

Computational Tool for Ensemble Averaging of Single-Molecule Data

Thomas Blackwell,¹ W. Tom Stump,¹ Sarah R. Clippinger,¹ and Michael J. Greenberg^{1,*}

¹Department of Biochemistry and Molecular Biophysics, Washington University School of Medicine, St. Louis, Missouri

ABSTRACT Molecular motors couple chemical transitions to conformational changes that perform mechanical work in a wide variety of biological processes. Disruption of this coupling can lead to diseases, and therefore there is a need to accurately measure mechanochemical coupling in motors in both health and disease. Optical tweezers with nanometer spatial and millisecond temporal resolution have provided valuable insights into these processes. However, fluctuations due to Brownian motion can make it difficult to precisely resolve these conformational changes. One powerful analysis technique that has improved our ability to accurately measure mechanochemical coupling in motor proteins is ensemble averaging of individual trajectories. Here, we present a user-friendly computational tool, Software for Precise Analysis of Single Molecules (SPASM), for generating ensemble averages of single-molecule data. This tool utilizes several conceptual advances, including optimized procedures for identifying single-molecule interactions and the implementation of a change-point algorithm, to more precisely resolve molecular transitions. Using both simulated and experimental data, we demonstrate that these advances allow for accurate determination of the mechanics and kinetics of the myosin working stroke with a smaller set of data. Importantly, we provide our open-source MATLAB-based program with a graphical user interface that enables others to readily apply these advances to the analysis of their own data.

SIGNIFICANCE Single-molecule optical trapping experiments have given unprecedented insights into the mechanisms of molecular machines. Analysis of these experiments is often challenging because Brownian-motion-induced fluctuations introduce noise that can obscure molecular motions. A powerful technique for analyzing these noisy traces is ensemble averaging of individual binding interactions, which can uncover information about the mechanics and kinetics of molecular motions that are typically obscured by Brownian motion. Here, we provide an open-source, easy-to-use computational tool, SPASM, with a graphical user interface for ensemble averaging of single-molecule data. This computational tool utilizes several conceptual advances that significantly improve the accuracy and resolution of ensemble averages, enabling the generation of high-resolution averages from a smaller number of binding interactions.

INTRODUCTION

Molecular motors generate force and movement in a wide array of cellular processes, including muscle contraction, packaging of DNA into viral capsids, intracellular transport, DNA damage repair, and cell motility (1). These motors have complex mechanochemical cycles in which chemical transitions are coupled to conformational changes in the protein structure that generate mechanical work. The kinetics and mechanics of these transitions are tuned to the specific molecular role of the motor in the cell, and subtle changes in these properties can lead to an array of dis-

eases (2). Therefore, there is a need for experimental and computational techniques for probing these relationships.

Single-molecule optical trapping techniques with nanometer spatial and millisecond temporal resolution have proven to be powerful tools for studying the mechanochemical coupling in motors. One widely used optical trapping technique is the three-bead assay (Fig. 1 A; (3,4)). In this assay, two beads are held in place by dual-beam optical tweezers. The motor's track (e.g., actin) is strung between these beads and then lowered onto a third, surface-bound bead. This third bead is sparsely coated with motor molecules (e.g., myosin) such that only a single motor interacts with the track at any given time. The positions of the two optically trapped beads are monitored to study the interactions between the motor and the track (Fig. 1 B), where motor binding to the track causes both displacement

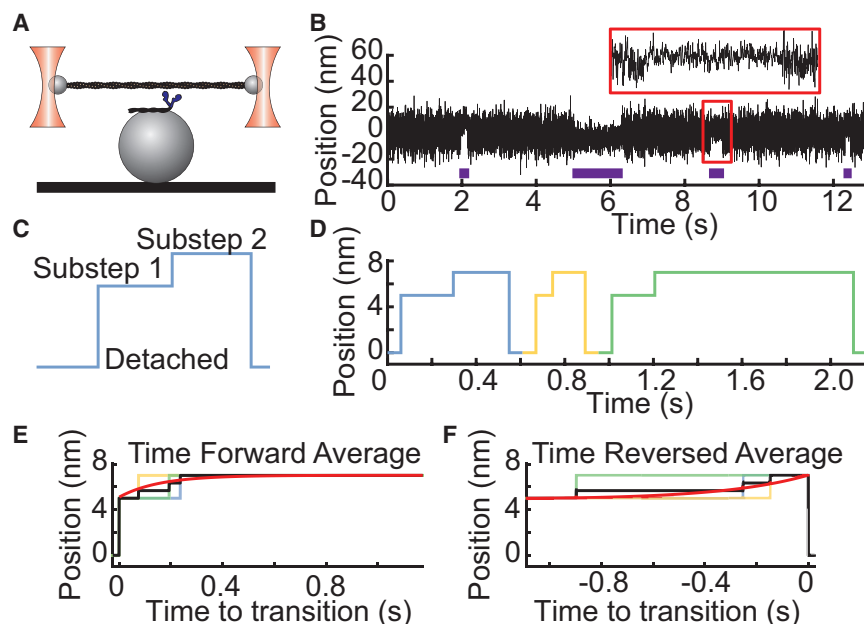
Submitted August 11, 2020, and accepted for publication October 9, 2020.

*Correspondence: greenberg@wustl.edu

Editor: David Thomas.

<https://doi.org/10.1016/j.bpj.2020.10.047>

© 2020 Biophysical Society.



average is fit with a single exponential function (red). The y offset and amplitude of this exponential provide estimates of the average size of the first and second substeps, respectively. The rate of this exponential gives the rate of transitioning from the first substep to the second substep. (F) The procedure for generating time-reversed ensemble averages from individual binding interactions is shown. Individual trajectories are aligned upon dissociation and averaged backward in time (black), and the average is fit with a single exponential function (red). The y offset and amplitude of this exponential provide estimates of the average size of the total step and the second substep, respectively. The rate of this exponential gives the rate of transitioning from the second substep to the detached state. To see this figure in color, go online.

of the beads as well as a reduction in the bead variance. This assay has been applied to study several motor and nonmotor systems, including dynein (5), the lac repressor (6), kinesins (7), and several myosin isoforms (8–14).

