

# KymoButler, a Deep Learning software for automated kymograph analysis

3

4 Maximilian A. H. Jakobs\*, Andrea Dimitracopoulos, Kristian Franz\*

5

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

8

\* To whom correspondence should be addressed: [mj455@cam.ac.uk](mailto:mj455@cam.ac.uk) (M.A.H.J.) or  
[kf284@cam.ac.uk](mailto:kf284@cam.ac.uk) (K.F.)

11

## 12 **Abstract**

13 Kymographs are graphical representations of spatial position over time, which are often used  
14 in biology to visualise the motion of fluorescent particles, molecules, vesicles, or organelles  
15 moving along a predictable path. Although in kymographs tracks of individual particles are  
16 qualitatively easily distinguished, their automated quantitative analysis is much more  
17 challenging. Kymographs often exhibit low signal-to-noise-ratios (SNRs), and available tools  
18 that automate their analysis usually require manual supervision. Here we developed  
19 KymoButler, a Deep Learning-based software to automatically track dynamic processes in  
20 kymographs. We demonstrate that KymoButler performs as well as expert manual data  
21 analysis on kymographs with complex particle trajectories from a variety of different  
22 biological systems. The software was packaged in a web-based "one-click" application for  
23 use by the wider scientific community. Our approach significantly speeds up data analysis,  
24 avoids unconscious bias, and represents another step towards the widespread adaptation of  
25 Machine Learning techniques in biological data analysis.

26

## 27 **Introduction**

28 Many processes in living cells are highly dynamic, and molecules, vesicles, and organelles  
29 diffuse or are transported along complex trajectories. Particle tracking algorithms represent  
30 powerful approaches to track the dynamics of such particles ((Jaqaman et al. 2008;  
31 Sbalzarini & Koumoutsakos 2005; Lee & Park 2018)). However, particularly in scenarios  
32 where particles follow a stationary path and move much faster than the confounding cell  
33 (e.g., as in molecular transport along neuronal axons and dendrites, retrograde actin flow, or  
34 cilia transport), kymographs provide an elegant solution to the visualisation and analysis of  
35 particle dynamics. Kymographs are generated by stacking the intensity profile along a  
36 defined path for each time point of a movie. In the resulting space-time image, each (usually  
37 fluorescently) labelled particle is shown as a line, whose slope, for example, represents the  
38 velocity of that particle (Figure 1A).

39 In many biological processes, multiple particles move along the same stationary path with  
40 little to no deviations, making kymographs a very useful representation of their dynamics.  
41 Hence, kymographs have been widely employed to visualise biological processes across  
42 different length scales, ranging from diffusion and transport of single molecules to whole cell  
43 movements (Twelvetrees et al. 2016; Barry et al. 2015). The analysis of these kymographs  
44 only requires tracing lines in 2D images, a rather simple task compared to the more general  
45 approach of particle tracking, where one has to identify the centre of the particles in each  
46 frame, and then correctly assign these coordinates to corresponding particles across frames.

47 Publicly available kymograph analysis software simplifies the tedious and time-consuming  
48 task of tracing kymographs, but most of these solutions require manual supervision, and  
49 they are mainly applicable to particles that follow a unidirectional motion, i.e. do not change  
50 their direction or velocity (Figure 1C, example 2) (Neumann et al. 2017; Mangeol et al. 2016;  
51 Chenouard et al. 2010; Zala et al. 2013). This category includes, for example, the dynamics  
52 of growing microtubule +ends and F-actin dynamics in retrograde actin flow (Lazarus et al.

53 2013; del Castillo et al. 2015; Alexandrova et al. 2008; Babich et al. 2012). In many other  
54 biological contexts, however, particles can stop moving, change direction, merge, cross each  
55 other's path, or disappear for a few frames. The kymographs obtained from these processes  
56 exhibit 'bidirectional' motion (Figure 1C, example 1); this category includes cellular transport  
57 processes, for example molecular or vesicle transport in neuronal axons and dendrites (Fait  
58 et al. 2016; Tanenbaum et al. 2013; Koseki et al. 2017). Thus, the problem of automatically  
59 and reliably tracking dynamic processes in kymographs is still largely unresolved, and given  
60 the limitations of currently available kymograph analysis software, most kymographs are still  
61 analysed by hand, which is slow and gives rise to unconscious bias.

62 In recent years, Machine Learning (ML), and particularly Deep Neural Networks, have been  
63 very successfully introduced to data processing in biology and medicine (Mathis et al. 2018;  
64 Weigert et al. 2017; Florian et al. 2017; Guerrero-Pena et al. 2018; Falk et al. 2019; Bates et  
65 al. 2017). ML-based image analysis has several advantages over other approaches: it is less  
66 susceptible to bias than manual annotation, it takes a much shorter time to analyse large  
67 datasets, and, most importantly, it comes closer to human performance than conventional  
68 algorithms (Mathis et al. 2018).

69 Most ML approaches to image analysis utilise **Fully Convolutional Deep Neural Networks**  
70 (FCNs) that were shown to excel at object detection in images (Dai et al. 2016; Szegedy et  
71 al. 2014; LeCun et al. 2008; Falk et al. 2019). Through several rounds of optimisation, FCNs  
72 select the best possible operations by exploiting a multitude of hidden layers. These layers  
73 apply image convolutions using kernels of different shapes and sizes, aiming to best match  
74 the output of the neural network to the provided training data labels, which were previously  
75 derived from manual annotation. This means that the network learns to interpret the images  
76 based on the available data, and not on *a priori* considerations. This approach has become  
77 possible due to the incredible improvements in computation times of modern CPUs and the  
78 adoption of GPUs that can execute an enormous number of operations in parallel. Currently,  
79 the most successful architecture for biological and medical image analysis is the U-Net,

80 which takes an input image to generate a binary map that highlights objects of interest based  
81 on the training data (Ronneberger et al. 2015).

82 Here we present KymoButler, a new stand-alone FCN software based on the U-Net  
83 architecture, to automatically and reliably extract particle tracks from kymographs. The  
84 software was packaged into an easy-to-use web interface and a downloadable software  
85 package, and it was benchmarked against traditional software and manual annotation on  
86 synthetic (i.e., ground truth) data. We show that KymoButler performs very well on  
87 challenging bidirectional kymographs, where particles disappear, reappear, merge, cross  
88 each other's path, move in any direction, change speed, immobilise, and reverse direction.  
89 KymoButler thus represents a substantial improvement in the automation of kymograph  
90 tracing, speeding up the experimental workflow, while preserving the accuracy of manual  
91 annotations.

## 92 **Results**

### 93 The KymoButler software package

94 For our FCN-based kymograph analysis software, we implemented a customised  
95 architecture based on the U-Net (Ronneberger et al. 2015). We first trained the FCN to  
96 segment kymographs, i.e. binarize the image into regions with particle tracks (foreground)  
97 and noise (background). Our training data consisted of manually annotated tracks in 487  
98 unidirectional and 79 bidirectional kymographs (unpublished data from our group and other  
99 laboratories, see Materials and Methods and Acknowledgements for details). Since no  
100 ground truth was available in the manually annotated kymographs, we also generated 221  
101 synthetic unidirectional and 21 synthetic bidirectional kymographs that were used for training  
102 (see Figure 1-figure supplement 3 for examples).

103 Our network takes an input kymograph to generate 2D maps that assign a “trackness” value  
104 between 0 and 1 to each pixel of the input image, with higher values representing a higher  
105 likelihood of pixels being part of a track. The training was performed with pixel-wise cross-  
106 entropy loss (see Methods for details) and implemented in Mathematica  
107 (<http://www.wolfram.com/mathematica>). We furthermore took advantage of the intrinsic  
108 differences in the appearance of unidirectional and bidirectional kymographs and trained two  
109 separate specialised networks, a unidirectional segmentation module, and a bidirectional  
110 segmentation module (Figure 1–figure supplement 1 and Figure 1–figure supplement 2).  
  
111 The unidirectional segmentation module generates separate trackness maps for tracks with  
112 negative and positive slopes (which could, for example, correspond to tracks of anterograde  
113 and retrograde transport processes, respectively), to remove line crossings from the output  
114 (Figure 1–figure supplement 1). The trackness maps are then binarized and morphologically  
115 thinned to yield separated lines in a skeletonized map (Figure 1–figure supplement 1). We  
116 found the binarization threshold to depend on the biological application and on the signal to  
117 noise ratio of the input image. For our synthetic data, we used a value of 0.2 and generally  
118 observed consistent results for both segmentation modules between 0.1-0.3 (Figure 1–figure  
119 supplement 4).  
  
120 In bidirectional kymographs, tracks show more complex morphologies, since they can  
121 change direction and cross each other multiple times. The bidirectional segmentation  
122 module therefore generates a single trackness map, which needs to be further processed in  
123 order to obtain individual particle tracks. After thresholding and morphologically thinning the  
124 trackness map, we obtained a skeletonised image with multiple track crossings (Figure 1–  
125 figure supplement 1). In these images, we detected starting points of tracks by  
126 morphological operations (Figure 1–figure supplement 1B) and moved along each line from  
127 one row (time point) to the next. Then, whenever a crossing point was encountered (with two  
128 or more possible pixels to advance to), the software calls a decision module to resolve the  
129 crossing. The decision module, again based on a modified version of the U-Net, is

130 specialised in solving these crossings and trained on our bidirectional kymograph data  
131 (Figure 1–figure supplement 1B and Figure 1–figure supplement 2). The inputs of the  
132 module consist of three 48 by 48 pixel crops: (1) the input kymograph, (2) the skeletonised  
133 trackness map, and (3) the skeleton of the current track (Figure 1–figure supplement 1B).  
134 The output of the module is a map that assigns a score between 0 and 1 to each pixel of the  
135 skeletonised trackness map (2). Then, the most likely skeleton segment to continue the  
136 current track (3) is selected from the decision score map and the average score saved as a  
137 measure for track confidence. If the predicted path is less than 3 pixels long, the track is  
138 resolved and terminated. Once all the tracks with starting points are resolved, they are  
139 removed from the skeletonised trackness map, which is then scanned again for starting  
140 points, and the steps above are repeated until no further starting points are found.  
141 Furthermore, long overlaps between tracks are assigned to the track with the highest  
142 confidence so that no large overlapping regions between tracks are found in the final result  
143 (see Materials and Methods).  
144 Finally, we implemented the class module, a simple convolutional network that classifies  
145 input kymographs into unidirectional or bidirectional classes (Figure 1–figure supplement 1B  
146 and Figure 1–figure supplement 2A). The class module was trained on both unidirectional  
147 and bidirectional data until the error rate on a validation dataset, which contained 72  
148 kymographs and their classes, became persistently 0%. We linked the class module to the  
149 unidirectional and bidirectional segmentation modules as well as to the decision module  
150 (Figure 1–figure supplement 1B), and packaged them into KymoButler, an easy-to-use, drag  
151 & drop browser-based app for quick and fully automated analysis of individual kymographs  
152 (<http://kymobutler.deepmirror.ai>).  
153 The only free parameter in KymoButler is the threshold for trackness map segmentation. The  
154 default threshold is 0.2, but users can freely adjust it between 0.1 and 0.3 (+1 and -1 in the  
155 cloud interface) for their specific application. After the computation, which only takes 1-20  
156 seconds per kymograph (depending on complexity), KymoButler generates several files

157 including a dilated overlay image highlighting all the tracks found in different colours, a CSV  
158 file containing all track coordinates, and another summary file with post processing data,  
159 such as average velocities and directionality (Figure 1B). Finally, we tested KymoButler on  
160 previously published kymographs from a variety of different biological data (Figure 1C and  
161 Figure 1-figure supplement 1A) and on unpublished data from collaborators (not shown).

## 162 Performance on unidirectional Kymographs

163 We quantitatively evaluated the performance of KymoButler on unidirectional kymographs,  
164 i.e. particles that move with mostly uniform velocities and with no change in direction (Figure  
165 1C, Figure 2, Figure 1-figure supplement 1A). The unidirectional module of KymoButler was  
166 compared to an existing kymograph analysis software, which is based on Fourier filters, and  
167 which provided the best performance among publicly available software in our hands  
168 (KymographDirect package (Mangeol et al. 2016)). Additionally, we traced kymographs by  
169 hand to obtain a control for the software packages.

170 First, we generated 10 synthetic movies depicting unidirectional particle dynamics with low  
171 signal-to-noise ratio (~1.2, see Materials and Methods) and extracted kymographs from  
172 those movies using the KymographClear (Mangeol et al. 2016) Fiji plugin. Each of the  
173 kymographs was then analysed by Fourier-filtering (KymographDirect), KymoButler, and by  
174 hand, and the identified trajectories overlaid with the ground truth (i.e., the known dynamics  
175 of the simulated data) (Figure 2A).

176 We then quantified the quality of the predicted traces. We first determined the best predicted  
177 track for each ground truth track (in case several segments were predicted to cover the  
178 same track) and then calculated the fraction of the length of the ground truth track that was  
179 correctly identified by that predicted track (“track recall”) (Figure 2B). Additionally, we  
180 determined the best overlapping ground truth track for each predicted track and then  
181 calculated the fraction of the length of the predicted track that was overlapping with the

182 ground truth track (“track precision”). Examples of low/high precision and low/high recall are  
183 shown in Figure 2B. We then calculated the geometric mean of the average track recall and  
184 the average track precision (the “track F1 score”, see methods) for each kymograph (Figure  
185 2E). The median F1 score of the manual control was 0.90, KymoButler achieved 0.93, while  
186 Fourier filtering achieved a significantly lower F1 score of 0.63 ( $p = 4 \cdot 10^{-5}$ , Kruskal-Wallis  
187 Test, Tukey post-hoc: manual vs KymoButler  $p = 0.6$ , manual vs Fourier Filtering  $p = 3 \cdot$   
188  $10^{-3}$ ).

