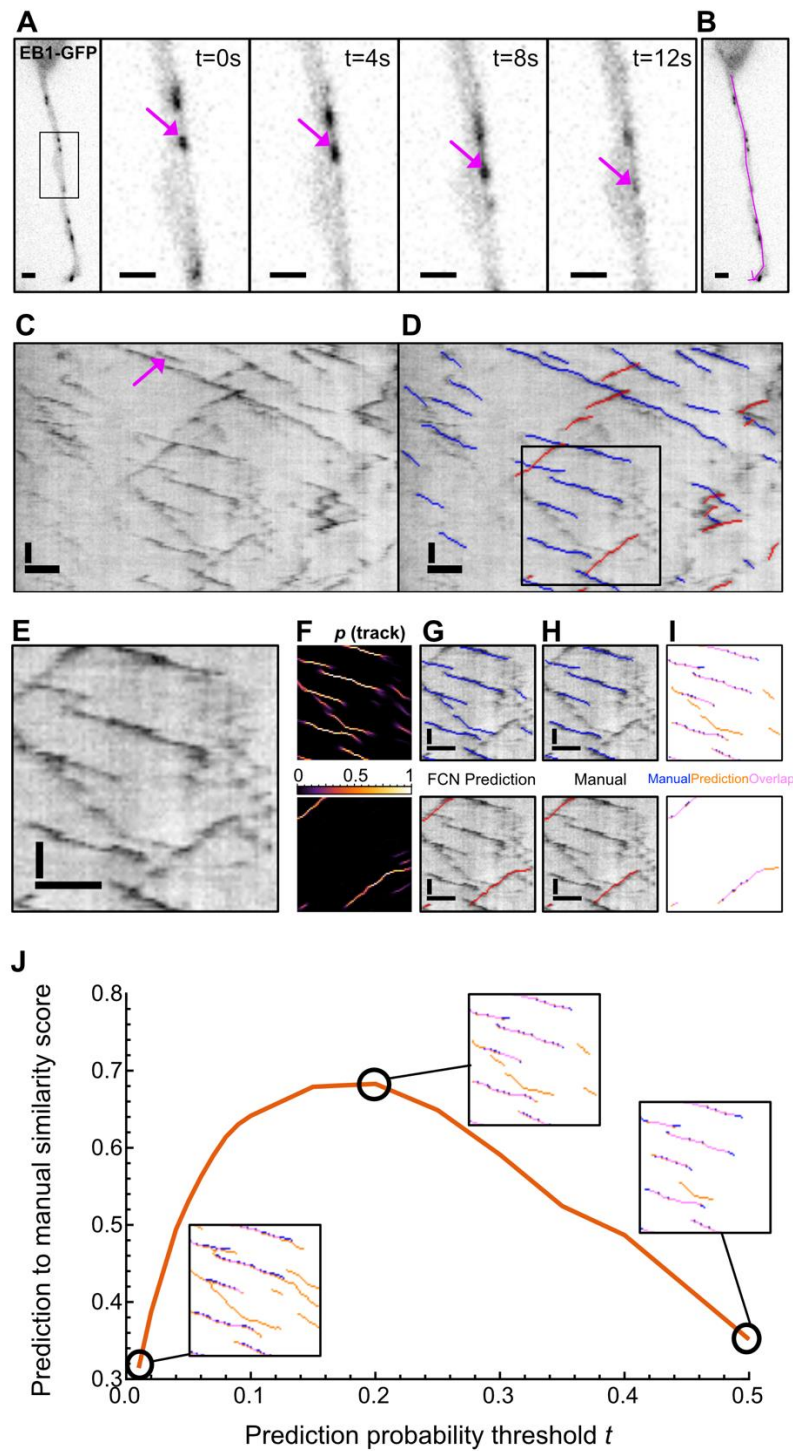
621

622 **Mov. S1: Example of an erroneously shortened EB1-GFP track.** The particle is detected
623 in the lower right corner in frame 1 (small red circle). The particle is then tracked for 7
624 consecutive frames (red line). While the particle does not disappear after frame 7 but rather
625 becomes a bit fainter in frame 8 to re-appear in frame 9 in the upper left corner of the movie,
626 the trace is finished after frame 7. The movie was generated with the plusTipTracker after
627 detection.

628

629 **Mov. S2: Example of an erroneously linked EB1-GFP track.** The particle is detected in the
630 lower right corner in frame 4 (small red circle). In frame 5, the particle moves slightly to the left
631 and gets correctly linked. However, in frame 6, a particle appearing in the upper left corner
632 becomes incorrectly linked to the track, increasing the estimated average velocity of the
633 particle to $\sim 15 \mu\text{m}/\text{min}$, about three-fold larger than the average velocity of EB1-GFP puncta
634 ($5 \mu\text{m}/\text{min}$, **Fig. 3**).