Analysis of the individual time-dependent trajectories of motor-induced displacements in the bead positions can provide information about both the mechanics and the kinetics of the motor's mechanochemical cycle. However, it can be difficult to resolve details of these trajectories because the amplitude of Brownian-motion-induced fluctuations in the bead position is frequently larger than the size of motor-induced displacements. One powerful method for extracting high spatial and temporal resolution information from noisy traces is postsynchronization ensemble averaging (13,15). In this method, trajectories from multiple individual binding interactions are aligned and then averaged together, thereby increasing the signal/noise ratio. This technique has been applied to successfully identify substeps of the myosin working stroke (12,13,16) and transitions in the ribosome (15) that likely would have been obscured using other analysis methods. Although this is a powerful tool for analyzing single-molecule data, there is no software in the public domain that is tailored to performing these calculations, and this has limited the adoption of these tools by many groups.

We have developed a MATLAB-based computational tool, Software for Precise Analysis of Single Molecules

(SPASM), with a graphical user interface for the identification and ensemble averaging of single-molecule trajectories. This computational tool utilizes several conceptual advances, including an optimized method for identifying binding interactions from noisy data and improved precision in determining the exact initiation and termination times of binding interactions using a change-point algorithm. Using both simulated and experimental data sets, we demonstrate that these advances permit the generation of accurate, high-resolution ensemble averages using fewer individual binding trajectories than were previously required. Our easy-to-use computational tool includes an intuitive graphical user interface and is offered both as open-source code and as a stand-alone program that does not require full installation of MATLAB. Finally, we provide a user guide, a separate tool for simulating data, and sample data sets to help other researchers apply this tool to their own single-molecule data.

MATERIALS AND METHODS

Implementation of the computational tool

The SPASM computational tool, which includes a graphical user interface and a tool for simulating data, was written in MATLAB (The MathWorks). The details of the implementation can be found in the [Supporting Material](#) (code availability: <https://github.com/GreenbergLab/SPASM>).

FIGURE 1 Ensemble averaging of optical trapping data enables the study of mechanochemical coupling. (A) A diagram of the three-bead assay, in which an actin filament strung between the two optically trapped beads is lowered onto a third surface-bound bead that is sparsely coated with myosin, is shown. (B) Single-molecule binding interactions between cardiac myosin and actin at $1 \mu\text{M}$ ATP recorded in the optical trap are shown. The average position between the optically trapped beads is plotted as a function of time, with blue horizontal bars indicating detected binding interactions. The mean position and variance of the beads change upon binding. A single binding interaction, shown in the red box, is expanded. Brownian motion obscures the second substep of the working stroke. (C) A schematic shows the two substeps of the myosin working stroke. (D) Idealized trace shows the position over time of a motor with a two-substep working stroke without Brownian motion. (E) The procedure for generating time-forward ensemble averages from individual binding interactions is shown. Individual trajectories are aligned at the initiation of binding and averaged forward in time (black line), and the

Design of optical trapping apparatus

Experiments were performed on a custom-built, microscope-free, dual-beam optical trap. The optical layout is described in the [Supporting Material](#) (Fig. S1). Porcine cardiac myosin and actin were purified from cryoground tissue (Pelfreez) as previously described (17,18). Trapping experiments were conducted as previously described (3). Details can be found in the [Supporting Material](#).

RESULTS AND DISCUSSION

Ensemble averaging of single-molecule binding interactions

Although ensemble averaging techniques are broadly applicable, we will focus in this article on their application to studying the interaction between myosin molecular motors and actin. Using ensemble averaging of optical trapping data, it has been shown that many myosin isoforms have a two-substep working stroke in which the first substep corresponds to the release of inorganic phosphate and the second substep corresponds to a transition associated with ADP release (Fig. 1, C and D; (8–10,12–14,16,19,20)). It is difficult to distinguish the second transition from raw data traces because of Brownian motion. However, ensemble averaging allows for easier visualization of this transition by increasing the signal/noise ratio.

One can collect information about both the kinetics and mechanics of the working stroke substeps from the post-synchronized ensemble-averaged trajectories of individual binding interactions (13,15). These interactions can be synchronized upon actomyosin attachment and then averaged forward in time or, alternatively, synchronized upon actomyosin detachment and then averaged backward in time

(Fig. 1, E and F). The magnitude of the initial displacement seen in the time-forward averages gives the size of the first substep of the myosin working stroke, a transition which occurs within the dead time of typical optical tweezer instruments (6). The amplitude of the subsequent exponential rise in displacement in the time-forward averages gives the size of the second substep of the working stroke. The rate of this exponential rise is the rate of transitioning from the first substep to the second substep, and it is associated with ADP release in myosins (13). For the time-reversed ensemble averages, the exponential rise in displacement before detachment has an amplitude equal to the size of the second substep, and the rate of this exponential gives the rate of transitioning from the second substep to the detached state, a transition that corresponds to ATP-induced actomyosin dissociation (13).

Generation of a covariance histogram to identify binding interactions

The first step in generating ensemble averages is the identification of binding interactions from single-molecule data traces. When optically trapped, the two beads in the bead-actin-bead dumbbell undergo fluctuations in their position due to Brownian motion (Fig. 2 A). The motion of these beads is mechanically coupled through the actin filament, as evidenced by the covariance between their positions (Fig. 2 B). When the surface-bound motor binds to the actin filament, it causes several pronounced changes: 1) it reduces the positional variance of each bead's position, 2) it reduces the coupled motion (covariance) of the two trapped beads, and 3) it displaces the mean position of each bead. The