189 Our synthetic data also included gaps of exponentially distributed lengths (see Materials and  
190 Methods), allowing us to quantify the ability of KymoButler to bridge gaps in kymograph  
191 tracks (Figure 2C, F), which are frequently encountered in kymographs extracted from  
192 fluorescence data (Applegate et al. 2011). Both KymoButler and manual annotation  
193 consistently bridged gaps that belonged to the same trajectory, while Fourier filtering was  
194 less accurate (89% of all gaps correctly bridged by KymoButler, 88% by manual, and 72%  
195 by Fourier filter analysis; median of all 10 synthetic kymographs,  $p = 10^{-4}$ , Kruskal-Wallis  
196 Test, Tukey post-hoc: manual vs KymoButler  $p = 0.9$ , manual vs Fourier Filtering  $p = 2 \cdot$   
197  $10^{-3}$ , Figure 2F).

198 We also quantified the ability of KymoButler to resolve track crossings. Again, both  
199 KymoButler and manual annotation performed significantly better than Fourier filtering (88%  
200 KymoButler, 86% manual, 60% Fourier filter; median percentage of correctly resolved  
201 crossings of all 10 synthetic kymographs,  $p = 10^{-4}$ , Kruskal-Wallis Test, Tukey post-hoc:  
202 manual vs KymoButler  $p = 0.9$ , manual vs Fourier Filtering  $p = 1 \cdot 10^{-3}$ , Figure 2G). In  
203 summary, KymoButler was able to reliably track particle traces in kymographs at low SNR,  
204 and it clearly outperformed currently existing software, while being as consistent as manual  
205 expert analysis.

206 KymoButler performance on bidirectional Kymographs

207 As in many kymographs obtained from biological samples trajectories are not unidirectional,  
208 we also tested the performance of KymoButler on complex bidirectional kymographs, i.e. of  
209 particles with wildly different sizes, velocities, and fluorescence intensities that frequently  
210 change direction, may become stationary and then resume motion again (see Figure 1B, C,  
211 Figure 3A, Figure 1-figure supplement 1A for examples). Available fully automated software  
212 that relied on edge detection performed very poorly on our synthetic kymographs (Figure 3-  
213 figure supplement 1). Therefore, we implemented a custom-written wavelet coefficient  
214 filtering algorithm in order to compare our FCN-based approach to a more traditional non-ML  
215 approach (Figure 3A, Figure 3-figure supplement 1, Materials and Methods). In short, the  
216 wavelet filtering algorithm generates a trackness map, similar to KymoButler, by applying a  
217 stationary wavelet transform to the kymograph to generate so-called “coefficient images” that  
218 highlight horizontal or vertical lines. These coefficient images are then overlaid and binarized  
219 with a fixed value (0.3), skeletonised, and fed into the KymoButler algorithm without the  
220 decision module, i.e. crossings are resolved by linear regression prediction.

221 We generated 10 kymographs from our synthetic movies with the KymographClear package  
222 (average signal-to-noise ratio was 1.4, since any lower signal generally obscured very faint  
223 and fast tracks). Each of the kymographs was then analysed by wavelet coefficient filtering,  
224 KymoButler, and manual annotation, and the predicted traces overlaid with the ground truth  
225 (Figure 3A). While the wavelet approach and KymoButler were able to analyse the 10  
226 kymographs in less than one minute, manual annotation by an expert took about 1.5 hours.  
227 Moreover, whereas the manual annotation and KymoButler segmentation overlaid well with  
228 the ground truth, the wavelet approach yielded numerous small but important deviations.

229 Similarly to the unidirectional case, we quantified track precision and recall (Figure 3B, E)  
230 and calculated the resolved gap fraction (Figure 3C, F) and crossing fraction (Figure 3D, G).  
231 The median of the track F1 scores per kymograph for manual annotation (0.82) was similar

232 to KymoButler (0.80), while the wavelet filter approach only gave 0.60 ( $p = 8 \cdot 10^{-5}$ ,  
233 Kruskal-Wallis Test, Tukey post-hoc: manual vs KymoButler  $p = 0.7$ , manual vs wavelet  
234 filtering  $p = 10^{-4}$ , Figure 3E). While gaps were resolved by KymoButler and manual  
235 annotation in 89% and 95% of cases, respectively, only 74% were resolved by the wavelet  
236 algorithm (median of all 10 synthetic kymographs,  $p = 3 \cdot 10^{-4}$ , Kruskal-Wallis Test, Tukey  
237 post-hoc: manual vs KymoButler  $p = 0.4$ , manual vs wavelet filtering  $p = 2 \cdot 10^{-4}$ , Figure 3F).  
238 Crossings were rarely resolved correctly by the wavelet algorithm (12%) but much more  
239 reliably by KymoButler (61%) and manual annotation (76%) (median of all 10 synthetic  
240 kymographs,  $p = 3 \cdot 10^{-5}$ , Kruskal-Wallis Test, Tukey post-hoc: manual vs KymoButler  
241  $p = 0.4$ , manual vs wavelet filtering  $p = 2 \cdot 10^{-5}$ , Figure 3G).  
242 Overall, these results showed that KymoButler performs well on both unidirectional and  
243 bidirectional kymographs, outperforms currently available automated analysis of kymographs,  
244 and it performs as well as manual tracing, while being much faster and not prone to  
245 unconscious bias.

## 246 **Discussion**

247 In this work, we developed software based on Deep Learning techniques to automate the  
248 tracking of dynamic particles along a stationary path in a noisy cellular environment.  
249 Convolutional neural networks (CNNs) are nowadays widely applied for image recognition.  
250 Since tracking is *a priori* a visual problem, we built a modular software utilising CNNs for  
251 identifying tracks in kymographs. We deployed our networks as KymoButler, a software  
252 package that takes kymographs as inputs and outputs all tracks found in the image in a  
253 matter of seconds. The network outperforms standard image filtering techniques on synthetic  
254 data as well as on kymographs from a wide range of biological processes, while being as  
255 precise as expert manual annotation.

256 The KymoButler software has only one adjustable parameter that is left to the user: a  
257 sensitivity threshold that, if low, allows more ambiguous tracks to be recognised, and if high  
258 discards them. For our synthetic data, the best value for the threshold lay between 0.1 and  
259 0.3 (Figure 1-figure supplement 4), and we observed a similar range for a variety of  
260 kymographs from published data. However, the threshold depends on the SNR of the input  
261 images, so that the correct threshold has to be chosen based on each biological application  
262 and imaging conditions. We strongly recommend to visually inspect the output of  
263 KymoButler for each new application, and to compare the output to manual annotation.

264 Most of the publicly available kymograph analysis software requires manual labelling to  
265 extract quantitative data (Chenouard et al. 2010; Neumann et al. 2017; Zala et al. 2013).  
266 Some automated approaches have been published in the context of specific biological  
267 questions, but since these programs are currently not publicly available it is not clear how  
268 well they would perform on kymographs from other applications (Mukherjee et al. 2011; Reis  
269 et al. 2012). Other approaches do not extract individual tracks but only macroscopic  
270 quantities, as for example velocities (Chan & Odde 2008). As KymoButler is fully automated  
271 and able to reliably analyze kymographs from a wide range of biological applications, it fills  
272 an important gap. Here we showed that KymoButler is able to quantify mitochondria  
273 movement in neuronal dendrites, microtubule growth dynamics in axons, and *in vitro*  
274 dynamics of single cytoplasmic dynein proteins (Figure 1 and Figure 1-figure supplement 1).  
275 We predict that it can furthermore be applied to most if not all other kymographs obtained  
276 from time-lapse fluorescence microscopy without the need of any modifications.

277 KymoButler outperformed Fourier filtering, edge detection, and customised wavelet  
278 coefficient selection on synthetic kymographs. While Fourier filtering 'only' performed ~30%  
279 worse than KymoButler on unidirectional kymographs, edge detection on bidirectional  
280 kymographs suffered greatly from background fluctuations and low SNR to such an extent  
281 that the extracted data was unusable (see Figure 3-figure supplement 1 for one example).  
282 Therefore, we designed a filtering algorithm based on wavelet coefficient image selection to

283 analyse complex bidirectional kymographs specifically for our synthetic data. KymoButler still  
284 performed 25% better than this approach (Figure 3). The main problem with either filtering  
285 approach compared to KymoButler was their inability to bridge track gaps and resolve line  
286 crossings, both of which occur frequently in biological data (Figure 2C, D and 3C, D). These  
287 challenges are met by KymoButler, which performed as well as expert annotation, but within  
288 a much shorter time (Figure 2 and 3).

289 Our results show that KymoButler is able to correctly identify individual full-length tracks in  
290 kymographs with an average track F1 score (geometric mean of track precision and recall)  
291 of 92% on unidirectional tracks and 80% on complex bidirectional tracks, without suffering  
292 from inconsistency, bias, and laborious tracing, that plague manual tracking. While  
293 KymoButler is already performing very well, we aim to significantly improve it over future  
294 iterations. Every time a researcher uses our webform, the corresponding kymograph is  
295 anonymously uploaded to our cloud. Once a large number of diverse kymographs is  
296 uploaded, these kymographs will be added to our training data, improving KymoButler even  
297 further.

298 The ultimate challenge will be to expand our approach to 2D or even 3D tracking problems.  
299 Here, we defined a 1D region of interest in 2D time-lapse movies, extracted 2D (space and  
300 time) images (kymographs), and finally tracked 2D lines in those images. A similar, albeit  
301 computationally heavier, approach could stack the frames of a 2D/3D movie on top of each  
302 other to generate a 3D/4D kymogram (2D space and time, or 3D space and time). Previously  
303 generated kymograms have led to intriguing results on whole-cell particle tracking problems  
304 with high SNR (Racine et al. 2007). The use of higher dimensional FCNs in the future has  
305 great potential to yield human-like performance on any biological and medical tracking  
306 problems.

307

## 308 **Material and Methods**

309 All code was written in the wolfram language in Mathematica  
310 <https://wolfram.com/mathematica> and, if not stated otherwise, can be found online under our  
311 GitHub: <https://github.com/deepmirror/KymoButler>

## 312 The KymoButler software package

313 The KymoButler software was implemented in Mathematica to take advantage of easy web  
314 form deployment and distribution. The workflow is shown in Figure 1-figure supplement 1B.  
315 Our approach was to first segment kymograph pixels that are part of particle tracks from  
316 pixels that were part of the background with our segmentation modules. From previous work  
317 we knew that kymographs that depict unidirectional movement only, can be filtered into  
318 tracks that have positive slope and those that have negative slope (Chenouard et al. 2010),  
319 while no such assumptions can be made about bidirectional kymographs. Hence, we  
320 decided to take advantage of this simplification of unidirectional kymograph analysis by  
321 training two modules: one that is specialized to segment unidirectional kymographs and  
322 another one that segments bidirectional ones. Note that the bidirectional module is able to  
323 analyze any kymograph, including unidirectional ones, but since it is not specialized it  
324 performs slightly worse than the unidirectional module on unidirectional kymographs. To  
325 further simplify software usability, we prepended a class module that classifies input  
326 kymographs as bidirectional or unidirectional, and then applies the corresponding  
327 segmentation module and decision module (for bidirectional kymographs only). Our  
328 downloadable software package on GitHub allows the user to call either segmentation  
329 module (unidirectional/bidirectional) directly, if they wish to do so.  
  
330 When the kymograph is classified as unidirectional by the class module, the unidirectional  
331 segmentation module generates two trackness score maps for particles with negative or  
332 positive slope (Figure 1-figure supplement 1B). Since the particles move with roughly the

333 same velocity, the resulting maps mostly do not exhibit any crossings. Thus, we binarize the  
334 maps with a threshold between 0.1-0.3 (see benchmarking section for more information  
335 about the threshold). The resulting binary maps are then thinned iteratively so that each  
336 trace is only one pixel wide at any point and pruned so that branches that are shorter than 3  
337 pixels are deleted. Subsequently, each trace is segmented and selected only if they are at  
338 least 3 frames long. In the final step, pixels that lie in the same row of the kymograph are  
339 averaged over so that the final track has only one entry per frame.

340 For bidirectional kymographs the software generates a trackness map, applies a binarization  
341 threshold (0.1-0.3, see benchmarking for more details), iterative thinning, and pruning  
342 (minimum length 3 pixels). However, since the resulting skeletonised map had a substantial  
343 number of crossings, and could not be easily segmented to yield individual tracks, we  
344 implemented a further module in the software. First, all lines in the skeletonised map are  
345 shortened so that each white pixel at a track end only has neighbouring pixels in different  
346 rows (time dimension). This was done so that we could detect track starting points ("seeds")

347 with a Hit-Miss transformation with kernel:  $\begin{pmatrix} -1 & -1 & -1 \\ -1 & 1 & -1 \\ 0 & 0 & 0 \end{pmatrix}$ . Application of this kernel yielded  
348 a binary map with 0 everywhere except at track seeds (Figure 1-figure supplement 1B, red  
349 dots). These seeds were then used to start tracing individual tracks in the kymograph by  
350 always advancing to the next white pixel. Once more than one potential future pixel is  
351 encountered, the decision module is called. The module generates three 48x48 crops of (1)  
352 the input kymograph, (2) the skeletonised trackness map, and (3) the skeleton of the current  
353 track and predicts a trackness map that has high values on the skeleton segment of the  
354 most likely future track (Figure 1-figure supplement 1B). This map is binarized with threshold  
355 0.5 and thinned. The precise threshold had little effect on the final output, so we fixed it at  
356 0.5 for all applications. Next, the largest connected component in the map is selected as the  
357 most likely future path and appended to the track if longer than 2 pixels. The average  
358 trackness value of this component (from the decision module prediction) is saved as a

359 measure of decision “confidence”. This process is repeated until no further possible pixels  
360 are found or no future path is predicted which is when the track is terminated. Once all seeds  
361 are terminated, the software subtracts all the found paths from the skeletonised trackness  
362 map and again looks for new seeds which are then again tracked in the full skeletonised  
363 image. The process is repeated until no further seeds are found, and then all tracks are  
364 averaged over their timepoints (rows in the kymograph image). Subsequently the software  
365 deletes tracks that are shorter than 5 pixels or part of another track and assigns overlaps  
366 that are longer than 10 pixels to the track with the highest average decision confidence.