635 **Figures**

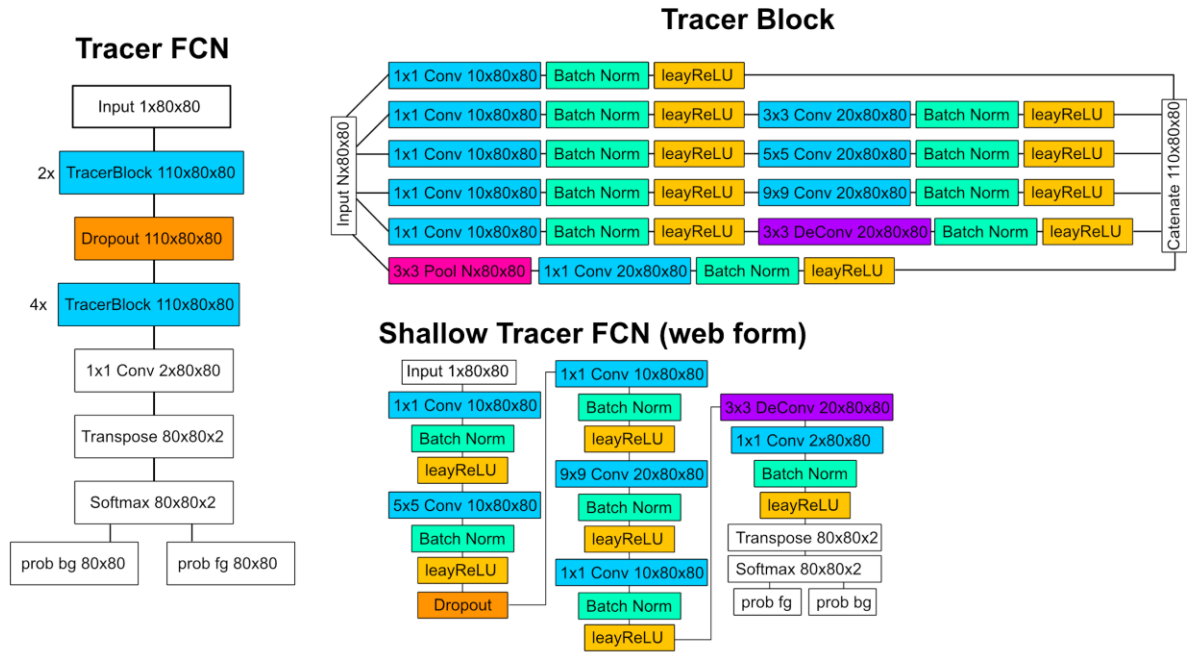


636

637

638

Figure 1

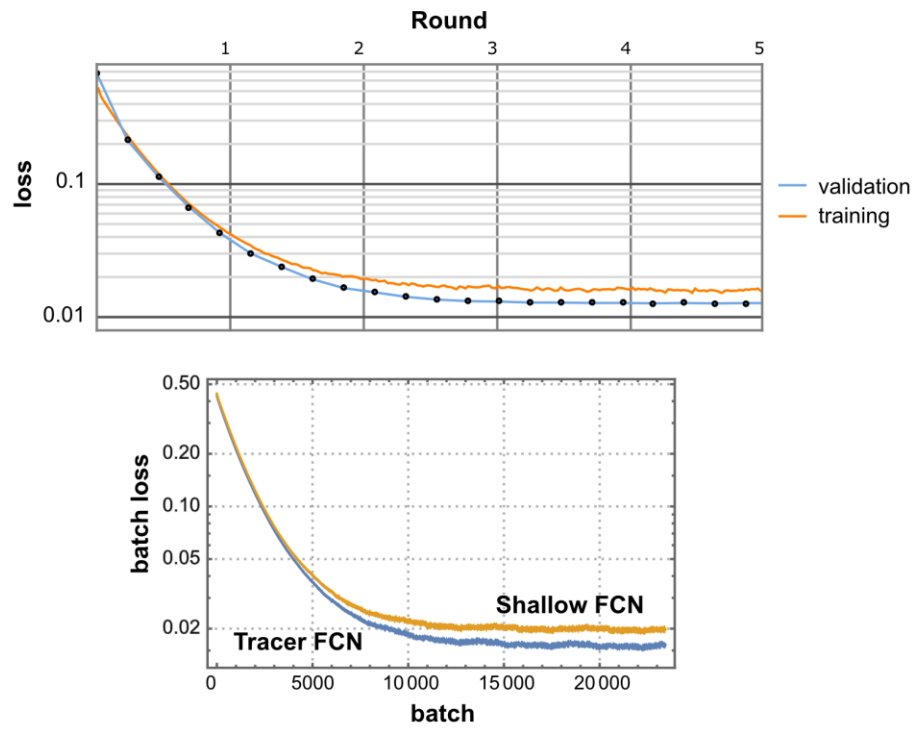


639

640

641

Figure 1–figure supplement 1