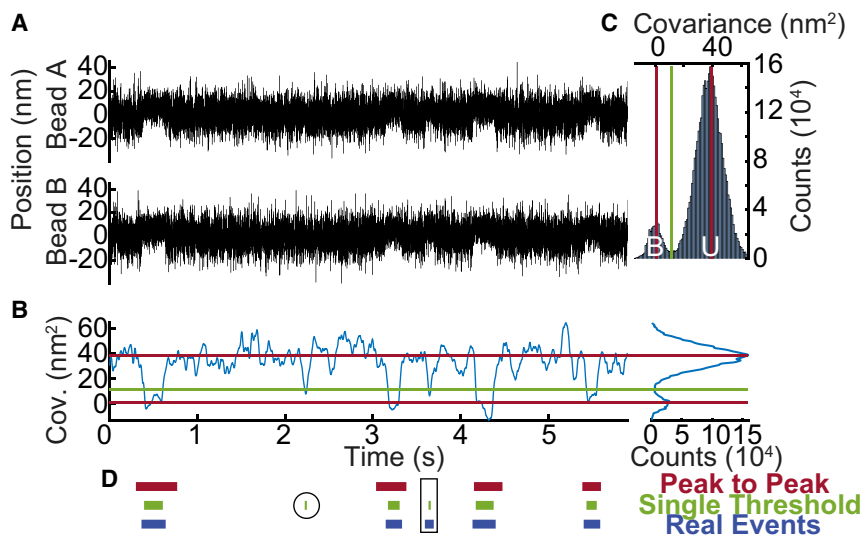


FIGURE 2 Detection of binding interactions using either the single or peak-to-peak covariance threshold method. (A) Simulated optical trapping data show the position of each optically trapped bead over time. (B) Covariance between the position of the optically trapped beads at each time point gives rise to a bimodal distribution. (C) A histogram of covariance values shows two distinct populations that correspond to the bound (B) and unbound (U) states. In the single-threshold method, a binding interaction is detected when the covariance crosses the value located at the minimum between the two populations (green). In the peak-to-peak method, two thresholds are placed, one at the peak of each population (red), and a binding interaction is detected when the covariance transitions from one threshold to the other threshold. (D) Simulated binding interactions detected by the peak-to-peak method (red), binding interactions detected by the single-threshold method (green), and actual simulated binding interactions (real, blue) are shown. The single threshold is more susceptible to false-positive interactions (circled). The peak-to-peak method is more susceptible to false-negative interactions (boxed). To see this figure in color, go online.

majority of analysis methods for identifying binding interactions utilize the changes in the mean position, variance, and/or covariance of the optically trapped beads upon binding of myosin to actin (11,21–23).

One popular method for selecting binding interactions is to set a threshold based on the variance or covariance of the beads. The choice of using a variance or covariance threshold for binding interaction identification will partially be dictated by the optical trap layout. For systems that only monitor the position of a single bead, one must use a variance threshold for the position of the single bead. For systems that monitor both bead positions, a covariance threshold is preferred because it is less sensitive to noisy fluctuations in the data. Although we focus on the use of our computational tool with a covariance threshold, the same approaches and conclusions will hold true for a variance threshold based on the position of one bead. A version of SPASM that uses a variance threshold is provided (see [Supporting Material](#)).

Our computational tool identifies binding interactions from the change in the covariance between the positions of the two trapped beads that occurs upon myosin binding to actin. SPASM first calculates the covariance over a sliding window in time and then smooths the covariance over a separate window. With properly chosen window lengths, the histogram of the covariance values reveals two populations (Fig. 2 C), with the higher covariance population corresponding to unbound states and the lower population corresponding to bound states (3). One can then select binding interactions based upon thresholds that distinguish between these two populations (see [Selection of Binding Interactions](#) below).

The success of this approach depends on the degree of separation between the two peaks in the covariance histogram. If the peaks are not well separated, the analysis is more susceptible to false and/or missed binding interactions. The ability to generate a histogram with two well-separated peaks depends partly on the selection of proper window lengths for the calculation and smoothing of the covariance. Optimal values for these parameters, in turn, depend on the kinetics of the myosin's interaction with actin, the compliance of the myosin and/or myosin-surface attachment, the pretension between the optically trapped beads, and the noise in the system. Therefore, the window lengths often need to be determined empirically. If the kinetics of the myosin's transitions are known from other experimental measurements, one can simulate data and select window lengths that optimize analysis of the simulated data (see [Supporting Material](#)). If kinetic information about the myosin's transitions is unknown, it may not be possible to generate meaningful simulated data. In these cases, the window lengths can be determined empirically through the computational tool's graphical user interface, which allows the user to vary the window lengths until a suitable bimodal covariance histogram is achieved.

Selection of binding interactions

Once a suitable covariance histogram with two well-defined peaks has been generated, the next step is to determine proper thresholds for the covariance that will be used to detect binding interactions. One possibility for distinguishing the bound state from the unbound state is to use a single covariance threshold located at the minimal value between the two peaks of the covariance histogram (10). Here, detected interactions start when the covariance drops below this threshold value, and they end when the covariance rises back above this threshold value (Fig. 2 D). Alternatively, one could identify the binding interactions using a set of two different covariance thresholds, located at the two peaks of the covariance histogram. In this “peak-to-peak” approach, a binding interaction is considered to start when the covariance drops from the threshold defined by the unbound peak to the threshold defined by the bound peak. Likewise, a binding interaction is considered to end when the covariance rises from the threshold defined by the bound peak to the threshold defined by the unbound peak (Fig. 2 D).

We tested the abilities of the single-threshold and peak-to-peak methods to accurately detect simulated binding interactions between actin and cardiac myosin. Interactions were simulated using a continuous-time Markov jump process with kinetics and mechanics based on previously measured parameters for ventricular cardiac myosin (8,24,25) (see [Materials and Methods](#) for details). With simulated data, the exact locations of the binding interactions are known, allowing for easy comparison between the simulated interactions and the interactions detected by the computational tool using either method (Fig. 2 D).

We generated 10 independent sets of simulated data, each containing 100 binding interactions (sets 1–10). For each data set, we used our computational tool to calculate the covariance histogram, locate the peaks and minimum of the histogram, and identify binding interactions using either the single-threshold method or the peak-to-peak method. When we used a single threshold to identify binding interactions, we correctly detected 80 ± 4 of the 100 binding interactions on average, and we incorrectly detected 4 ± 1 false-positive binding interactions per 100 s of data, on average (Table S1). The reported errors are standard deviations. When we used the peak-to-peak method to identify binding interactions, we correctly detected 65 ± 5 of the 100 binding interactions on average, and we did not detect any false-positive binding interactions. Although the peak-to-peak method misses a greater number of binding interactions, the false-positive rate is lower for this method ($p < 0.001$).