367 Both the unidirectional and the bidirectional module output a coloured overlay in which each  
368 track is drawn in a different randomly assigned colour and dilated with factor 1 for better  
369 visibility (see Figure 1B-C and Figure 1-figure supplement 1A). Additionally, the software  
370 generates one CSV file that contains all the track coordinates and a summary CSV file that  
371 gives derived quantities, such as track direction and average speed.

372 The software was deployed from Mathematica as a cloud based interface  
373 (<http://kymobutler.deepmirror.ai>) and a Mathematica package  
374 (<https://github.com/deepmirror/KymoButler>)

## 375 Network architectures

376 Our networks were built from convBlocks (a convolutional layer with 3x3 kernel size, padding,  
377 and arbitrary number of output channels followed by a batch normalisation layer and a ‘leaky’  
378 ramp (leakyReLU) activation function ( $\text{leakyReLU}(x) := \max(x, 0) - 0.1 \max(-x, 0)$ ). Batch  
379 normalisation is useful to stabilise the training procedure as it rescales the inputs of the  
380 activation function (leakyReLU), so that they have zero mean and unit variance. The  
381 leakyReLU prevents the so-called “dead ReLu’s” by applying a small gradient to values  
382 below 0. These building blocks were previously used for image recognition tasks in Google’s  
383 inception architecture and in the U-Net architecture (Szegedy et al. 2014; Falk et al. 2019).

384 The module architectures we settled on are shown in Figure 1–figure supplement 1-2. All  
385 modules used the same core building blocks while having different input and output ports.  
386 The classification module takes a resized kymograph of size 64x64 pixels and generates two  
387 output values that correspond to the class probabilities for unidirectional/bidirectional  
388 kymographs (Figure 1–figure supplement 2A). The unidirectional segmentation module takes  
389 one input kymograph and generates two output images that correspond to the trackness  
390 scores of particles with positive or negative slopes (Figure 1–figure supplement 2B). The  
391 bidirectional segmentation module takes one input kymograph and generates one trackness  
392 score map highlighting any found particle tracks (Figure 1–figure supplement 2C). Finally,  
393 the decision module takes three inputs of size 48x48 pixels to generate one trackness map  
394 (Figure 1–figure supplement 2D). All modules share the same core network that is  
395 essentially a U-Net with padded convolutions and with 64 (in the top level) to 1024 (in the  
396 lowest level) feature maps. We experimented with more complex architectures (parallel  
397 convolution modules instead of blocks, different number of feature maps) but could only  
398 observe minor increase in accuracy at a large expense in computation time. Due to the U-  
399 Net architecture, each dimension of the inputs to the segmentation modules needs to be a  
400 multiple of 16. Thus, inputs were resized when they did not match the dimension  
401 requirements, and then the binarized output images from the segmentation modules were  
402 resized to the original input image size before proceeding further.

#### 403 Network training

404 To train the networks we quantified the difference between their output  $o$  and the desired  
405 target output  $t$  through a cross entropy loss layer ( $CEloss(t, o) = -(t \cdot \ln(o) + (1 - t) \cdot$   
406  $\ln(1 - o))$ . The loss was averaged over all output entries (pixels and classes) of each  
407 network. While we tried other loss functions, specifically weighted cross entropy loss and  
408 neighbour dependent loss as described in (Bates et al. 2017), we persistently obtained  
409 higher precision and recall with the basic cross entropy loss above.

410 Our training data comprised a mixture of synthetic data and manually annotated unpublished  
411 kymographs, kindly provided by the research groups mentioned in the acknowledgements.

412 Most of the manual annotation was done by M. A. H. J. and A. D. In total, we used 487  
413 (+200 synthetic) unidirectional, and 79 (+21 synthetic) bidirectional kymographs, with 95% of  
414 the data used for network training, and ~5% of retained for network validation. All network  
415 training was performed on a workstation, using a nVidia 1080 Ti or a nVidia 1070 GPU.

416 The class module depicted in Figure 1-figure supplement 2A was trained with batches of  
417 size 50 (with 25 unidirectional and 25 bidirectional kymographs to counter class imbalance)  
418 with random image transformations that included image reflections, rotations, resizing,  
419 colour negation, gaussian noise, random noise, and random background gradients. The final  
420 input image was randomly cropped to 64x64 pixels (see examples Figure 1-figure  
421 supplement 3A) and the class module was trained using stochastic gradient descent (ADAM  
422 optimiser (Kingma & Optimization n.d.), initial learning rate 0.001), until the validation set  
423 error rate was consistently 0%.

424 The unidirectional segmentation module (Figure 1-figure supplement 2B) was trained with  
425 batches comprising 20 randomly selected kymographs from our training set (example in  
426 Figure 1-figure supplement 3B). We applied the following image transformations: Random  
427 reflections along either axis, random 180-degree rotations, random cropping to 128x80  
428 pixels (approximately the size of our smallest kymograph), random gaussian and uniform  
429 noise, and random background gradients. Note that we did not apply any resizing to the raw  
430 kymograph since that generally decreased net performance. Additionally, we added Dropout  
431 Layers (10-20%) along the contracting path of our custom U-Net to improve regularisation.

432 Each kymograph in this training set was generated by hand with KymographTracker  
433 (Chenouard et al. 2010), but to increase dataset variability we took the line profiles from  
434 KymographTracker and generated kymographs with a custom Mathematica script that  
435 applied wavelet filtering to the plotted profiles. The resulting kymographs have a slightly  
436 different appearance than the ones created with KymographTracker and are thus useful to

437 regularize our training process. Several modules were trained until convergence and the  
438 best performing one (according to the validation score) was selected (ADAM optimiser, initial  
439 learning rate of 0.001, learning rate schedule = *If*[batch < 4000, 1, .5]).

440 The bidirectional segmentation module (Figure 1-figure supplement 2C, example data Figure  
441 1-figure supplement 3C) was trained in the same way as the unidirectional segmentation  
442 module, with the exception of a slightly different learning rate schedule (*If*[batch < 3000, 1,  
443 .5]). Additionally, since we did not have access to many of the original movies from which the  
444 kymographs were generated, we could not generate kymographs with different algorithms as  
445 done for the unidirectional module.

446 Training data for the decision module (Figure 1-figure supplement 2D) was obtained from the  
447 bidirectional (synthetic + real) kymographs by first finding all the branch points in a given  
448 ground truth or manually annotated image. Then, each track was separated into multiple  
449 segments, that go from its start point to a branching point or its end point. For each  
450 branchpoint encountered while following a track, all segments that ended within 3 pixels of  
451 the branchpoint were selected. Then, (1) a 48x48 pixel crop of the raw kymograph around  
452 the branchpoint, (2) a binary map representing the track segment upstream of the branching  
453 point (centred with its end in pixel coordinates 25,25, with image padding applied if the end  
454 was close to an image corner), and (3) the corresponding 48x48 pixel region in the binary  
455 image representing all possible paths were used as inputs to the decision module. The  
456 binary image representing the ground truth or annotated future segment downstream of the  
457 branchpoint was used as the target image (see Figure 1-figure supplement 3D for an  
458 example training set). Thus, the training set comprised three input images and one output  
459 image which we used to train the decision module. To increase the module's focus on the  
460 non-binary raw kymograph crop, we applied 50% dropout to the full skeletonised input and  
461 5% dropout to the input segment. As explained above, we used random image augmentation  
462 steps like reflections, rotations, gaussian + uniform noise. Additionally, we employed random  
463 morphological thinning to the binary input/output images to simulate artefacts. Several

464 networks were trained until convergence (pixel wise cross entropy loss, ADAM optimiser,  
465 initial learning rate 0.001, batch size 50, learning rate schedule *If* [ $batch < 8000$ , 1, .5]), and  
466 the best performing one was selected.

467 **Synthetic Data**

468 Synthetic data was generated by simulating individual particles on a stationary path of length  
469 300 pixels for 300 frames to generate 300x300 pixel kymographs. To obtain unidirectional  
470 particles we seeded 30+30 particles with negative or positive slope at random  
471 timepoints/positions. Next, a random velocity between 1-3 pixels/frame was chosen for all  
472 particles in the movie, with a random noise factor to allow slight changes in velocity, and a  
473 particle PSF between 3-6 pixels. Each particle was assigned a survival time drawn from an  
474 exponential distribution with scale 0.01, after which it would disappear. Gaps of random  
475 length (exponentially distributed) were subsequently assigned to each track individually.  
476 From these tracks we then generated a kymograph with gaussian noise, used for neural  
477 network training, and a 20x300 pixel movie with 300 frames for benchmarking. The resulting  
478 kymographs and movies had an average signal-to-noise ratio of 1.2 (calculated as the  
479 average intensity of the signal, divided by the average intensity of the background). Finally,  
480 we removed tracks that overlapped for the whole duration of their lifetime.

481 To obtain synthetic data of complex bidirectional particle movements, we generated datasets  
482 with either 15 tracks (for benchmarking) or 30 tracks (for training) per movie. The maximum  
483 velocity was set to 3 pixels/frame, as above this velocity it became hard to manually  
484 segment tracks from kymographs. Each movie was assigned a random velocity noise factor  
485 between 0 and 1.5 pixels/frame, a random switching probability between 0 and 0.1 (to switch  
486 between stationary and directed movement) and a random velocity flipping factor between 0  
487 and 0.1 (to flip the direction of the velocity). Individual particles were simulated by first  
488 calculating their lifetime from an exponential distribution with scale 0.001. Then, a random  
489 initial state, moving or stationary, was selected as well as a random initial velocity and a

490 particle size between 1-6 pixel. In the simulation, particles could randomly switch between  
491 different modes of movement (stationary/directed), flip velocities and were constantly  
492 subjected random velocity noise (movie specific). Finally, tracks that were occulted by other  
493 tracks were removed, and a movie (used for benchmarking) and a kymograph (used for  
494 training) were generated. The resulting kymographs and movies had an average signal-to-  
495 noise ratio of 1.4.

496 **Benchmarking**

497 In order to benchmark the performance of software and manual predictions, we implemented  
498 a custom track F1 score which was calculated as the geometric mean of track recall and  
499 track precision. To calculate track recall, each ground truth track was first compared to its  
500 corresponding predicted track, and the fractional overlap between them was calculated.  
501 Since predicted tracks do not necessarily follow the exact same route through a kymograph,  
502 but frequently show small deviations from the ground truth (see Figure 3 and Figure 3-figure  
503 supplement 1) we allowed for a 3.2-pixel deviation from the ground truth (2 diagonal pixels).  
504 The maximum fractional overlap was then selected and stored as the track recall. The recall  
505 was thus 1 when the full length of a ground truth track was predicted, and 0 if the track was  
506 not found in the prediction. We would like to highlight that this criterion is very strict: if a  
507 ground truth track is predicted to be 2 tracks (for example, by failing to bridge a gap along  
508 the track), the recall fraction would decrease by up to 50%, even if most of the pixels are  
509 segmented correctly and belong to predicted tracks.

510 Track precision was calculated by finding the largest ground truth track that corresponded,  
511 i.e. had the largest overlap, to each prediction, and then calculating the fraction of the  
512 predicted track that overlapped to the ground truth track. Therefore, a track precision of 1  
513 corresponded to a predicted track that was fully part of a ground truth track while a precision  
514 of 0 meant that the predicted track was not found in the ground truth. In general, increasing

515 precision leads to a lower recall and vice versa, so that taking the track F1 score as the  
516 geometric mean between the two is a good measure of overall prediction performance.

517 To quantify gap performance, we searched for track segments within 3 pixels of the gap for  
518 each frame, to allow for predictions that deviated slightly from the ground truth. Once each  
519 frame of the gap was assigned to a corresponding predicted segment, the gap was deemed  
520 resolved. If one or more frames of the gap had no overlapping segment to the prediction, the  
521 gap was labelled unresolved. Our synthetic tracks had 954 gaps in the 10 kymographs of  
522 unidirectional data, and 840 gaps in the 10 kymographs of bidirectional data, and the largest  
523 gap size was 6 pixels. For each kymograph, we then calculated the fraction of gaps resolved.

524 To quantify KymoButler performance on crossings, we first generated binary images for  
525 each ground truth track and calculated overlaps with other ground truth tracks by multiplying  
526 those images with each other. The resulting images had white dots wherever two tracks  
527 crossed. Those dots were then dilated by a factor of 16 to generate circles and overlaid with  
528 the original single-track binary image to generate binary maps that contain segments of  
529 ground truth tracks that cross/merge with other tracks. Next, we generated dilated (factor 1)  
530 binary maps for each predicted track and multiplied them with each of those cross segments  
531 to obtain the largest overlapping track for each segment. We then visually inspected a few  
532 examples and determined that an overlap of 70% corresponds to a correctly resolved  
533 crossing and allowed for slight variations in predicted tracks when compared to ground truth.  
534 Finally, we calculated the fraction of crossings resolved per kymograph.

