



Rapid discrimination among individual DNA hairpin molecules at single-nucleotide resolution using an ion channel

Wenonah Vercoutere¹, Stephen Winters-Hilt², Hugh Olsen¹, David Deamer¹, David Haussler², and Mark Akeson^{1,2*}

RNA and DNA strands produce ionic current signatures when driven through an α -hemolysin channel by an applied voltage. Here we combine this nanopore detector with a support vector machine (SVM) to analyze DNA hairpin molecules on the millisecond time scale. Measurable properties include duplex stem length, base pair mismatches, and loop length. This nanopore instrument can discriminate between individual DNA hairpins that differ by one base pair or by one nucleotide.

Single DNA molecules can be examined using atomic force microscopy¹, video fluorescence microscopy², and force-measuring laser tweezers³. Ion channels ("nanopores") have also been used to examine individual DNA or RNA strands, leading to the suggestion that an engineered nanopore could be used for sequencing at high rates⁴.

The prototype nanopore we use is formed by an α -hemolysin channel inserted in a lipid bilayer^{5–8} (Fig. 1, bottom panel, A). X-ray diffraction analysis of the assembled protein reveals a heptamer with a 2.6 nm aperture leading into a slightly wider vestibule that abruptly narrows to a transmembrane domain with a 1.5 nm constriction⁹. In 1.0 M KCl (pH 8.0), a 120 mV applied potential produces a steady open channel current (I_0) of 120 ± 5 pA (Fig. 1, top panel, A). Translocation of single-stranded linear DNA reduces this current to $I \approx 14$ pA ($I/I_0 = 12\%$)^{5,7}. Each monomer within single-stranded DNA traverses the length of the 10 nm pore in 1–3 μ s. Here we show that a prototype nanopore instrument can resolve single-nucleotide or single-base pair differences between otherwise identical, unmodified DNA hairpin molecules.

Results and discussion

Use of DNA hairpins to model duplex DNA interaction with the pore. We chose DNA hairpins as model duplexes because they can be formed from short, highly pure oligonucleotides that can be designed to adopt one base-paired structure in 1.0 M salt at room temperature. The initial experiments involved a well-characterized DNA hairpin with a 6 base pair stem and a 4-deoxythymidine (4dT) loop¹⁰. When captured within an α -hemolysin nanopore, this molecule caused a partial current blockade (or "shoulder") lasting hundreds of milliseconds (Fig. 1, top panel, B), followed by a rapid downward spike (Fig. 1, top panel, C). This "shoulder–spike" signature is consistent with two sequential steps: (1) capture of a hairpin stem in the vestibule (Fig. 1, bottom panel, B), where the molecule rattles in place because the duplex stem cannot fit through the 1.5 nm-diameter limiting aperture of the pore, and (2) simultaneous dissociation of the 6 base pairs in the hairpin stem, thus allowing the extended single strand to traverse the channel (Fig. 1, bottom panel, C). This type of signature describes ~60% of blockade events caused by the 6 base pair hairpin. The remaining events varied in amplitude and were less than 1 ms in duration. In the proposed model, these

fast events would be explained by interactions of the hairpin loop with the mouth of the pore without entry into the vestibule.

We tested our explanation of the shoulder–spike signature using a series of blunt-ended DNA hairpins with stems that ranged in length from 3 to 9 base pairs (Table 1). If the model depicted in Figure 1 is accurate, we would expect a substantial increase in blockade shoulder lifetime for each additional base pair and a modest linear increase in the lifetime of the downward spike at the end of the event. We would also expect the shoulder amplitude to decrease as the stem length increased. These predictions proved to be correct. Each base pair addition resulted in a measurable increase in median blockade shoulder lifetime that correlated with the calculated ΔG° of hairpin formation (Fig. 2). Increasing stem length resulted in a 10 μ s increase in median duration of the terminal spike. A downward trend in shoulder current amplitude was also observed from I/I_0 equal to 68% for a 3 base pair stem to I/I_0 equal to 32% for a 9-base pair stem (Table 1). Our results are consistent with greater obstruction of ionic current as the hairpin stem extends farther into the vestibule with each additional base pair. Examination of a 71 base pair DNA duplex confirmed this trend, giving I/I_0 values <30% (data not shown). Control oligonucleotides with the same base compositions as the DNA hairpins, but scrambled, caused blockade events that were on average much shorter than the hairpin events and that did not conform to the shoulder–spike pattern.