A single threshold could work well for selecting binding interactions if the two populations of the histogram are sufficiently distinct. However, it is often not possible to obtain sufficient separation between the peaks because of factors

that lower the signal/noise ratio (e.g., system noise, insufficient pretension between the beads, fast binding kinetics). In these cases, this single-threshold approach is prone to identifying false-positive interactions in which the covariance crosses the threshold even though the actomyosin has remained in an unbound state. These false-positive binding interactions do not generate a net displacement in the optical trap, so their inclusion in the ensemble averages is expected to lead to an underestimation of the true size of the working stroke. A methodology has been developed that attempts to correct for these false-positive interactions through the use of normalization factors (10). Alternatively, because the vast majority of these false-positive interactions arise as a result of either Brownian-motion-induced (or system-noise-induced) rapid downward spikes in the covariance (which lead to very short detected interactions) or rapid upward spikes in the covariance (which lead to multiple detected interactions in quick succession), it is possible to avoid these false-positive interactions through the use of temporal filters that exclude interactions that are too short or pairs of interactions that are too close to one another. However, it is not always easy to determine appropriate values for these temporal filters. Further, the use of these temporal filters may lead to the exclusion of many correctly identified binding interactions. When we used optimized values for these filters to exclude all of the false-positive interactions in the simulated data that were detected by the single-threshold method, we were left with fewer interactions than were detected by the peak-to-peak method (Fig. S2). The optimal parameter values to remove false-positive binding interactions will depend on the signal/noise ratio of the data (Fig. S3), and the user will need to optimize these values for their own data using SPASM.

With the peak-to-peak method, the criteria for detecting a binding interaction is much stricter than with the single-threshold method, and the number of identified false-positive binding interactions is expected to decrease while the number of missed, short-lived binding interactions increases. Unlike the inclusion of false-positive interactions, the exclusion of these missed binding interactions does not adversely affect the size or shape of the ensemble averages. Although we demonstrate that the peak-to-peak method performs better in data traces with moderate separation between the peaks of the covariance histogram, some experimental data might have better peak separation. In this case, the single-threshold method would be preferable because it maximizes the number of captured binding interactions. The computational tool allows the user to try both methods, and it automatically determines appropriate values for the thresholds.

Alignment of binding interactions using covariance thresholds

After binding interactions are identified, they must be precisely aligned at the transitions between the bound and

unbound states to generate accurate ensemble averages. The most critical step in aligning these interactions is the careful determination of when exactly a transition occurs. Inadequate determination of these transitions will lead to inaccurate measurements of the substep sizes and/or kinetics. Several methods have been applied to locate transitions in single-molecule data traces, including hidden Markov models (23) and step-finding algorithms (26), but a frequently used method for postsynchronization is to align the binding interactions based on the same thresholds used to identify the binding interactions (3,10,13).

To test the abilities of the single-threshold and peak-to-peak methods to accurately identify the transitions, we used the same 10 simulated data sets containing 100 transitions each as described previously (sets 1–10). When we used a single threshold to identify transition times, we found that the detected attachment times occurred 28.2 (95% confidence intervals: +13.8, –21.7) ms after the actual attachment times, on average (Table S2), and the detected detachment times occurred 28.6 (+11.9, –19.1) ms before the actual detachment times, on average. On the other hand, when we used the peak-to-peak method to identify transitions, we found that the detected attachment times occurred 55.5 (+195.5, –69.0) ms before the actual attachment times, on average, and the detected detachment times occurred 50.4 (+188.1, –64.9) ms after the actual detachment times, on average. Taken together, the single-threshold method has better temporal resolution when identifying transitions between the bound and unbound states.

When binding interactions are aligned based on the covariance thresholds, it is assumed that the covariance drops and rises in conjunction with transitions between the bound and unbound states. With the single-threshold method, this is a fairly reasonable assumption that explains why it outperforms the peak-to-peak method. Each true transition point separates more highly correlated bead motion (i.e., the unbound state) from less highly correlated bead motion (i.e., the bound state). The covariance is calculated over a window so that when the covariance window is centered at a transition point, the window will include equal amounts of more highly and less highly correlated data. The covariance at the transition point should then lie at some intermediate value between the two peaks of the covariance histogram. However, the single-threshold method is not perfect at locating the transition points. First, although the value of the covariance at a transition point will likely be near the minimal value between the two peaks of the covariance histogram, there is no guarantee that it will lie exactly at this minimal value. Additionally, synchronized large-scale movement of both beads due to the myosin's power stroke can produce transient spikes in the covariance value during transitions, and these spikes can potentially decrease the accuracy of the single-threshold method in identifying exact transition times.

The peak-to-peak method produced poorer alignment than the single-threshold method. When the peak-to-peak

method is used to identify transitions, it is assumed that transitions occur when the covariance crosses the upper threshold, defined by the position of the unbound peak. This is inherently less accurate for estimating transition points than the single-threshold method. A window of data that has a covariance value that is similar to the value of the unbound peak contains primarily correlated data, and therefore it is unlikely that the center of this window is near the actual transition point. In fact, the calculated transition point using the peak-to-peak method would be expected to deviate from the actual transition point by at least half the window size.

Taken together, our data show that when binding interactions are synchronized using a single covariance threshold, the resulting ensemble averages are expected to have better alignment of binding interactions. However, as noted previously, the use of a single covariance threshold to detect binding interactions is more susceptible to false-positive binding interactions, which would lead to an underestimation of the true substep sizes. The peak-to-peak method is better for binding interaction detection without including false positives, but it lacks the necessary temporal resolution to accurately align the detected interactions.

Change-point algorithm for aligning interactions

Rather than relying on the covariance when estimating transition times, we tested the use of separate methods for detecting and synchronizing binding interactions. To improve our ability to locate the transition times of each binding interaction, we implemented a change-point algorithm (see [Materials and Methods](#) for details). Change-point algorithms have been used in step finding for transitions in biological processes in which the algorithm identifies the most likely times when there was a change in a parameter such as motor position or rotation of the myosin lever arm (26,27). We have adapted the change-point algorithm for the three-bead assay, in which we search for the most likely transition times based on changes in both the mean and the variance of the bead positions because both of these parameters differ between the bound and unbound states (Fig. 3 A). For each binding interaction identified by the covariance threshold method (Fig. 3 B), our algorithm examines the positions of the trapped beads in a window surrounding that interaction and finds the two points (i.e., binding initiation and detachment) within this window that most likely represent transitions in the mean and variance of the data (Fig. 3 C; see [Materials and Methods](#) for details).