535 All statistical analysis was carried out in MATLAB (<http://mathworks.com>).

536 Module performance evaluation

537 To benchmark the unidirectional segmentation module of KymoButler, we generated 10  
538 synthetic movies of the dynamics of particles that move with uniform speed and do not  
539 change direction as described in the section about synthetic data generation. We then

540 imported these movies into ImageJ (<http://imagej.nih.gov>) via the Kymograph Clear package  
541 (Mangeol et al. 2016), drew a profile by hand and generated kymographs from them. These  
542 kymographs were then imported into the KymographDirect software package (also (Mangeol  
543 et al. 2016)), Fourier filtered and thresholded to extract individual particle tracks. This  
544 approach required manual selection of the threshold for each individual kymograph. We  
545 additionally traced the same kymographs by hand in ImageJ to compare software  
546 performance to expert analysis. To find a suitable range of binarization thresholds for our  
547 unidirectional segmentation module we calculated the track wise F1 score on the 10  
548 kymographs for thresholds between 0.05 and 0.5 (Figure 1-figure supplement 4). We  
549 observed the highest scores between 0.1 and 0.3 for both our synthetic data and other  
550 unpublished kymographs and also deemed these thresholds best by visual inspection of  
551 predicted kymograph tracks. Hence, we chose 0.2 as the segmentation map threshold to  
552 benchmark our predictions at.

553 In order to benchmark the bidirectional segmentation module and the decision module we  
554 generated 10 synthetic movies of the dynamics of complex bidirectional particles. These  
555 movies were imported into ImageJ with the KymographClear package and kymographs  
556 extracted. We subsequently tried to use the edge detection option in KymographDirect to  
557 extract individual tracks but were unable to obtain meaningful tracks (Figure 3-figure  
558 supplement 1). We also tried other options in the package but could not get good results on  
559 our synthetic data without substantial manual labor for each kymograph, defeating the goal  
560 of a fully automated analysis. Therefore, we wrote a custom script to carry out automated  
561 bidirectional kymograph analysis. We experimented with a few different approaches (for  
562 example fourier-filtering and customized edge detection) and settled on wavelet coefficient  
563 filtering as it gave the highest F1 score on our test dataset. This algorithm applied a  
564 stationary wavelet transformation with Haar Wavelets (Mathematica wavelet package) to  
565 each kymograph to decompose the image into different coefficient images that highlight  
566 different details (for example vertical or horizontal lines). We then selected only those

567 coefficient images that recapitulated particle traces in our synthetic kymographs. These  
568 images are overlaid and thresholded with an optimized threshold to generate binary maps  
569 that can be iteratively thinned to obtain a skeletonized “trackness” map similar to the outputs  
570 of our segmentation modules. This map was then traced with the same algorithm as in our  
571 decision module. However, while the KymoButler decision module used a neural network to  
572 predict path crossings, the wavelet filtering algorithm performed simple linear prediction by  
573 taking the dilated (factor 1) binary segment of a track and rotating it by 180 degrees. Then  
574 the “prediction” was multiplied with the skeletonized trackness map and the largest  
575 connected component selected as the future path. In contrast to the original decision module,  
576 this approach does not yield any information about decision “confidence”. Thus, to resolve  
577 track overlaps at the end of the algorithm, we randomly assigned each overlap to one track  
578 and deleted them from the others. Note that the wavelet approach was heavily optimized on  
579 our synthetic kymographs and performed poorly on generic real kymographs. We also traced  
580 the same 10 kymographs by hand in ImageJ. To find a suitable range of binarization  
581 thresholds for our bidirectional segmentation module we calculated the track wise F1 score  
582 for thresholds between 0.05 and 0.5 (Figure 1-figure supplement 4) and observed the same  
583 optimal range as the unidirectional segmentation module (0.1-0.3) for both our synthetic data  
584 and other unpublished kymographs. Hence, we chose 0.2 as the threshold score to  
585 benchmark our predictions.

586

587

588 **Key resources table**

Resource	Designation.	Source.	Identifiers.	Additional Information.
Software,	MATLAB	MATLAB	RRID: <a href="#">SCR_0</a>	Used for statistical

algorithm			<a href="#">01622</a>	analysis
Software, algorithm	Fiji	Fiji is Just ImageJ ( <a href="https://fiji.sc">https://fiji.sc</a> )	RRID: <a href="#">SCR_002285</a>	Used to generate and analyse kymographs with KymographClear/Direct <a href="https://sites.google.com/site/kymographanalysis/">https://sites.google.com/site/kymographanalysis/</a>
Software, algorithm	Wolfram Mathematica	Wolfram Mathematica	RRID:SCR_014448	Code available under <a href="https://github.com/deepmrror/KymoButler">https://github.com/deepmrror/KymoButler</a>

589

## 590 **Acknowledgements**

591 We would like to thank Eva Pillai for scientific input, proofreading, and logo design, Hannes  
592 Harbrecht for fruitful discussions about neural networks; Hendrik Schuermann, Ishaan  
593 Kapoor, and Kishen Chahwala for help with kymograph tracing; and the Mathematica stack  
594 exchange community (<https://mathematica.stackexchange.com>) without whom this project  
595 would have taken several decades longer. Unpublished kymographs to train KymoButler  
596 were provided by Caroline Bonnet (neurocampus, University of Bordeaux), Dr. Jean-Michel  
597 Cioni (San Raffaele Hospital, Milan), Dr. Julie Qiaojin Lin (University of Cambridge), Prof.  
598 Leah Gheber and Dr. Himanshu Pandey (Ben-Gurion University of the Negev), Dr. Carsten  
599 Janke and Satish Bodakuntla (Institut Curie, Paris), and Dr. Timothy O'Leary and Adriano  
600 Bellotti (University of Cambridge). Additionally, we would like to thank the Bordeaux Imaging  
601 Center, part of the national infrastructure France Biolimaging (ANR-10INBS-04-0), for  
602 valuable feedback on our software. We would also like to thank eLife and PLOS whose open  
603 access policy enabled us to show a variety of kymographs in this manuscript. The authors  
604 acknowledge funding by the Wellcome Trust (Research Grant 109145/Z/15/Z to M.A.H.J.),  
605 the Herchel Smith Foundation (Fellowship to A.D.), Isaac Newton Trust (Research Grant  
606 17.24(p) to K.F.), UK BBSRC (Research Project Grant BB/N006402/1 to K.F.), and the ERC  
607 (Consolidator Award 772426 to K.F.).

## 608 **Competing Interests**

609 We launched deepmirror.ai as a platform to promote the use of AI-based technologies for  
610 biological data analysis. We will be publishing tutorials and sample code to help people get  
611 started with developing their own machine learning software. We also intend to publish our  
612 work on KymoButler and future publications of our AI-based software on the website. All of  
613 this will be free of charge and available to all. Further in the future, we plan to also start  
614 offering paid professional services for customers that want to set up custom AI-based

615 software for applications, in case they are not covered by our research. This software may or  
616 may not be made available on [deepmirror.ai](http://deepmirror.ai), depending on our clients' requests.

617 **Software**

618 Quick and easy cloud platform: <http://www.kymobutler.deepmirror.ai>

619 Mathematica notebook with examples on how to use the software offline:

620 <https://github.com/deepmirror/KymoButler>

621

622

623 **References**

624 625 626 Alexandrova, A.Y. et al., 2008. Comparative Dynamics of Retrograde Actin Flow and Focal Adhesions: Formation of Nascent Adhesions Triggers Transition from Fast to Slow Flow C.-P. Heisenberg, ed. *PLoS one*, 3(9).

627 628 629 Applegate, K.T. et al., 2011. plusTipTracker: Quantitative image analysis software for the measurement of microtubule dynamics. *Journal of Structural Biology*, 176(2), pp.168–184.

630 631 Babich, A. et al., 2012. F-actin polymerization and retrograde flow drive sustained PLC $\gamma$ 1 signaling during T cell activation. *The Journal of Cell Biology*, 197(6), pp.775–787.

632 633 Barry, D.J. et al., 2015. Open source software for quantification of cell migration, protrusions, and fluorescence intensities. *The Journal of Cell Biology*, 209(1), pp.163–180.

634 635 Bates, R. et al., 2017. Extracting 3D Vascular Structures from Microscopy Images using Convolutional Recurrent Networks. *arXiv.org*, cs.CV.

636 637 Chan, C.E. & Odde, D.J., 2008. Traction Dynamics of Filopodia on Compliant Substrates. *Science*, 322(5908), pp.1687–1691.

638 639 640 Chenouard, N. et al., 2010. Curvelet analysis of kymograph for tracking bi-directional particles in fluorescence microscopy images. In 2010 17th IEEE International Conference on Image Processing (ICIP 2010). IEEE, pp. 3657–3660.

641 642 Dai, J. et al., 2016. R-FCN: Object Detection via Region-based Fully Convolutional Networks. pp.379–387.

643 644 645 del Castillo, U. et al., 2015. Interplay between kinesin-1 and cortical dynein during axonal outgrowth and microtubule organization in Drosophila neurons V. Allan, ed. *eLife*, 4, p.e10140.

646 647 Faits, M.C. et al., 2016. Dendritic mitochondria reach stable positions during circuit development. *eLife*, 5, p.e11583.

648 649 Falk, T. et al., 2019. U-Net: deep learning for cell counting, detection, and morphometry. *Nature methods*, 16(1), pp.67–70.

650 651 Florian, F. et al., 2017. Gp-Unet: Lesion detection from weak labels with a 3D regression network.

652 653 Guerrero-Pena, F.A. et al., 2018. Multiclass Weighted Loss for Instance Segmentation of Cluttered Cells. *arXiv.org*, cs.CV, pp.2451–2455.

654 655 Jaqaman, K. et al., 2008. Robust single-particle tracking in live-cell time-lapse sequences. *Nature methods*, 5(8), pp.695–702.

656 657 Kingma, D.P. & Optimization, J.B.A.A.M.F.S., *DP Kingma and J. Ba, Adam: A method for stochastic optimization*, *arXiv*: 1412.6980,

658 659 Koseki, H. et al., 2017. Selective rab11 transport and the intrinsic regenerative ability of CNS axons. *eLife*, 6, p.5546.

660 Lazarus, J.E. et al., 2013. Dynactin subunit p150(Glued) is a neuron-specific anti-  
661 catastrophe factor. D. Pellman, ed. *PLoS biology*, 11(7), p.e1001611.

662 LeCun, Y. et al., 2008. Backpropagation Applied to Handwritten Zip Code Recognition.  
663 [dx.doi.org](https://doi.org/10.1101/405183), 1(4), pp.541–551.

664 Lee, B.H. & Park, H.Y., 2018. HybTrack: A hybrid single particle tracking software using  
665 manual and automatic detection of dim signals. *Scientific reports*, 8(1), p.212.

666 Mangeol, P., Prevo, B. & Peterman, E.J.G., 2016. KymographClear and KymographDirect:  
667 two tools for the automated quantitative analysis of molecular and cellular dynamics  
668 using kymographs. *Molecular biology of the cell*, 27(12), pp.1948–1957.

669 Mathis, A. et al., 2018. Markerless tracking of user-defined features with deep learning.  
670 [arXiv.org](https://doi.org/10.1101/405183), cs.CV.

671 Mukherjee, A. et al., 2011. Automated kymograph analysis for profiling axonal transport of  
672 secretory granules. *Medical Image Analysis*, 15(3), pp.354–367.

673 Neumann, S. et al., 2017. KymoAnalyzer: a software tool for the quantitative analysis of  
674 intracellular transport in neurons. *Traffic*, 18(1), pp.71–88.

675 Racine, V. et al., 2007. Visualization and quantification of vesicle trafficking on a three-  
676 dimensional cytoskeleton network in living cells. *Journal of microscopy*, 225(Pt 3),  
677 pp.214–228.

678 Reis, G.F. et al., 2012. Molecular motor function in axonal transport *in vivo* probed by  
679 genetic and computational analysis in *Drosophila*. Y. Zheng, ed. *Molecular biology of the*  
680 *cell*, 23(9), pp.1700–1714.

681 Ronneberger, O., Fischer, P. & Brox, T., 2015. U-Net: Convolutional Networks for  
682 Biomedical Image Segmentation. In *Medical Image Computing and Computer-Assisted*  
683 *Intervention – MICCAI 2015*. Lecture Notes in Computer Science. Cham: Springer,  
684 Cham, pp. 234–241.

685 Sbalzarini, I.F. & Koumoutsakos, P., 2005. Feature point tracking and trajectory analysis for  
686 video imaging in cell biology. *Journal of Structural Biology*, 151(2), pp.182–195.

687 Szegedy, C. et al., 2014. Going Deeper with Convolutions. [arXiv.org](https://doi.org/10.1101/405183), cs.CV.

688 Tanenbaum, M.E., Vale, R.D. & McKenney, R.J., 2013. Cytoplasmic dynein crosslinks and  
689 slides anti-parallel microtubules using its two motor domains. *eLife*, 2, p.e00943.

690 Twelvetrees, A.E. et al., 2016. The Dynamic Localization of Cytoplasmic Dynein in Neurons  
691 Is Driven by Kinesin-1. *Neuron*, 90(5), pp.1000–1015.

692 Weigert, M. et al., 2017. Content-Aware Image Restoration: Pushing the Limits of  
693 Fluorescence Microscopy. *bioRxiv*, p.236463.

694 Zala, D. et al., 2013. Vesicular glycolysis provides on-board energy for fast axonal transport.  
695 *Cell*, 152(3), pp.479–491.