The model in Figure 1 also assumes that the hairpin loop is unable to enter the vestibule. Using a molecular dynamics simulation (AMBER field)¹¹, we found that the 4dT loop of these hairpin molecules adopted conformations that would prevent ready entry into the pore vestibule. We tested this using a DNA "dumbbell" with 4dT loops at either end (Table 1). If the loop cannot enter the vestibule, interaction of dumbbell hairpins with the pore would not result in the shoulder–spike signature. When we examined the dumbbell hairpins experimentally, fast blockades (<1 ms) were observed, but shoulder–spike blockades were not.

Single-base pair resolution among individual DNA hairpins.

The strength of the nanopore instrument is discrimination among individual DNA molecules at high speed and high resolution. This is illustrated by an I/I_0 versus duration plot of 3 to 8 base pair hairpin blockade events that were screened manually for adherence to the shoulder–spike signature (Fig. 3A). The identity of a large majority of

¹Department of Chemistry and Biochemistry, and ²Center for Biomolecular Science and Engineering, University of California, Santa Cruz, CA 95064.

*Corresponding author (makeson@chemistry.ucsc.edu).




Figure 1. Blockade of the α -hemolysin nanopore by a DNA hairpin. The upper panel shows a current trace caused by capture and translocation of a 6-bp DNA hairpin through the pore. The lower panel shows a molecular model of these events. (A) Cross section of the α -hemolysin heptamer inserted in a lipid bilayer. A 120 mV applied voltage across the open pore produces an ~ 120 pA ionic current in 1 M KCl at room temperature. (B) Capture of a 6-bp DNA hairpin in the channel causes an abrupt current reduction to an intermediate level ($I/I_0 = 52\%$). Because only linear single-stranded DNA can traverse the 1.5 nm limiting aperture⁷, the stem duplex holds the molecule in the vestibule (760 ms median duration). The 4- deoxythymidines of the hairpin loop (yellow) are shown spanning the pore entrance, and the 6 bp of the stem (white and red) are shown extended into the vestibule. Note the increase in low-frequency noise during hairpin occupancy of the vestibule relative to the open channel. (C) Translocation of the DNA through the limiting aperture of the channel. The partial hairpin blockade ends with a sharp downward spike to ~ 14 pA ($I/I_0 = 12\%$) that lasts about 60 μ s. In our model, this corresponds to simultaneous dissociation of the 6 bp in the hairpin stem, which allows translocation of the extended strand.

DNA hairpin molecules could be visually discerned and the single-base pair difference between individual molecules was easily resolved.

The high level of discrimination visible in Figure 3A was confirmed using an automated procedure for recognition of signal

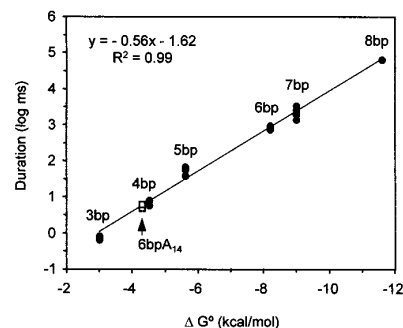


Figure 2. Standard free energy of hairpin formation versus shoulder blockade duration. Standard free energy of hairpin formation (see Table 1) correlated with median duration of hairpin shoulder blockades (solid circles). Each point represents the median blockade duration for a given hairpin length acquired using a separate α -hemolysin pore on a separate day. Median blockade durations and ΔG° for the equivalent of the 6-bp hairpin with a single mismatch (6bpA₁₄, Table 1) are represented by open squares.