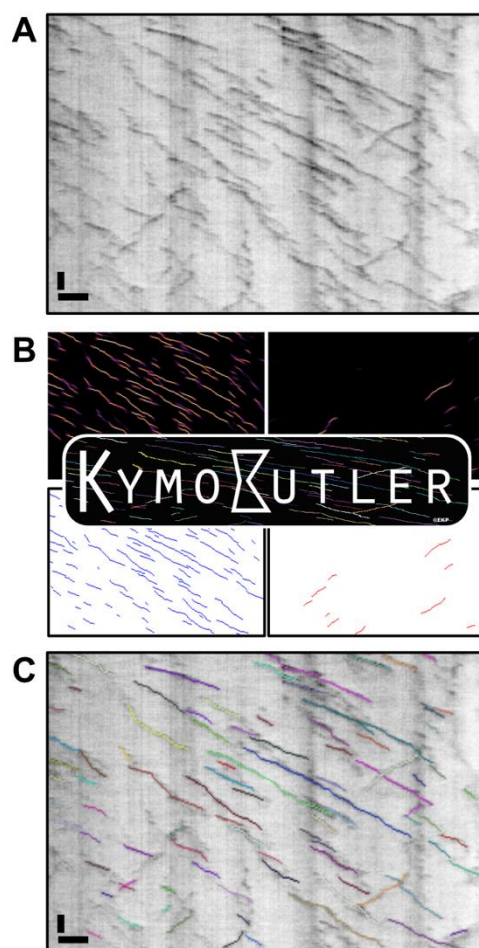


642

643

644

Figure 1–figure supplement 2




645

646

647

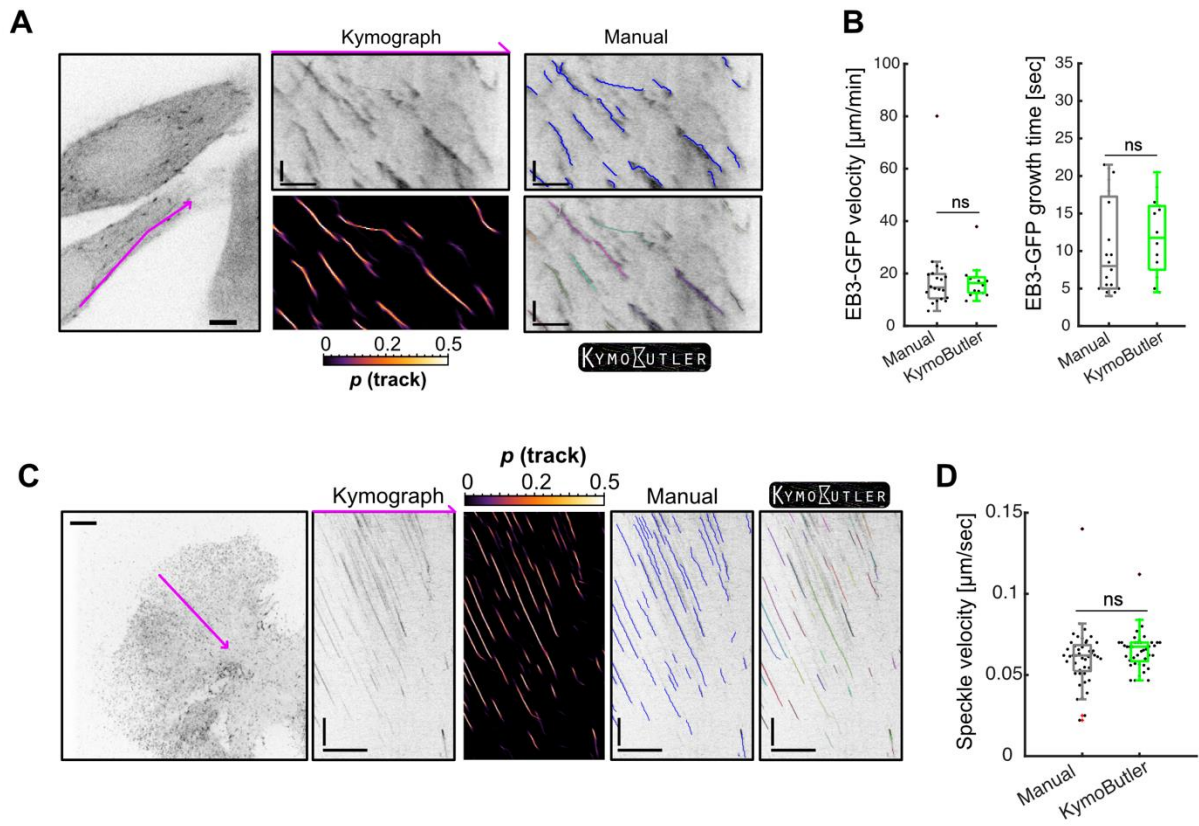
Figure 2



648

649

Figure 3



650

651

652

Figure 4