696  
697

698 **Figure legends**

699 **Figure 1: Kymograph generation and KymoButler**

700 **(A)** Schematic of kymograph generation from live imaging data. A cell and four particles are  
701 shown at 3 different timepoints (top row). A temporal projection of this cell highlights how  
702 each particle moves along a stationary path. It is possible to track the path (magenta line),  
703 and then extract the intensity of the particle in subsequent frames in a 2D kymograph image,  
704 where the horizontal and vertical axes represent space and time, respectively. Individual  
705 lines in a kymograph represent several particles moving along the same path. **(B)**  
706 Functionality of KymoButler. A kymograph, here the motion of mitochondria along neuronal  
707 dendrites adapted from (Faits et al. 2016), is uploaded via drag & drop to the cloud interface  
708 at <http://www.kymobutler.deepmirror.ai>, where the noise-dependent sensitivity can be  
709 manually adjusted. The outputs are: an overlay highlighting all the tracks found in different  
710 (random) colours, a .csv file with the time and space coordinates for each track, and a .csv  
711 file containing the summary of the direction and velocity of each track. **(C)** KymoButler image  
712 outputs from two example kymographs. Left: dynamics of fluorescently labelled Rab11a in  
713 rat cortical axons (adapted from (Koseki et al. 2017), bidirectional movement as Rab11a can  
714 move both ways in the axon or become stationary). Right: dynamics of fluorescently labelled  
715 microtubule plus-ends in mouse dorsal root ganglion axons (adapted from (Lazarus et al.  
716 2013), unidirectional movement since microtubule growth is continuous). The top row depicts  
717 the raw kymographs as taken from the published manuscripts. The middle row shows the  
718 identified tracks as dilated coloured lines. The bottom row depicts an overlay of the raw  
719 kymograph with the KymoButler prediction. Further examples from published work are  
720 shown in Figure 1-figure supplement 1A.

721 **Figure 1-figure supplement 1: Example kymographs and software workflow**

722 **(A)** Three example kymographs from published manuscripts. Example 1: *In vitro* dynamics of  
723 single cytoplasmic dynein proteins adapted from (Tanenbaum et al. 2013). Example 2: EB1-

724 GFP labelled growing microtubule plus-ends in mouse dorsal root ganglion axons (Lazarus  
725 et al. 2013). Example 3: Mitochondria dynamics in mouse retinal ganglion cell dendrites  
726 (Faits et al. 2016). Each dilated coloured line depicts an identified track. **(B)** KymoButler  
727 software workflow. First, a classification module is applied to each kymograph to determine  
728 whether the kymograph is unidirectional or bidirectional. If the kymograph is deemed  
729 unidirectional the unidirectional segmentation module is applied to the image to generate two  
730 trackness maps that assign each pixel a score between 0-1, approximating the likelihood  
731 that this pixel is part of a track with negative slope (left image) or positive slope (right image).  
732 Subsequently, the trackness maps are binarized, skeletonised, and segmented into their  
733 respective connected components. Finally, those components are averaged over each row  
734 to generate individual tracks, and a dilated representation of each track is plotted in a  
735 random colour. If the kymograph is classified as bidirectional, another segmentation module  
736 is applied to the kymograph, which generates a trackness map that does not highlight any  
737 particular slope. This map is binarized with a user-defined threshold and subsequently  
738 skeletonised, resulting in a binary map that exhibits multiple track crossings. To resolve  
739 these crossings, we first apply a morphological operation that detects the starting points of  
740 tracks in the binary map (red dots). Then, the algorithm tracks each line from its starting  
741 point until a crossing is encountered. At each crossing, the decision module is called, whose  
742 inputs are (i) the raw kymograph in that region, (ii) the previous track skeleton, and (iii) all  
743 possible tracks in that region. The decision module then generates another trackness map  
744 that assigns high values to the most likely future path from the crossing. This map is then  
745 again binarized and thinned with a fixed threshold of 0.5. If the predicted path is longer than  
746 2 pixels, the path tracking continues. Once all starting points have been tracked until an end  
747 (either no prediction or no further pixels available), the algorithm again looks for starting  
748 points in the skeletonised trackness map excluding the identified tracks, and repeats the  
749 steps outlined above until all pixels are occupied by a track. The resulting tracks are then  
750 drawn with each track in a random colour.

751 **Figure 1-figure supplement 2: The software modules in detail**

752 **(A)** The class module. This module resizes any input kymograph to 64x64 pixels. It  
753 subsequently applies two convBlocks with no padding and 64 output feature maps to the  
754 image. ConvBlocks comprise a convolutional layer with 3x3 kernels followed by a  
755 BatchNormalisation Layer and a leaky **R**ectified **L**inear **U**nit (ReLU) activation function (leak  
756 factor 0.1). The convBlocks are followed by 2x2 max pooling to halve the feature map sizes.  
757 This is repeated another 2 times while steadily increasing the number of feature maps until  
758 the last convBlock generates 256 feature maps of size 9x9. These maps are then pooled  
759 with a final 2x2 max pool operation followed by a 4x4 mean pool operation to generate a  
760 vector of 256 features. These features are then classified with a fully connected layer with  
761 output nodes followed by another leaky Ramp and finally another fully connected layer  
762 generates 2 output values that correspond to the probability of being a  
763 unidirectional/bidirectional kymograph. **(B)** The unidirectional segmentation module takes  
764 and an input kymograph of arbitrary size. Subsequently two convBlocks with 64 output  
765 feature maps are applied to the image followed by max pooling. This is repeated three times  
766 while doubling the number of feature maps with each pooling operation forming the  
767 “contracting path”. To obtain an image of the same size as the input image the small feature  
768 maps at the lowest level of the network have to be deconvolved 4 times each time halving  
769 the number of feature maps and applying further convBlocks. After each 2x2 deconvolution  
770 the resulting feature maps are catenated with the feature maps of the same size from the  
771 contracting path so that the network only learns residual alterations of the input image. The  
772 final 64 feature maps are linked to two independent convolutional layers that generate  
773 outputs that correspond to the trackness scores for positive and negative sloped lines. **(C)**  
774 The bidirectional segmentation module has the same architecture as the unidirectional one  
775 but only generates one output that corresponds to the trackness map for any lines in the  
776 image. **(D)** The decision module architecture is the same as the bidirectional segmentation  
777 module but takes three input images instead of one.

778 **Figure 1-figure supplement 3: Synthetic training data examples**

779 **(A)** Class module training data consisted of 64x64 pixel images that were either classified as  
780 unidirectional (example 1) or bidirectional (example 2). **(B)** Synthetic training data for the  
781 unidirectional segmentation module comprised 300x300 pixel kymographs with two binary  
782 ground truth maps, corresponding to particle motion with negative and positive slopes. **(C)**  
783 Synthetic bidirectional segmentation module training data comprises 300x300 pixel  
784 kymographs with only one ground truth image containing all ground truth tracks. **(D)** The  
785 decision module was trained with 48x48 pixel image crops of the raw kymograph, the  
786 previous skeletonised path, and all the skeletonised paths in the cropped region. The ground  
787 truth is simply the known future segment of the given path.

788 **Figure 1-figure supplement 4: Geometric mean of track recall and precision for  
789 different trackness thresholds**

790 **(A)** 10 synthetic unidirectional and bidirectional kymographs were analysed with varying  
791 trackness thresholds, and recall and precision were calculated. The geometric mean of recall  
792 and precision does not exhibit much variation between 0.1 and 0.3 but decreases at lower  
793 and higher values.

794

795 **Figure 2: Benchmark of KymoButler against unidirectional synthetic data**

796 **(A)** An example synthetic kymograph and its corresponding ground truth, manual control, the  
797 prediction by KymoButler, and the prediction by Fourier filtering. The top row depicts  
798 individual tracks in different colours and the bottom row shows the prediction overlay  
799 (magenta) with the ground truth (green) for all approaches. Discrepancies are thus  
800 highlighted in magenta (false positive) and green (false negative), while matching ground  
801 truth and prediction appears white. **(B)** Schematic explaining the concept of recall and  
802 precision. The top row depicts the possible deviations of the prediction from the ground truth.

803 The middle and bottom rows show example overlays, again in green and magenta, from the  
804 synthetic data. In the left column, the prediction is larger than the ground truth (magenta is  
805 visible) leading to false positive pixels and low track precision, but a small number of false  
806 negatives and thus high track recall. An example prediction overlay of the Fourier filter  
807 approach is shown, which tends to elongate track ends. The right column shows a shorter  
808 prediction than the ground truth, leading to green segments in the overlay. While this  
809 prediction has high track precision (low number of false positive pixels), track recall is low  
810 due to the large number of false negatives. Again, a cut-out from the Fourier filter prediction  
811 is shown, where multiple gaps are introduced in tracks, thus severely diminishing track recall  
812 (see Material and Methods for a detailed explanation of recall and precision). The middle  
813 column shows the same two cut outs analysed by KymoButler. No magenta or green  
814 segments are visible, thus leading to high recall and precision. **(C)** Synthetic kymograph  
815 region with four gaps highlighted (arrow heads): in one or more kymograph image rows the  
816 signal was artificially eliminated but kept in the ground truth to simulate real fluorescence  
817 data. While KymoButler efficiently connects tracks over gaps, the Fourier filter is unable to  
818 do so and breaks up those tracks into segments or incorrectly shortens these tracks (red  
819 arrow heads). Yellow arrow heads depict correct gap bridging events. **(D)** A synthetic  
820 kymograph with several line crossings. While KymoButler efficiently resolved all crossings,  
821 i.e. lines that cross other lines are not broken up into two segments, the Fourier filter  
822 correctly identifies the line crossing at the yellow arrow head but erroneously terminates the  
823 red and yellow tracks at the red arrow head. **(E)** The geometric means of recall and precision  
824 (“track F1 score”) for KymoButler, the Fourier filter approach, and manual control. Each dot  
825 represents the average track F1 score of one synthetic kymograph ( $p = 4 \cdot 10^{-5}$ , Kruskal-  
826 Wallis Test, Tukey post-hoc: manual vs KymoButler  $p = 0.6$ , manual vs Fourier Filtering  
827  $p = 3 \cdot 10^{-3}$ ). **(F)** Quantification of gap bridging performance for KymoButler (89%), manual  
828 control (88%), and Fourier filter (72%); lines: medians of all 10 synthetic kymographs,  
829  $p = 10^{-4}$ , Kruskal-Wallis Test, Tukey post-hoc: manual vs KymoButler  $p = 0.9$ , manual vs

830 Fourier Filtering  $p = 2 \cdot 10^{-3}$ . **(G)** The fraction of correctly identified crossings for  
831 KymoButler, manual annotation, and the Fourier filter (88% KymoButler, 86% manual, 60%  
832 Fourier filter; lines: medians of all 10 synthetic kymographs,  $p = 10^{-4}$ , Kruskal-Wallis Test,  
833 Tukey post-hoc: manual vs KymoButler  $p = 0.9$ , manual vs Fourier Filtering  $p = 1 \cdot 10^{-3}$ ).

834

835 **Figure 3: Benchmark of KymoButler against complex bidirectional synthetic data**

836 **(A)** Example synthetic kymograph and its corresponding ground truth, manual control, the  
837 prediction by KymoButler, and the prediction via wavelet coefficient filtering. The top row  
838 depicts individual tracks in different colours and the bottom row shows the prediction overlay  
839 (magenta) with the ground truth (green) for all approaches. Discrepancies are highlighted in  
840 magenta (false positive) and green (false negative), while the match of ground truth and  
841 prediction appears white. **(B)** Example recall and precision of KymoButler and wavelet  
842 filtering. While KymoButler shows high recall and high precision, the wavelet filter approach  
843 yields significant deviations from the ground truth (green and magenta pixels). **(C)** Synthetic  
844 kymograph region with three artificial gaps highlighted (arrow heads). While KymoButler  
845 efficiently connects tracks over gaps, the wavelet filter is unable to do so and breaks up  
846 those tracks into segments (red arrow heads). The yellow arrow heads depict correct gap  
847 bridging events. **(D)** A synthetic kymograph with several line crossings. While KymoButler  
848 efficiently resolved all crossings, i.e. lines that cross other lines are not broken up into  
849 segments, the wavelet filter only resolves one crossing correctly (yellow arrow head). **(E)**  
850 The geometric means of track recall and track precision (track F1 score) for KymoButler,  
851 manual control, and the wavelet filter. Each dot represents the average F1 score of one  
852 synthetic kymograph ( $p = 8 \cdot 10^{-5}$ , Kruskal-Wallis Test, Tukey post-hoc: manual vs  
853 KymoButler  $p = 0.7$ , manual vs wavelet filtering  $p = 10^{-4}$ ). **(F)** Quantification of gap  
854 performance for KymoButler, manual annotation, and wavelet filter ( $p = 3 \cdot 10^{-4}$ , Kruskal-  
855 Wallis Test, Tukey post-hoc: manual vs KymoButler  $p = 0.4$ , manual vs wavelet filtering

856  $p = 2 \cdot 10^{-4}$ ). **(G)** The fraction of resolved crossings for KymoButler, manual control, and the  
857 wavelet filter ( $p = 3 \cdot 10^{-5}$ , Kruskal-Wallis Test, Tukey post-hoc: manual vs KymoButler  
858  $p = 0.4$ , manual vs wavelet filtering  $p = 2 \cdot 10^{-5}$ ). KymoButler identifies tracks in complex  
859 kymographs as precisely as manual annotation by an expert.

860 **Figure 3-figure supplement 1: Performance of different skeletisation techniques on a**  
861 **synthetic bidirectional kymograph**

862 **(A)** Example of a synthetic bidirectional kymograph and its corresponding ground truth, the  
863 predictions by manual annotation, KymoButler, wavelet coefficient filtering, and tracks  
864 detected through edge filtering. The top row depicts individual tracks in different colours and  
865 the bottom row shows the prediction overlay (magenta) with the ground truth (green) for both  
866 approaches. Discrepancies are highlighted in magenta (false positive) and green (false  
867 negative), while a match of ground truth and prediction appears white.

868 **Figure 2-source data 1: Table of presented data.** A CSV file that contains: the average  
869 track F1 score, the average gap score, and the average crossing score for each  
870 unidirectional synthetic kymograph.

871 **Figure 2-source data 2: Synthetic kymographs and movies.** A ZIP file containing all  
872 analysed synthetic unidirectional movies, their kymographs, results from KymographClear  
873 based analysis and manually annotated ImageJ rois.