To test the ability of the change-point algorithm to accurately identify transition times, we again analyzed the same 10 sets of simulated data described above (sets 1–10). We found that the attachment times detected by the change-point algorithm occurred 0.5 (+9.0, –5.5) ms after the actual attachment times, on average (Table S2), and the detachment times detected by the change-point algorithm




FIGURE 3 The change-point algorithm more precisely identifies transitions between bound and unbound states. (A) Simulated optical trapping data show the average position between the optically trapped beads over time during a binding interaction. Data obtained during the bound state (light purple) are drawn from a normal distribution with a shifted mean and a lower variance when compared with data obtained during the unbound state (black). The change-point algorithm seeks to find the time points that best separate the two distributions. The locations of the actual simulated transitions are marked with blue vertical lines. (B) The calculated covariance of the bead positions during the simulated binding interaction in (A) is shown. The attachment and detachment times identified by the single-threshold (green) and the peak-to-peak (red) methods are shown with dashed vertical lines. The actual transitions are marked with solid blue vertical lines. (C) The change-point algorithm determines the likelihood that any two points within an extended search window are the two transition points. (left) A plot of the likelihood assigned to each pair of points within the search window, viewed from the side (see [Materials and Methods](#) for details), is shown. The change points, which occur when the likelihood is maximized, are shown with dashed yellow vertical lines, and the actual transitions are marked with solid blue vertical lines. (right) The likelihood viewed from above is shown. Regions of yellow correspond to higher likelihood, and regions of dark blue correspond to lower likelihood. The two change points are marked with solid black lines. To see this figure in color, go online.

occurred 0.7 (+4.8, –4.2) ms after the actual detachment times, on average (Table S2). Statistical testing demonstrates that the change-point algorithm outperforms both the single-threshold method ($p_{\text{start}} < 0.001$, $p_{\text{end}} < 0.001$) and the peak-to-peak method ($p_{\text{start}} < 0.001$, $p_{\text{end}} < 0.001$) in identifying transition times. As our simulated data were generated with a sampling frequency of 2 kHz, these average errors of ~ 0.5 ms indicate that the change-point algorithm was typically correct within one point. It is possible that a higher sampling frequency would further increase the accuracy.

To explore the ability of these three methods to accurately identify transition points, we generated cumulative distributions of the differences between the detected transition times and the actual simulated transition times for both the initiation and termination of the binding interactions (Fig. 4). Here, the width of the distribution reveals the precision of the corresponding method, whereas the sign and magnitude of the average error reveal the systematic bias of that method.




FIGURE 4 The change-point algorithm minimizes the error when detecting the locations of transitions. The error was calculated as the difference between the detected binding times and the actual simulated binding times for simulated data (sets 1–10). (left) Cumulative distributions of the errors in determining the binding initiation times using the peak-to-peak method (red), the single-threshold method (green), and the change-point algorithm (yellow) are shown. Blue shows the actual transition. Statistical comparisons can be found in Table S2. (right) Cumulative distributions of the errors when determining the binding termination times using the peak-to-peak method (red), the single-threshold method (green), and the change-point algorithm (yellow) are shown. Blue shows the actual transition. To see this figure in color, go online.

As expected, the cumulative distributions of errors generated from the peak-to-peak method are wide, indicating low precision at identifying the transitions, whereas the distributions generated from the single-threshold method are narrower, indicating higher precision. The distributions generated from the change-point algorithm are very narrow, and the mean error is close to 0. This indicates that the change-point algorithm is very precise and has lower systematic bias than either the single-threshold or peak-to-peak method.

Comparison of ensemble averages generated using different methods

To test our predictions about the relative accuracy of the ensemble averages when using each method of analysis, we generated ensemble averages from the 10 sets of simulated data studied previously (sets 1–10). First, we generated ensemble averages using the actual locations of all 1000 simulated binding interactions to align the binding interactions (Fig. 5, A and B, real). We also generated ensemble averages for each of the 10 sets of data using the actual locations of the 100 simulated binding interactions within each set. Exponential curves were fit to each of these averages to estimate the substep sizes and rates of the simulated myosin working stroke (Fig. 5, C–F, real; Table S3). The magnitude of substep 1 estimated from the time-forward averages was 4.7 (+0.4, –0.4) nm, on average, whereas the magnitude of the total step estimated from the time-forward averages was 6.4 (+0.2, –0.2) nm, on average. The magnitude of substep 1 estimated from the time-reversed averages was 5.7 (+0.2, –0.3) nm, on average, whereas the magnitude of the total step estimated from the time-reversed averages was 6.5 (+0.1, –0.2) nm, on average. The estimated rate of transitioning from the first substep

to the second substep (k_f) was 68.7 (+15.8, –20.9) s^{-1} , and the estimated rate of transitioning from the second substep to the detached state (k_r) was 4.3 (+2.2, –1.9) s^{-1} .

We then used either the single-threshold method or the peak-to-peak method to detect binding interactions within each data set. When the single-threshold method was used to detect binding interactions, we applied a filter to ignore any detected interactions that were shorter than 77 ms or within 63 ms of another detected interaction to avoid including false-positive interactions (Figs. S2 and S4 show the effect of including these false-positive binding interactions). These parameters were selected to optimize the analysis of the simulated data with its signal/noise ratio; however, the optimal values for these parameters will vary depending on the signal/noise ratio of the experimental data (Fig. S3), and the program enables the user to adjust these values to suit their data. That being said, techniques have been developed to accurately determine attachment lifetimes from data with pronounced experimental dead times (28).

To identify transitions between the bound and unbound states for each interaction, we either included or omitted the change-point algorithm. For each of these analysis methods, we used the binding interactions and transitions detected over all 10 data sets to generate ensemble averages (Fig. 5, A and B). As before, we also generated ensemble averages from the binding interactions detected within each of the 10 sets of data, and exponential curves were fit to each average to estimate the substep sizes and rates of the simulated myosin working stroke (Fig. 5, C–F; Table S3). As expected, using the change-point algorithm to align the binding interactions resulted in the most accurate estimates.