regions followed by signal classification using a support vector machine (SVM)^{12,13}. Support vector machines provide a scalable means to represent data in a higher dimensional space where discrimination can be achieved by a hyperplane dividing that space. This provides a uniform method to classify individual blockade signatures acquired by the nanopore instrument. In brief, blockade events from each data file were recognized and cut out using a customized finite-state automaton (FSA)¹⁴. The FSA was less stringent than the shoulder-spike criterion used in Figure 3A, and included any event that exceeded 200 μ s in duration and $I/I_0 < 85\%$. Approximately twice as many signals were passed by the FSA as met the shoulder-spike diagnostic evaluated manually (Fig. 3A), whereas $< 0.1\%$ of the shoulder-spike signals were rejected. Signal features (see Experimental Protocol) were extracted and grouped as a "feature vector" for each event. The feature vectors were used by the SVM to classify each event.

Figure 3B shows an example where blockade events caused by 6 base pair hairpins were classified against blockades caused by 3, 4, 5, 7, and 8 base pair hairpins. The FSA passed 529 of the 6 base pair hairpin events to the SVM and 3,185 of all other events. Because

Table 1. DNA hairpins used in this study

	TT	TT	TT	TT	TT	TT	TT	TT	TT	TT
	T T	T T	T T	T T	T T	T T	T T	T T	T T	T T
Predicted hairpin	G:C	G:C	G:C	G:C	G:C	G:C	G:C	G:C	G:C	G:C
secondary structures	C:G	C:G	C:G	C:G	C:G	C:G	C:G	C:G	C:G	C:G
	5' 3'	A:T	A:T	A:T	A:T	A:T	A:T	A:T	A A	A:T
		C:G	A:T	A:T	A:T	A:T	A:T	A:T	A A	A:T
		5' 3'	G:C	G:C	G:C	G:C	G:C	G:C	G:C	G:C 3'
			5' 3'	C:G	T:A	C:G	C:G	5' 3'	C:G	C:G 5'
				5' 3'	C:G	T:A	T:A	5' 3'	5' 3'	T:A
					5' 3'	G:C	T:A			T:A
						5' 3'	C:G			G:C
							5' 3'			C:G
										T T
										TT
Identity	3bp	4bp	5bp	6bp	7bp	8bp	9bp	5bp3dT	6bpA ₁₄	Dumbbell
ΔG° (kcal/mol)^a	-3.0	-4.5	-5.6	-8.2	-9.0	-11.4	-12.8	-4.2	-4.3	-11.3
I/I_0 (%)^b	68	64	60	52	47	35	32	62	53	NA

^a ΔG° values for hairpin formation were calculated using the DNA mfold server (<http://bioinfo.math.rpi.edu/~folder/dna/form1.cgi>) based on data from SantaLucia¹⁷. Assume 22°C and 1 M KCl.

^b I is the current average for an event shoulder (in pA). I_0 is the current average for the open channel (in pA).




Figure 3. Discrimination between DNA hairpins at single-base pair resolution. (A) Event diagram for DNA hairpins with 3- to 8-bp stems. Events were selected for adherence to the shoulder–spike signature. Each point represents the duration and amplitude of a shoulder blockade caused by one DNA hairpin captured in the pore vestibule. The data for each hairpin are from at least two different experiments run on different days. Median I/I_0 values for each type of hairpin varied by at most 2%. (B) Classification of the 6-bp hairpin (solid bars) versus all other hairpins (open bars) by SVM. Note the log scale on the y-axis. The dashed lines mark the limits of the rejection region. The boundaries of the rejection region were determined by independent data, not post hoc, on the data shown.

selectivity was relaxed at the FSA, there were many ambiguous signals with scores near zero. Using an additional set of independent data, the SVM can be trained to exclude these by introducing a rejection region for the low scores (the region between dashed lines in Fig. 3B). The events that were rejected were primarily fast blockades similar to those caused by loops on the dumbbell hairpin (Table 1) or acquisition errors caused by the low selectivity threshold of the FSA. When 20% of the events were rejected in this manner, the SVM scores for the 6 base pair hairpin discrimination achieved a sensitivity of 98.8% and a specificity of 98.8% (see Experimental Protocol). Similar results were obtained for each class of hairpins depicted in Figure 3A. Overall the SVM achieved an average sensitivity of 98% and average specificity of 99%. Thus, the stem length of an individual DNA hairpin can be determined at single-base pair resolution using a machine learning algorithm.