874 **Figure 3-source data 1: Table of presented data.** A CSV file that contains: the average  
875 track F1 score, the average gap score, and the average crossing score for each bidirectional  
876 synthetic kymograph.

877 **Figure 3-source data 2: Synthetic kymographs and movies.** A ZIP file containing all  
878 analysed synthetic bidirectional movies, their kymographs, and manually annotated ImageJ  
879 rois.

880

881

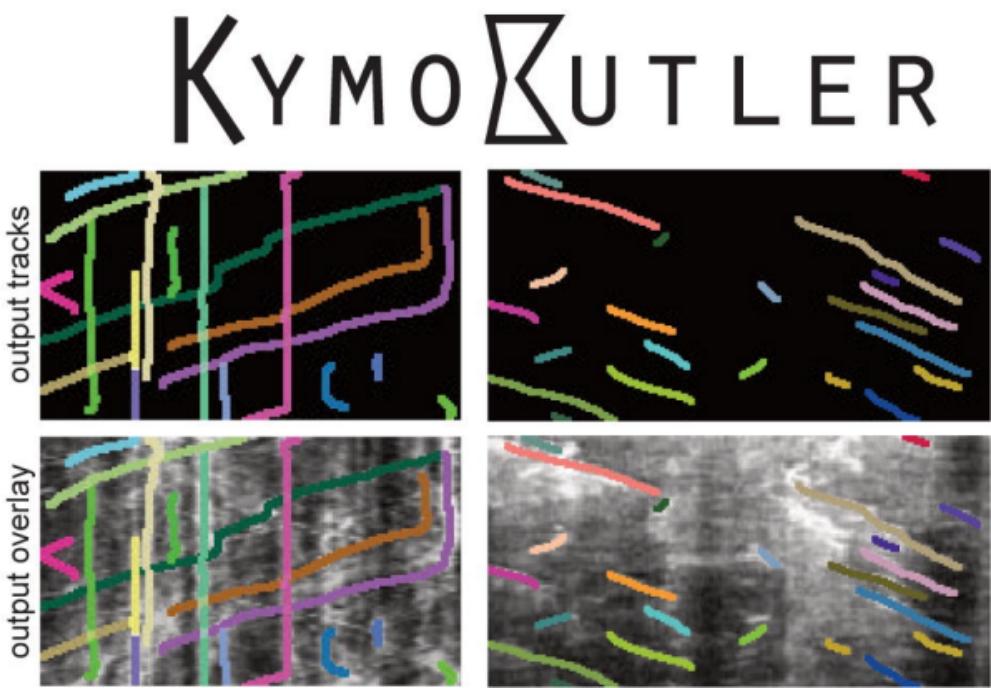
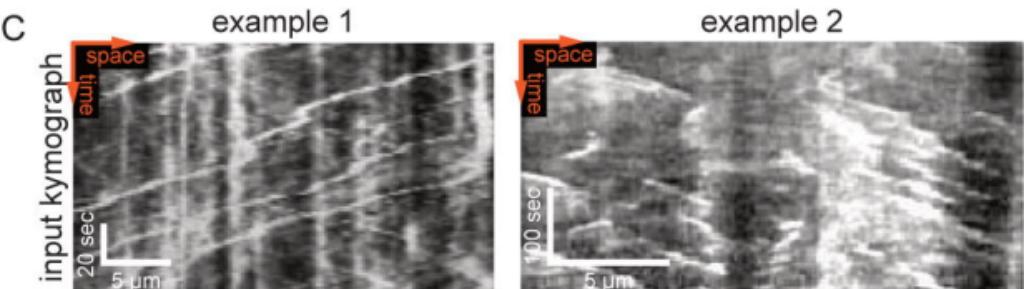
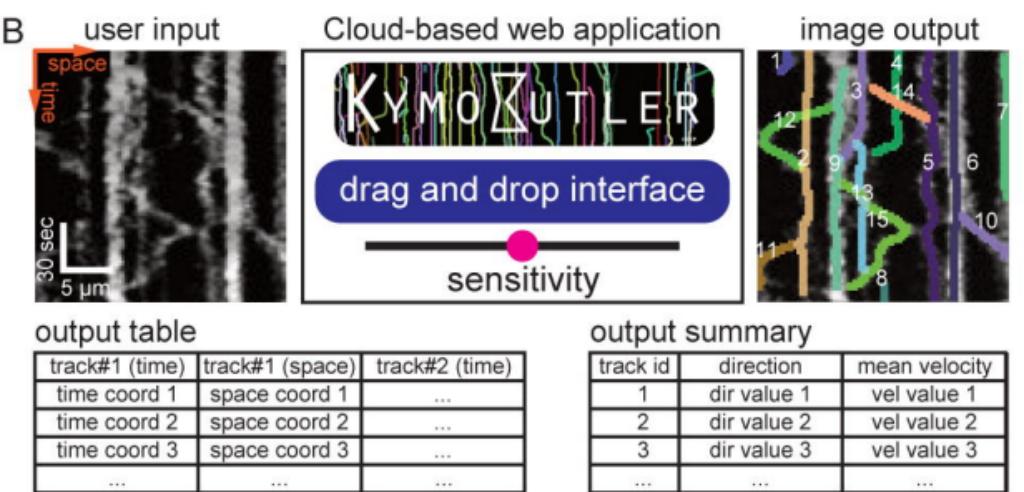
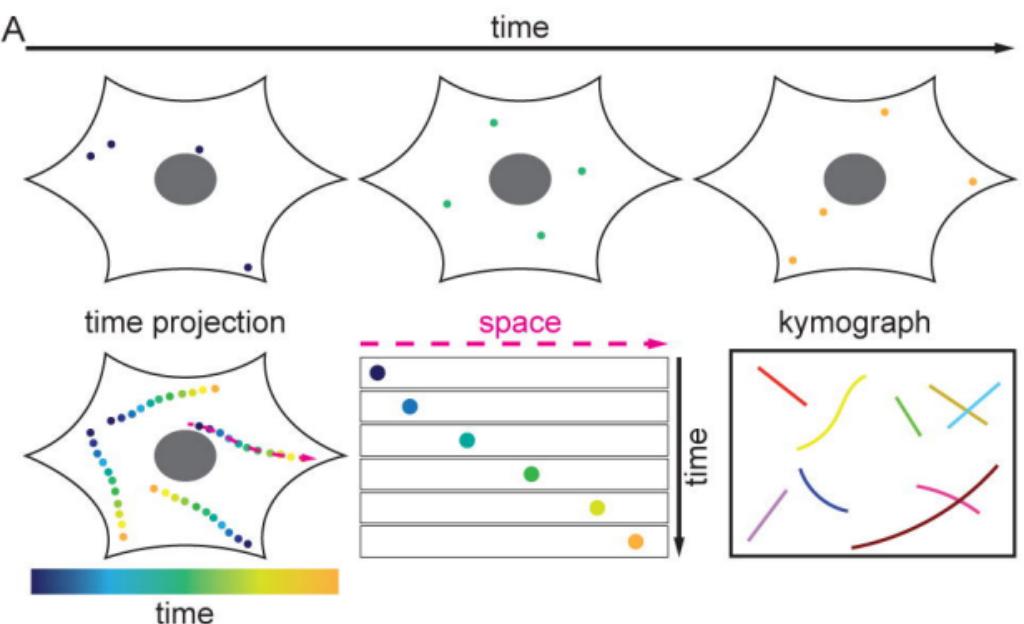
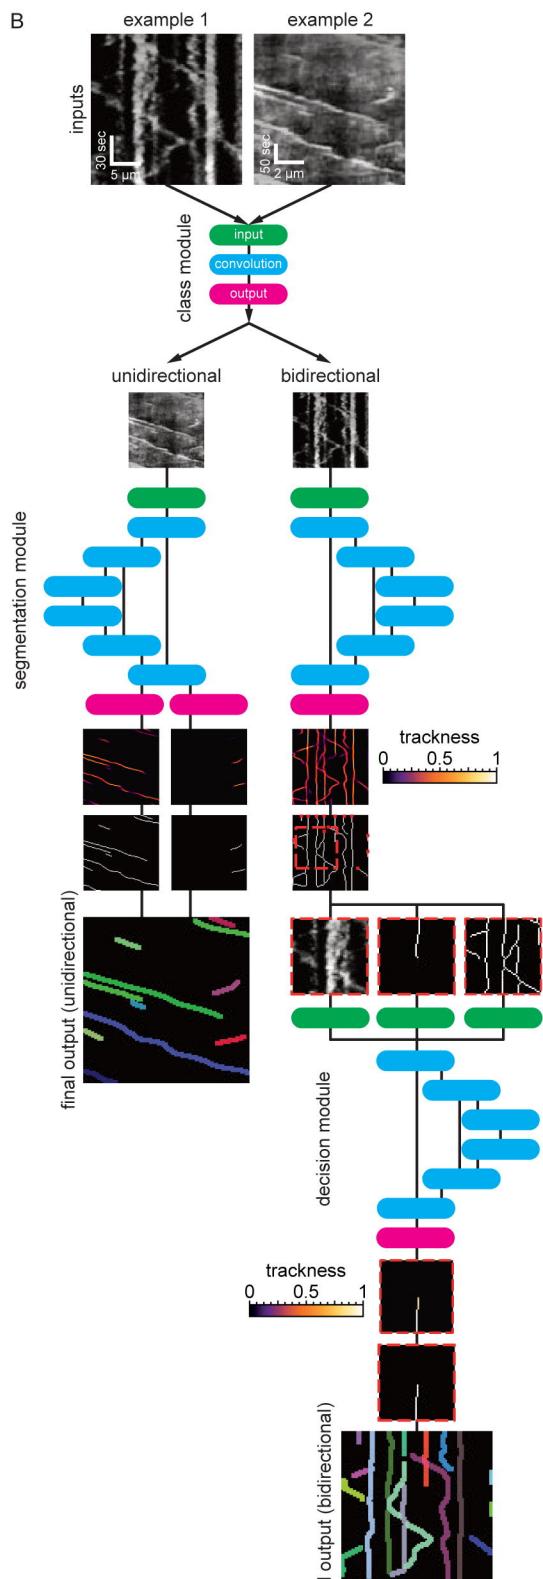
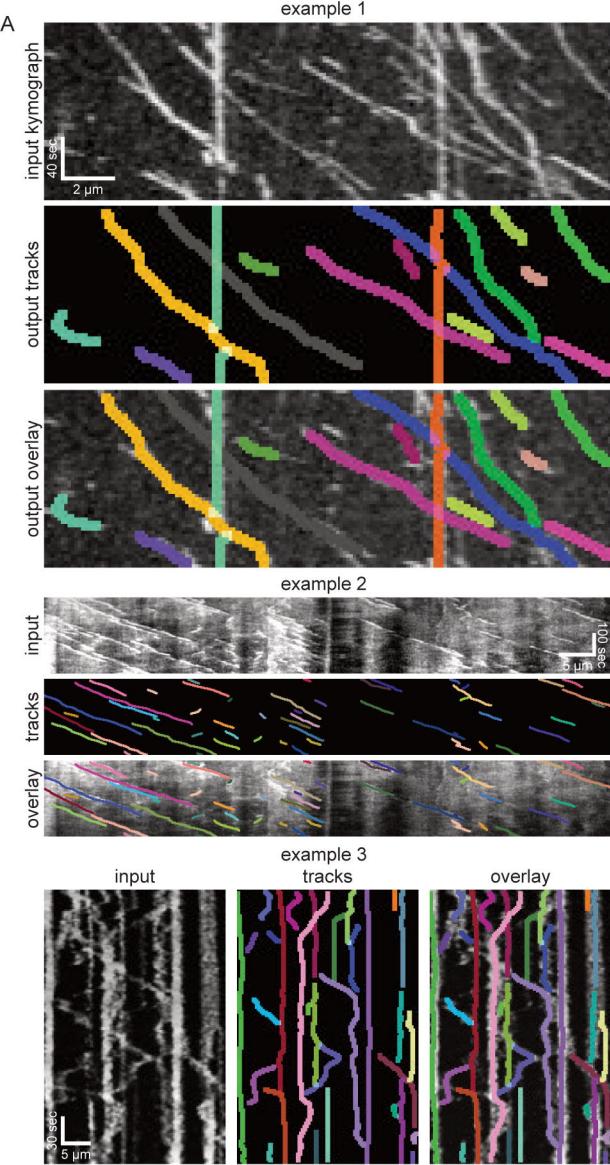
**Figure 1**