When the peak-to-peak method was used to both detect and align the binding interactions, the ensemble averages were misshapen (Fig. 5, peak-to-peak). The time-forward average, for example, includes the characteristic increase in displacement but then drops. This drop is due to the fact that the binding interaction termination times detected by the peak-to-peak method often came after the actual termination times, leading to the inclusion of baseline data at the end of the time-forward average. Total step size estimates were generated by extrapolation of the exponential fits. The time-forward average also appears to start too late, as the peak-to-peak method typically guesses that binding initiation times occur before they actually do (Fig. 4). Exponential curves were very poorly fit to these ensemble averages.

When the single-threshold method was used to both detect and align the binding interactions, the ensemble averages had better overall shape (Fig. 5, single-threshold). However, similar to the averages generated with the peak-to-peak method, misalignment among the individual trajectories resulted in very gradual transitions between the bound and unbound states. The time-forward average, for example, appears to start too early, as the single-threshold method typically guesses that binding initiation times occur after they actually do (Fig. 4).

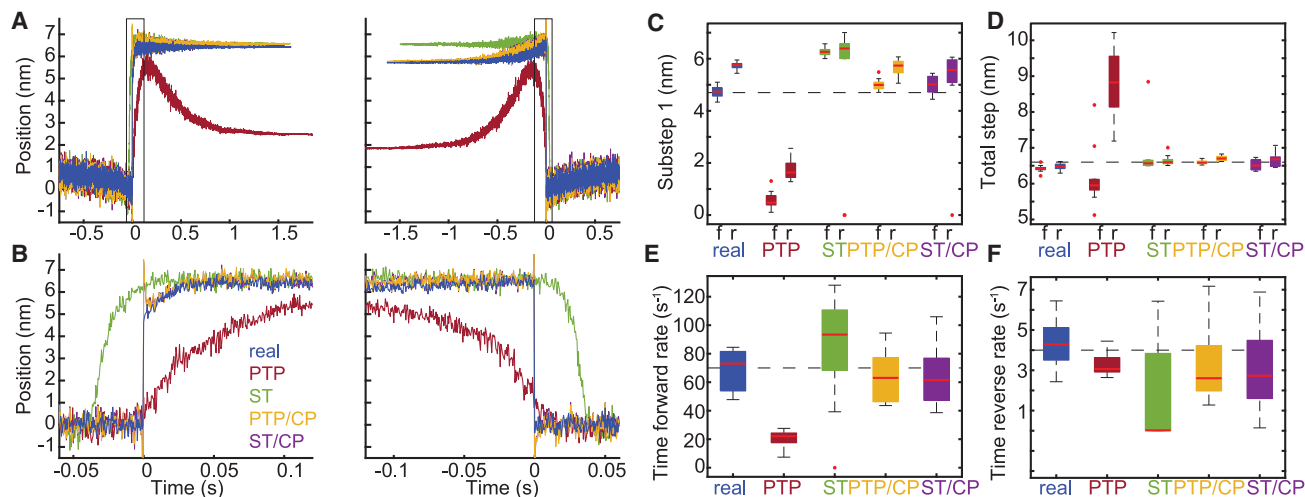


FIGURE 5 Ensemble averages of simulated binding interactions. In total, 10 sets of data were simulated, each containing 100 binding interactions (sets 1–10). Interactions were detected using either the peak-to-peak (PTP) or the single-threshold (ST) method, and interactions were aligned using either the transitions estimated by the covariance threshold method or the change points identified by the change-point algorithm (CP). (A) (*left*) For each analysis method, all detected binding interactions were aligned at the estimated initiation times and averaged together to generate time-forward ensemble averages. (*right*) For each analysis method, all detected binding interactions were aligned at the estimated termination times and averaged together to generate time-reversed ensemble averages. Also shown are the time-forward and time-reversed ensemble averages generated from the known locations of the actual simulated binding interactions (*blue, real*). (B) A zoomed-in view of the boxed segments of the ensemble averages in (A) highlights the misalignment in the averages when the change-point algorithm is omitted. (C–F) For each of the 10 simulated sets of data containing 100 binding interactions, ensemble averages were generated and fit with single exponential functions. The substep sizes and rates of the simulated myosin working stroke were estimated from the exponential fits. Box plots show the estimated parameters for each analysis method. Outliers are indicated by red dots. The substep sizes were estimated from both the time-forward (f) and the time-reversed (r) ensemble averages. Horizontal dashed lines show the values of the simulated parameters. Statistical analysis for each parameter can be found in Table S3. To see this figure in color, go online.

When the change-point algorithm was used to align the binding interactions, the ensemble averages featured much sharper transitions (Fig. 5, peak-to-peak/change-point algorithm and single-threshold/change-point algorithm). However, very sharp spikes in displacement occur at the transition times (Fig. 5, A and B, peak-to-peak/change-point algorithm and single-threshold/change-point algorithm). Brownian-motion-driven fluctuations in the bead positions can cause changes in the data from one point to the next that are not due to transitions between the bound and unbound states. If such noise happens to occur near a real transition point, it offers an attractive candidate for the change point, and the change-point algorithm may choose that point instead of the less pronounced, yet correct, transition time. However, we have shown that the transition times estimated by the change-point algorithm are within one to two points of the actual simulated transition times, on average (Fig. 4; Table S2), and the resulting ensemble averages are very accurate. Appropriate fits can be obtained by omitting these spikes from the fitted data.

The time-reversed ensemble average generated from the actual locations of the simulated binding interactions led to an overestimate of the magnitude of substep 1 (Fig. 5, B and C; Table S3). To generate the time-reversed ensemble average, short-lived binding interactions are extended in time to match the duration of the longest-lived binding interaction, and the value of this extension equals the average po-

sition of the beads during the first 5 ms of the binding interaction. The rate of transitioning from the first substep to the second substep in our simulated data was 70 s^{-1} , matching the rate of ADP release for beta cardiac myosin (24). Because of this fast rate, a large number of transitions to the second substep occur before the 5 ms used to generate the extensions, leading to inaccurate extension values. The proportion of binding interactions that is expected to transition to the second substep within the first 5 ms is given by the integral of the probability density function:

$$\text{proportion of substeps missed} = \int_0^{0.005} ke^{-kt} dt.$$

For a rate of 70 s^{-1} , this proportion is equal to $\sim 30\%$, and this will lead to an overestimation of the size of the first substep. A possible fix is to shorten the 5-ms window used for calculating the extensions, but it then becomes crucial that the binding initiation times are determined with high accuracy. Neither the single-threshold method nor the peak-to-peak method has sufficient resolution to accurately determine the exact initiation times (Fig. 4). Even the change-point algorithm, which we have shown to have an average error of $\sim 0.5 \text{ ms}$, would be insufficient for generating the time-reversed ensemble averages of interactions with very fast kinetics. It is possible that this could be

improved with faster data sampling. In the case of transitions with slower kinetics, this problem is easily avoided. When we simulated 1000 binding interactions using much slower rates (k_f of 5 s^{-1} and k_r of 2 s^{-1} , sets 11–20), we were able to generate time-forward and time-reversed ensemble averages with accurate step sizes using multiple methods (Fig. S5).