Detection of single-nucleotide differences. We found that single-nucleotide differences between otherwise identical DNA hairpins could be detected using the nanopore instrument. Here we present two examples. The first example involved the hairpin loop. A 5 base pair hairpin with a 3-deoxythymidine loop (5bp3dT in Table 1) caused pore blockades in which the shoulder amplitude was increased ~2 pA and the median shoulder duration (21 ms) was reduced threefold relative to the same hairpin stem with a 4dT loop (5bp in Table 1). Typical events are illustrated in Figure 4A. The FSA acquired 3,500 possible 5 base pair hairpin signals from 10 min of recorded data. The SVM classification for this data set (Fig. 4B) gave sensitivity and specificity values of 99.9% when 788 events were rejected as the unknown class. The second example involved the stem of a 6 base pair hairpin (T₁₄→A₁₄; 6bpA₁₄ in Table 1) caused an




Figure 4. Detection of single-nucleotide differences between DNA hairpins. (A) Comparison of typical current blockade signatures for a 5-bp hairpin and a 5-bp hairpin with a 3-dT loop. The standard 5-bp hairpin event has a 2% deeper blockade than the 5bp3dT hairpin. (B) Histogram of SVM scores for 5-bp hairpins (solid bars) versus 5-bp hairpins with 3-dT loops (open bars). (C) Comparison of typical current blockade signatures for a standard 6-bp hairpin and a 6-bp hairpin with a single dA₃-dA₁₄ mismatch in the stem. The 6bpA₁₄ event is expanded to show the fast downward spikes. These rapid, near-full blockades and the much shorter shoulder durations are the main characteristics identified and used by SVM to distinguish 6bpA₁₄ hairpin events from 6-bp hairpin events. (D) Histogram of SVM scores for 6-bp hairpins (solid bars) versus 6bpA₁₄ hairpins (open bars).

approximately 100-fold decrease in the median blockade shoulder duration relative to a hairpin with a perfectly matched stem (6bp in Table 1). Typical events are shown in Figure 4C. This difference in duration is consistent with the effect of a mismatch on ΔG° of hairpin formation (Fig. 2), and it permitted a 90% separation of the two populations using the manually applied shoulder–spike diagnostic. When analysis was automated, the FSA acquired 1,031 possible events from 10 min of recorded data (Fig. 4D). With the aid of wavelet features¹⁵ that characterize the low-frequency noise within the shoulder current, the SVM was able to discriminate the standard 6 base pair hairpin from the mismatched 6bpA₁₄ hairpin with sensitivity 97.6% and specificity 99.9% while rejecting only 42 events.

Voltage-pulse routine. For longer hairpin stems (or for native duplex DNA), very long shoulder blockades preclude rapid identification of each captured molecule. For example, the shoulder duration for a hairpin with as few as 8 base pairs ranged up to 300 s, resulting in a very small number of measurable events in a 30 min experiment (Fig. 3A). To overcome this limitation, we modified the acquisition protocol from a fixed +120 mV potential to a voltage pulse routine that alternated from +120 mV for 249.5 ms to -40 mV for 0.1 ms. In essence, the routine was designed to capture and examine each hairpin stem for a finite amount of time under standard conditions then eject the hairpin rather than pulling it through the pore. Representative blockades for 7, 8, and 9 base pair stems using this acquisition protocol are shown in Figure 5. Shoulder blockades caused by the 8 base pair and 9 base pair hairpins toggled between two conductance states. The greater of these states corresponded to the average conductance for the 7 base pair hairpin. The lesser conductance states for the 8 base pair and 9 base pair hairpins were nearly equal with one another; however, transitions between the two states were significantly more frequent for the 8 base pair hairpin than for the 9 base pair hairpin. These two conductance states may represent transient interaction of the terminal base pair of the 8 and 9 base pair hairpins with amino residues in the vestibule wall near the limiting aperture. Alternatively, toggling between these two conductance states may be caused by transitions between two hairpin