Figure 1 S1



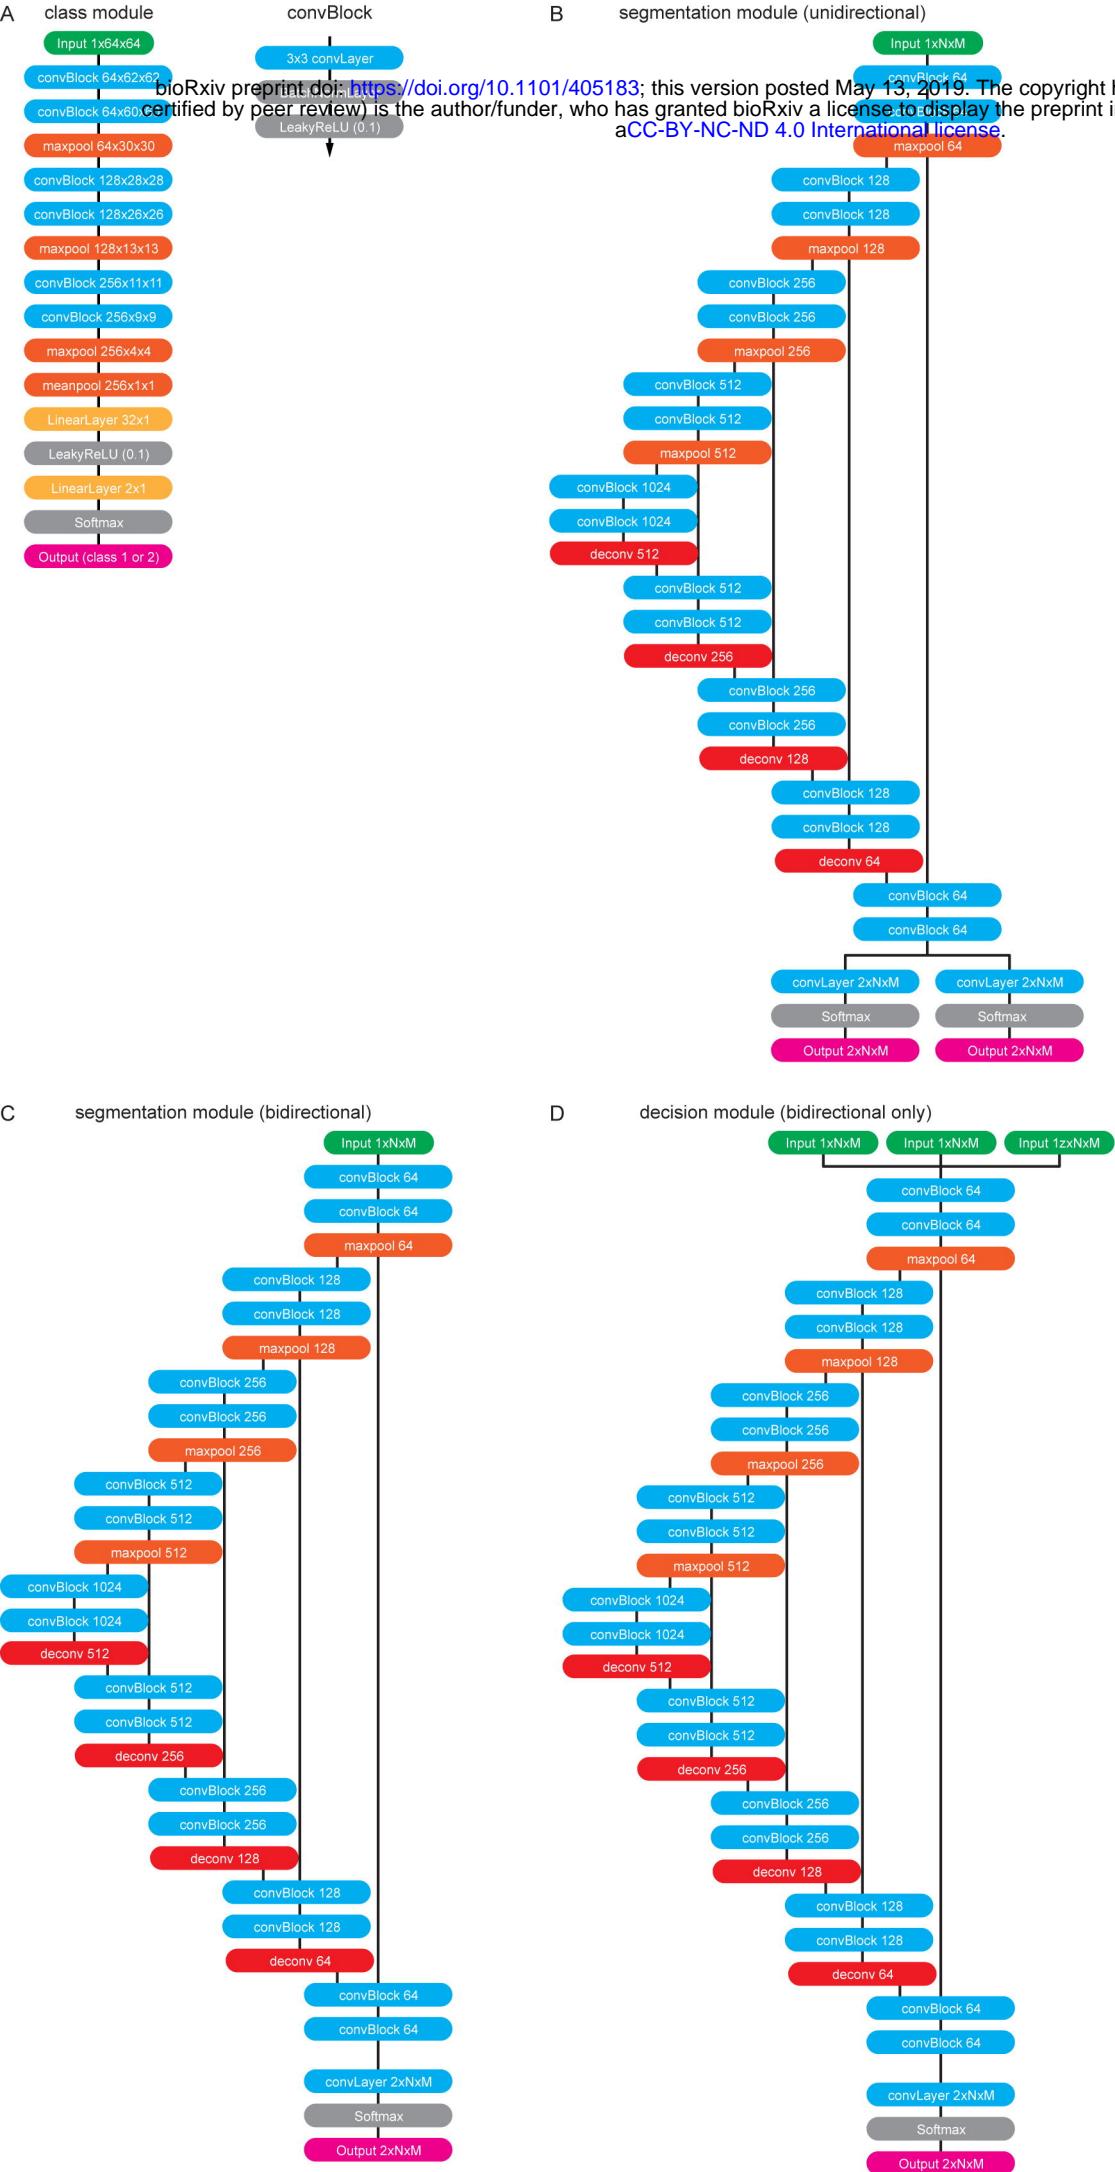
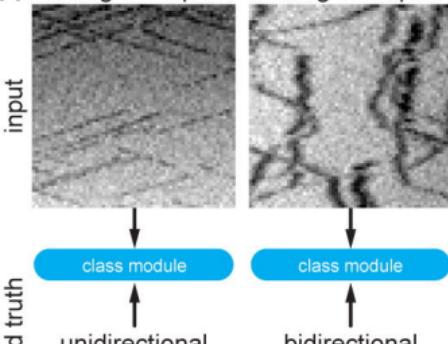
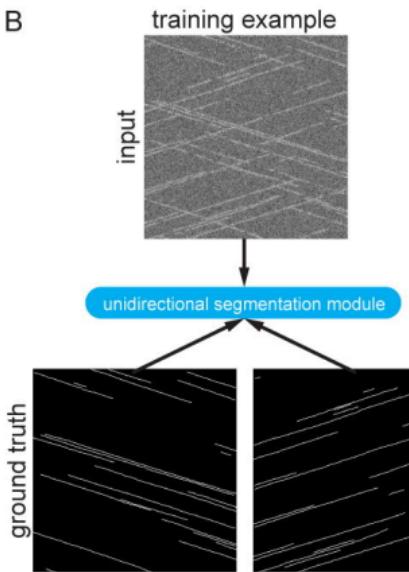