Performance of the computational tool to analyze experimental data

To test the ability of the computational tool on real experimental data, we conducted optical trapping experiments using ventricular myosin at $1 \mu\text{M}$ ATP (Fig. 6). We intentionally collected a small data set consisting of 66 binding interactions from five molecules. Binding interactions were identified using the peak-to-peak method, and transition points were identified using the change-point algorithm. The SPASM computational tool was used to generate cumulative distributions of individual binding interactions (Fig. 6B). The cumulative distributions of the attachment durations are well fit by a single exponential function. This exponential rate gives the rate of actomyosin detachment, and it has a value of 4.7 s^{-1} , which is




FIGURE 6 Ensemble averages of experimental optical trapping data. The kinetics and mechanics of cardiac myosin in $1 \mu\text{M}$ ATP were measured using the three-bead assay. (A) Experimental data trace shows the displacement (D) and covariance (C). (B) Cumulative distributions for the (left) binding interaction durations and (right) total working stroke displacements are shown. The peak-to-peak method was used to detect binding interactions. Red lines show the cumulative fits based on (left) exponential and (right) normal distributions. The characteristic rate obtained from the fit to the distribution of attachment durations gives a detachment rate equal to 4.7 s^{-1} , which is consistent with the expected rate of ATP-induced actomyosin dissociation at $1 \mu\text{M}$ ATP. The distribution of total step sizes has a mean of 6.3 nm and a standard deviation of 9.2 nm . (C) The change-point algorithm was used to align the interactions identified using the peak-to-peak method. A total of 66 binding interactions from five molecules were analyzed. The resulting ensemble averages have estimated substep sizes of 4.4 and 2.0 nm . The estimated time-forward rate is 74 s^{-1} , and the estimated time-reversed rate is 3.2 s^{-1} . These values are consistent with previous measurements using a much larger data set, and they agree well with the previously measured rates of ADP release and ATP-induced dissociation $1 \mu\text{M}$ ATP. To see this figure in color, go online.

consistent with the expected rate of ATP-induced actomyosin dissociation at $1 \mu\text{M}$ ATP (24). The cumulative distribution of total working stroke displacements is well fit by a single normal distribution (indicating likely single-molecule conditions), with a mean displacement of 6.3 nm and a standard deviation of 9.2 nm . This is consistent with previous measurements of the cardiac myosin working stroke (8,25). Ensemble averages (Fig. 6C) reveal that, consistent with previous measurements (8,25), ventricular cardiac myosin has a two-substep working stroke with a first substep of 4.4 nm and a total displacement of 6.4 nm . The time-forward averages have a rate of 74 s^{-1} , which is consistent with the rate of ADP release, and the time-reversed averages have a rate of 3.2 s^{-1} , which is consistent with the rate of ATP-induced actomyosin dissociation at $1 \mu\text{M}$ ATP (24). Taken together, our computational tool can generate accurate ensemble averages with sharp transitions from a relatively small set of experimental data.

Broader applicability of the approach

The methods presented in this work were applied to study actomyosin. As noted previously, the three-bead assay has been used to explore many different single-molecule systems, including dynein, the lac repressor, and kinesins. Moreover, the general ideas behind our computational tool are broadly applicable to any set of data, not just optical trapping data, containing well-defined populations that can be distinguished through some aspect of the data. One such possibility is data obtained from single-molecule fluorescence resonance energy transfer (FRET) experiments. In the Supporting Material, we describe how to adapt the change-point algorithm to systems in which the desired change points occur in data with different underlying distributions.

Limitations

There are a number of limitations accompanying our computational tool and the methods we use to analyze our data. Although the covariance between the position of each trapped bead in the three-bead assay is very helpful for locating binding interactions under many circumstances, it does have drawbacks. The covariance is calculated over a window, and therefore it does not always drop enough during short-lived binding interactions to register as a genuine binding interaction. Furthermore, depending on the quality of the data, it may be difficult or even impossible to obtain a covariance histogram with two distinct populations. This could stem from system compliances. One benefit of the peak-to-peak method is that the covariance histogram populations do not need to be completely separated to avoid false-positive binding interactions, but a certain degree of separation is needed to

make the covariance useful. Additionally, analysis is dependent on many parameters, including the window sizes used to calculate and smooth the covariance, and it can be difficult to choose appropriate values for these parameters for a given experimental system. The computational tool includes features that allow the user to correct for these drawbacks when they are encountered. Finally, as evidenced by the ensemble averages generated from our simulated data (Fig. 5), ensemble averaging has limitations for estimating the rates and substep sizes for transitions with very fast kinetics.

Summary

Here, we developed a computational tool, SPASM, for the detection and alignment of single-molecule binding interactions and for the generation of ensemble averages that can reveal characteristics about the data that are often obscured by noise. We show that it can be advantageous to use separate techniques for the detection and alignment of binding interactions. Specifically, we show that the addition of a change-point algorithm to identify transition times can generate precise ensemble averages with improved alignment. We offer the computational tool, with an intuitive graphical user interface, along with a user guide so that the reader can apply these methods to their own data.

SUPPORTING MATERIAL

Supporting Material can be found online at <https://doi.org/10.1016/j.bpj.2020.10.047>.

AUTHOR CONTRIBUTIONS

T.B. wrote the computational tool, simulated data, and analyzed data. W.T.S. built the optical trap and wrote software for data acquisition. S.R.C. collected optical trapping data. M.J.G. wrote code for the simulator and analyzed data. T.B. and M.J.G. wrote the first draft of the article, and all authors contributed to the final draft.

ACKNOWLEDGMENTS

Funding for this project was provided by National Institutes of Health (R01HL141086 to M.J.G., T32EB018266 to S.R.C.).

REFERENCES

1. Goldman, Y. E., and E. M. Ostap. 2012. 4.1 Introduction. *In Comprehensive Biophysics*. E. H. Egelman, ed. Elsevier, p. 1.
2. Spudich, J. A. 2014. Hypertrophic and dilated cardiomyopathy: four decades of basic research on muscle lead to potential therapeutic approaches to these devastating genetic diseases. *Biophys. J.* 106:1236–1249.
3. Greenberg, M. J., H. Shuman, and E. M. Ostap. 2017. Measuring the kinetic and mechanical properties of non-processive myosins using optical tweezers. *Methods Mol. Biol.* 1486:483–509.

4. Finer, J. T., R. M. Simmons, and J. A. Spudich. 1994. Single myosin molecule mechanics: piconewton forces and nanometre steps. *Nature*. 368:113–119.
5. Walter, W. J., M. P. Koonce, ..., W. Steffen. 2012. Two independent switches regulate cytoplasmic dynein's processivity and directionality. *Proc. Natl. Acad. Sci. USA*. 109:5289–5293.
6. Capitanio, M., M. Canepari, ..., F. S. Pavone. 2012. Ultrafast force-clamp spectroscopy of single molecules reveals load dependence of myosin working stroke. *Nat. Methods*. 9:1013–1019.
7. Pyrpassopoulos, S., H. Shuman, and E. M. Ostap. 2020. Modulation of Kinesin's load-bearing capacity by force geometry and the microtubule track. *Biophys. J.* 118:243–253.
8. Greenberg, M. J., H. Shuman, and E. M. Ostap. 2014. Inherent force-dependent properties of β -cardiac myosin contribute to the force-velocity relationship of cardiac muscle. *Biophys. J.* 107:L41–L44.
9. Greenberg, M. J., T. Lin, ..., E. M. Ostap. 2012. Myosin IC generates power over a range of loads via a new tension-sensing mechanism. *Proc. Natl. Acad. Sci. USA*. 109:E2433–E2440.
10. Laakso, J. M., J. H. Lewis, ..., E. M. Ostap. 2008. Myosin I can act as a molecular force sensor. *Science*. 321:133–136.
11. Guilford, W. H., D. E. Dupuis, ..., D. M. Warshaw. 1997. Smooth muscle and skeletal muscle myosins produce similar unitary forces and displacements in the laser trap. *Biophys. J.* 72:1006–1021.
12. Veigel, C., F. Wang, ..., J. E. Molloy. 2002. The gated gait of the processive molecular motor, myosin V. *Nat. Cell Biol.* 4:59–65.
13. Veigel, C., L. M. Coluccio, ..., J. E. Molloy. 1999. The motor protein myosin-I produces its working stroke in two steps. *Nature*. 398:530–533.
14. Veigel, C., M. L. Bartoo, ..., J. E. Molloy. 1998. The stiffness of rabbit skeletal actomyosin cross-bridges determined with an optical tweezers transducer. *Biophys. J.* 75:1424–1438.
15. Chen, C., M. J. Greenberg, ..., H. Shuman. 2012. Kinetic schemes for post-synchronized single molecule dynamics. *Biophys. J.* 102:L23–L25.
16. Capitanio, M., M. Canepari, ..., R. Bottinelli. 2006. Two independent mechanical events in the interaction cycle of skeletal muscle myosin with actin. *Proc. Natl. Acad. Sci. USA*. 103:87–92.
17. Barrick, S. K., S. R. Clippinger, ..., M. J. Greenberg. 2019. Computational tool to study perturbations in muscle regulation and its application to heart disease. *Biophys. J.* 116:2246–2252.
18. Clippinger, S. R., P. E. Cloonan, ..., M. J. Greenberg. 2019. Disrupted mechanobiology links the molecular and cellular phenotypes in familial dilated cardiomyopathy. *Proc. Natl. Acad. Sci. USA*. 116:17831–17840.
19. Lewis, J. H., M. J. Greenberg, ..., E. M. Ostap. 2012. Calcium regulation of myosin-I tension sensing. *Biophys. J.* 102:2799–2807.
20. Takagi, Y., R. E. Farrow, ..., J. E. Molloy. 2014. Myosin-10 produces its power-stroke in two phases and moves processively along a single actin filament under low load. *Proc. Natl. Acad. Sci. USA*. 111:E1833–E1842.
21. Molloy, J. E., J. E. Burns, ..., D. C. White. 1995. Movement and force produced by a single myosin head. *Nature*. 378:209–212.
22. Mehta, A. D., J. T. Finer, and J. A. Spudich. 1997. Detection of single-molecule interactions using correlated thermal diffusion. *Proc. Natl. Acad. Sci. USA*. 94:7927–7931.
23. Smith, D. A., W. Steffen, ..., J. Sleep. 2001. Hidden-Markov methods for the analysis of single-molecule actomyosin displacement data: the variance-Hidden-Markov method. *Biophys. J.* 81:2795–2816.
24. Deacon, J. C., M. J. Bloemink, ..., L. A. Leinwand. 2012. Identification of functional differences between recombinant human α and β cardiac myosin motors. *Cell. Mol. Life Sci.* 69:2261–2277.
25. Woody, M. S., M. J. Greenberg, ..., E. M. Ostap. 2018. Positive cardiac inotrope omecamtiv mecarbil activates muscle despite suppressing the myosin working stroke. *Nat. Commun.* 9:3838.

26. Kerssemakers, J. W., E. L. Munteanu, ..., M. Dogterom. 2006. Assembly dynamics of microtubules at molecular resolution. *Nature*. 442:709–712.
27. Beausang, J. F., Y. E. Goldman, and P. C. Nelson. 2011. Change-point analysis for single-molecule polarized total internal reflection fluorescence microscopy experiments. *Methods Enzymol.* 487:431–463.
28. Woody, M. S., J. H. Lewis, ..., E. M. Ostap. 2016. MEMLET: an easy-to-use tool for data fitting and model comparison using maximum-likelihood estimation. *Biophys. J.* 111:273–282.