Figure 1 S3

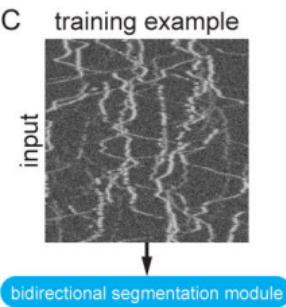
A training example 1 training example 2



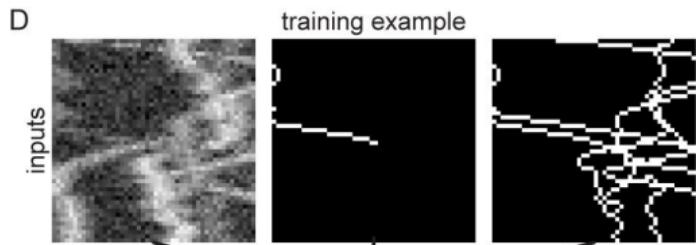
B



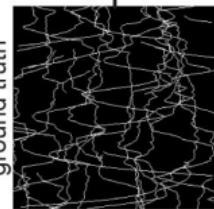
C training example



D



ground truth



ground truth

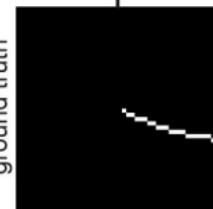


Figure 1 S4

A

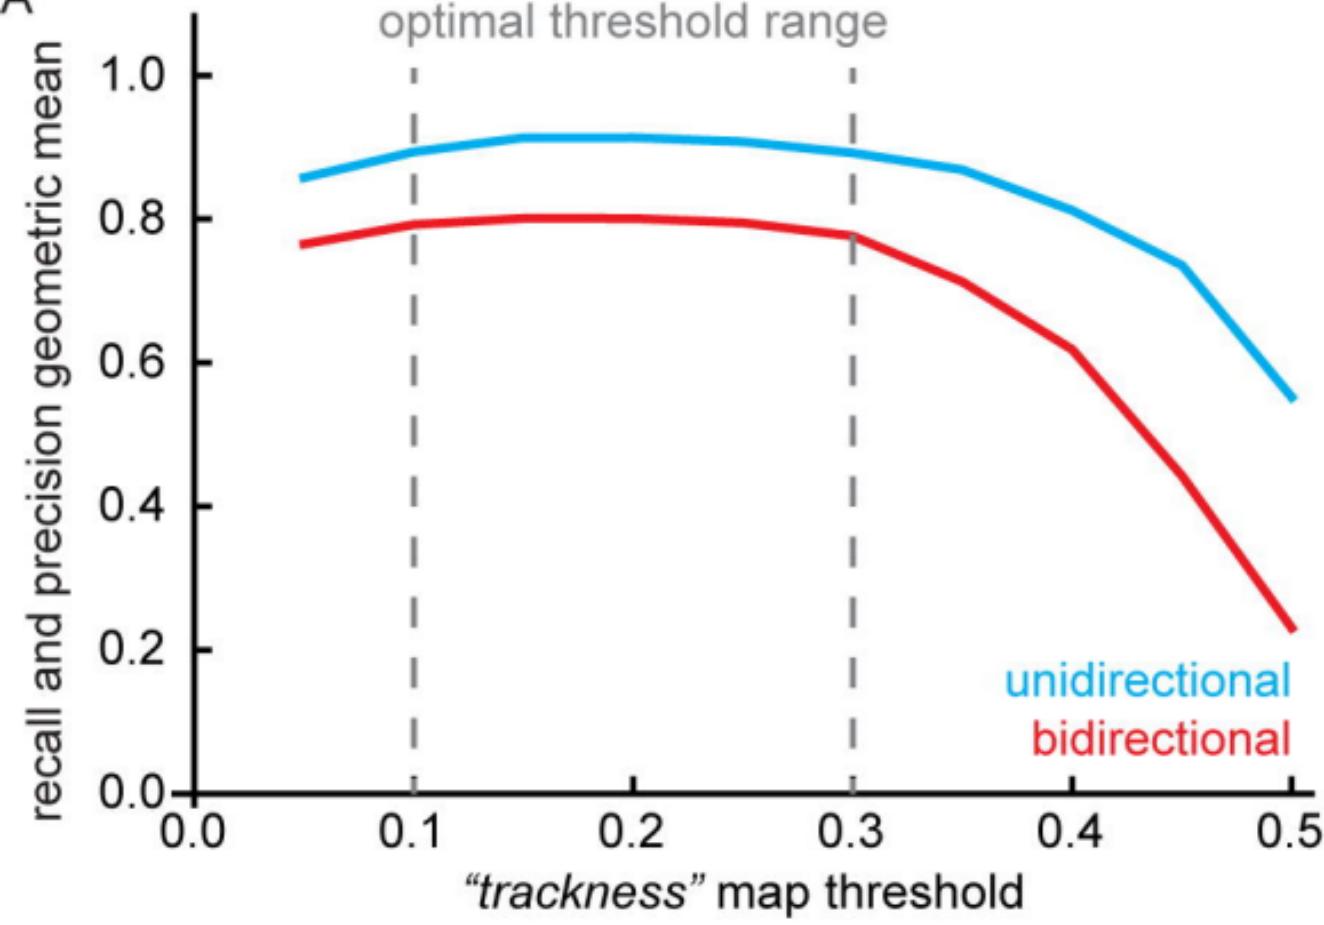


Figure 2

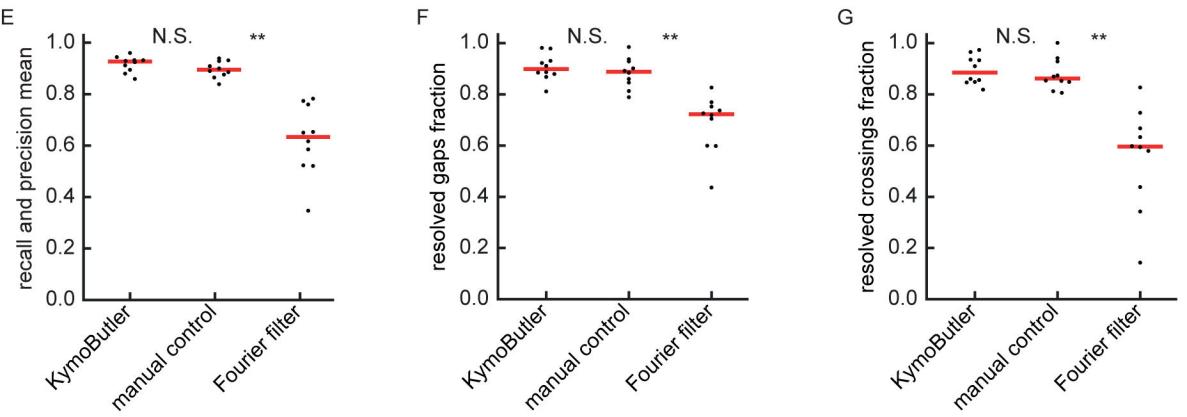
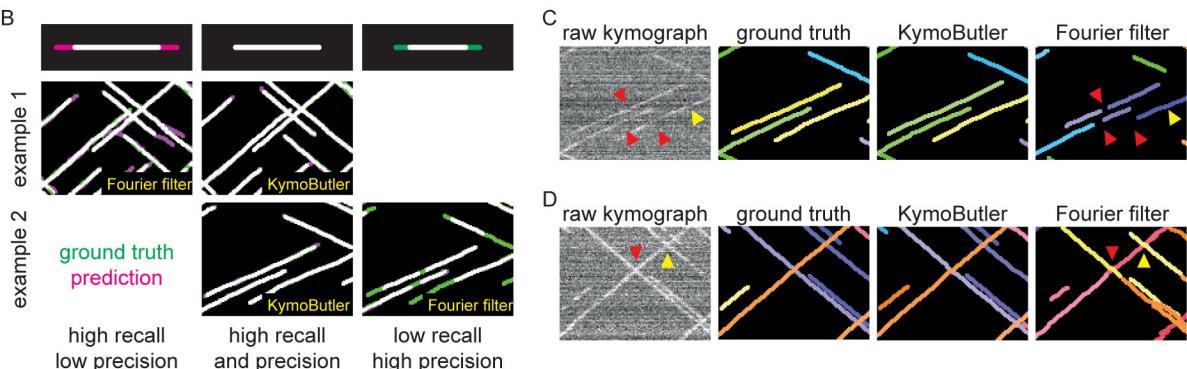
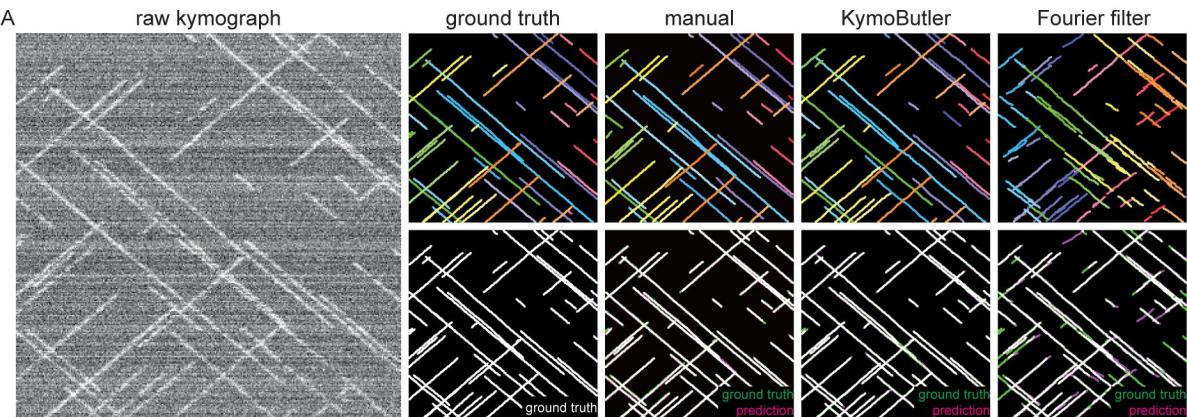


Figure 3

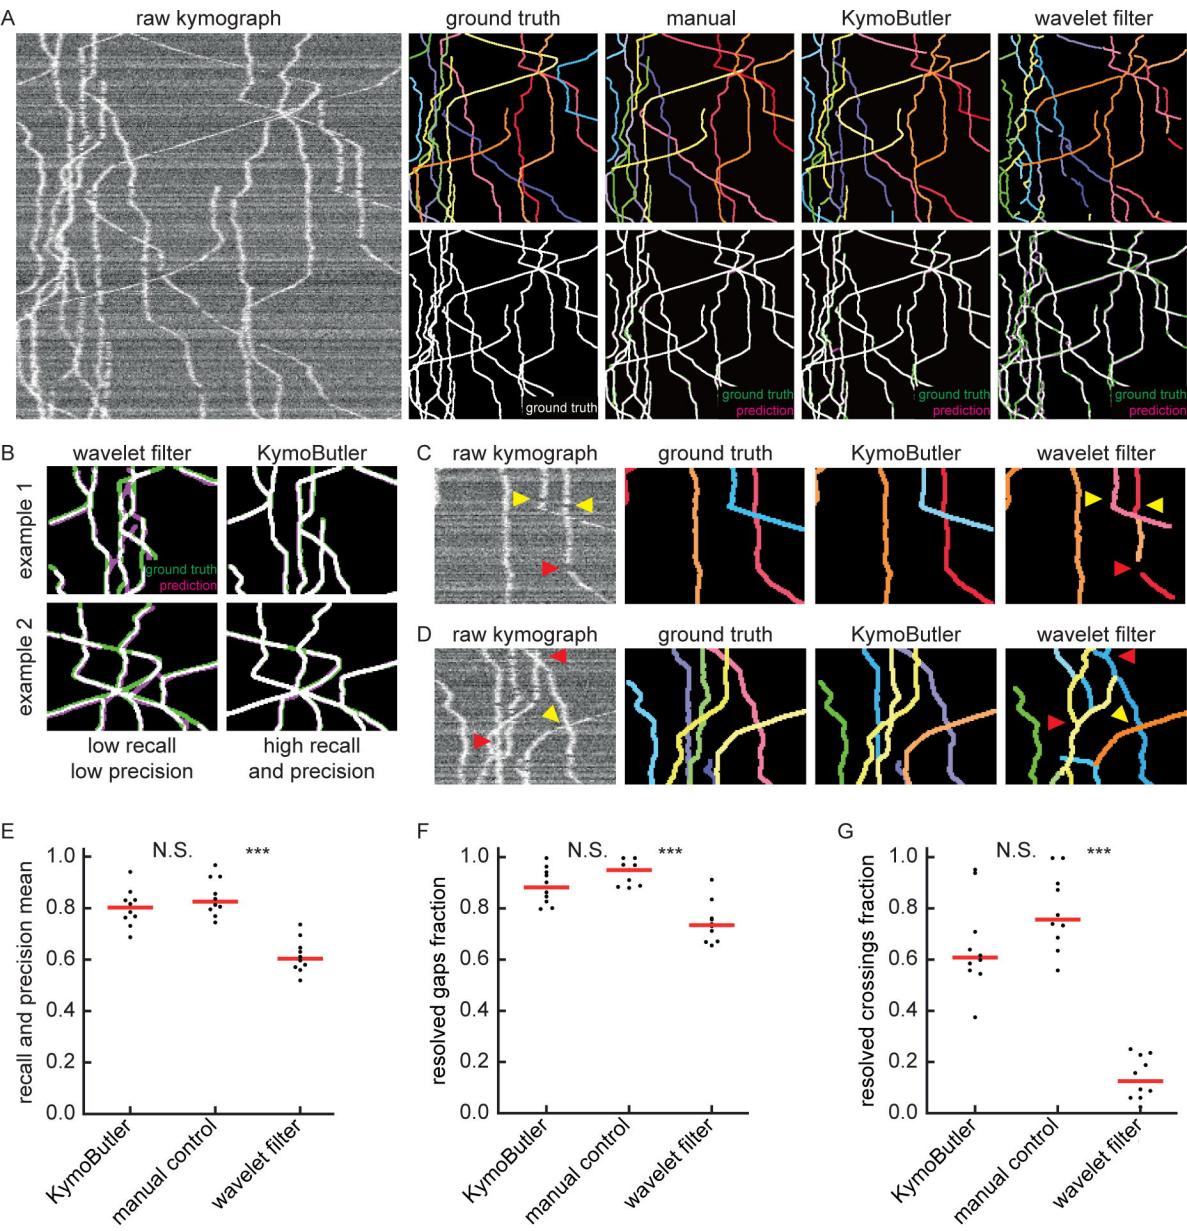
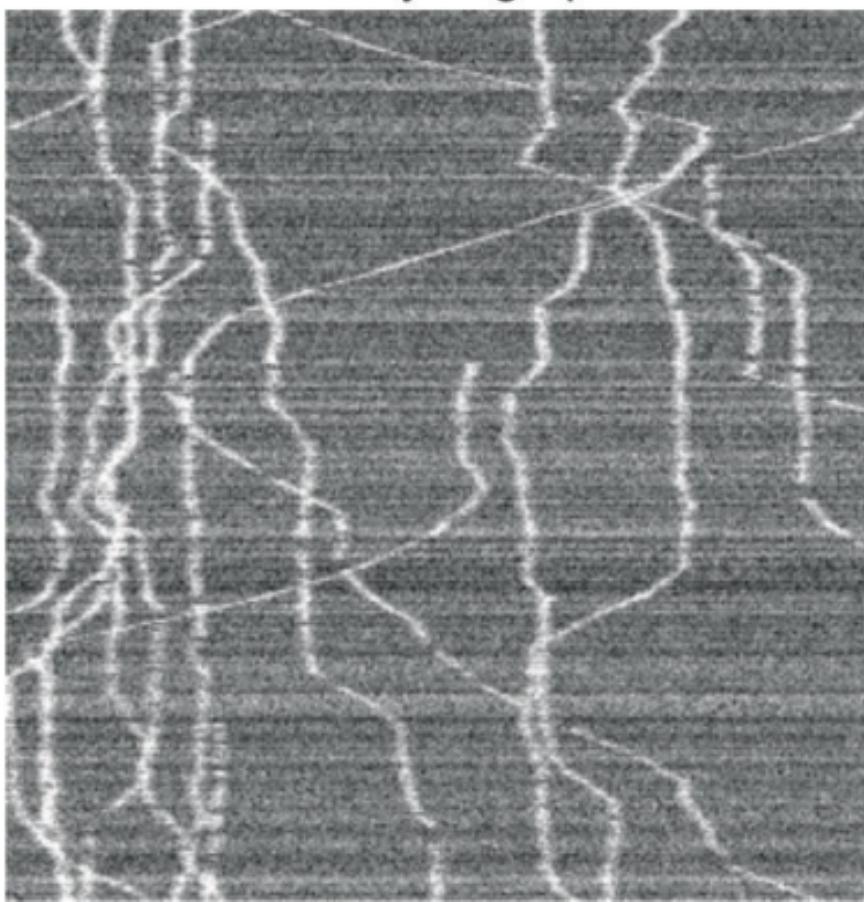


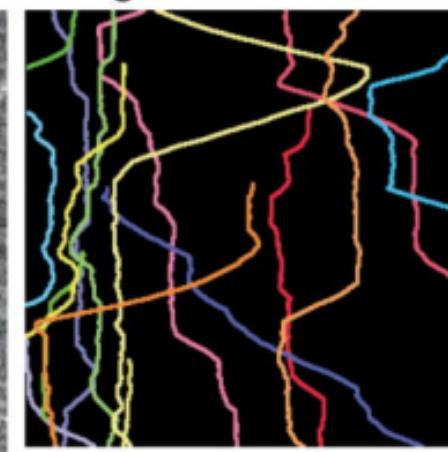
Figure 3 S1

A

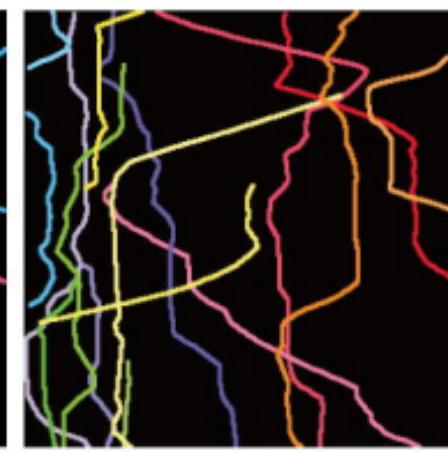
raw kymograph



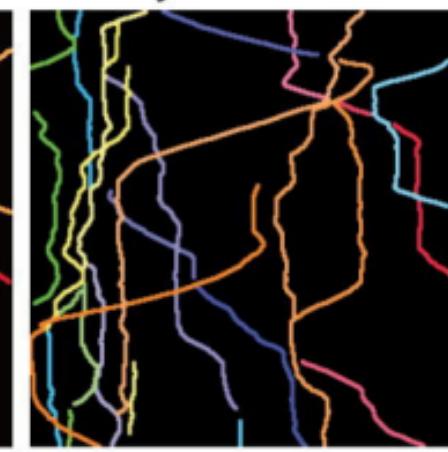
ground truth



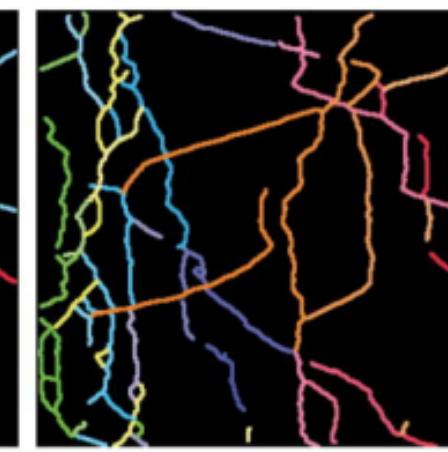
manual



KymoButler



wavelet filter



edge detection