Figure 5. Typical current blockade signatures caused by 7, 8, and 9 bp hairpins obtained using a voltage pulse routine. The top trace represents the voltage waveform applied across a single α -hemolysin channel. The bottom trace represents ionic current through the channel in response to this voltage during a single experiment sampling a mixture of 7, 8, and 9 bp hairpins. Each current sweep begins with a capacitance transient followed by a steady current of 122 pA through the open channel. Capture of a hairpin in the pore vestibule (arrows) results in a partial blockade. This ends when the voltage briefly reverses to -40 mV, releasing the hairpin. The blockade events shown for each hairpin length are representative of thousands of events acquired using a single α -hemolysin pore prepared separately on at least three occasions.

conformations that are dependent upon the length of the duplex stem. Toggling was also observed in experiments with 71 base pair duplex DNA, indicating that this may be a common characteristic in longer duplex DNA molecules.

When analysis of this data set was automated, signals for individual 7, 8, and 9 base pair hairpin molecules were distinguishable from one another in a three-way mixture with average sensitivity 99% and average specificity 96%. In ongoing work we are using higher order wavelets and hidden Markov models to boost the SVM discrimination¹⁶.

Conclusions. We conclude that a prototype nanopore detector coupled with machine learning algorithms can resolve single-nucleotide or single-base pair differences between otherwise identical DNA hairpin molecules. Unlike other single-DNA molecule assays, this nanopore instrument examines DNA molecules in solution without chemical modification, amplification, or adsorption to a solid surface. Thousands of molecules can be examined, classified, and quantified in minutes.

Single-base resolution among DNA hairpins by the prototype nanopore could be put to direct use in the measurement of sequence-specific DNA damage caused by radiation, or in the measurement of duplex stability changes caused by nucleotide modifications (Fig. 2). An even greater impact will likely come in design and testing of the next generation of nanopore devices. For example, a number of fabricated nanopores are under development^{4,8}; however, little is known about the heterogeneity of these structures at angstrom resolution or about their noncovalent interactions with DNA. The DNA hairpins we have examined with the prototype nanopore constitute precise molecular gauges that could be used to measure each of those properties. The next generation of nanopore devices will also require scalable computational methods that provide discrimination among ionic current signatures in real time. Implementation of this machine learning software will initially require data sets for FSA and SVM training that contain distinct, reproducible signals for which the identities can be independently verified. The data sets we introduce in this paper fulfill these requirements, establishing DNA hairpin analysis with the α -hemolysin pore as a tool for engineering at the nanoscale.

Experimental protocol

DNA hairpin design, synthesis, and purification. DNA oligonucleotides were synthesized using an ABI 392 Synthesizer (ABI, Foster City, CA), purified by polyacrylamide gel electrophoresis (PAGE), and stored at -70°C in Tris-EDTA buffer. The length of the DNA hairpin stems was varied by adding or removing base pairs from a well-characterized 6-bp hairpin¹⁰. Unless otherwise noted, the hairpin loops were composed of 4dT, and both ends of

the hairpin stems were closed with G:C or C:G base pairs. The prediction that each hairpin would adopt one base-paired structure was tested and confirmed using the DNA mfold server (<http://bioinfo.math.rpi.edu/~folder/dna/form1.cgi>), which is based in part on data from SantaLucia¹⁷. Control linear DNA strands had the same base compositions as the hairpins, but the primary sequences were scrambled so that stable duplex stem regions could not form.

DNA duplex design, synthesis, and purification. A 71-bp DNA duplex was amplified by PCR from a 1.8 kb *Xenopus* elongation factor template (Ambion, Austin, TX) using PCR SuperMix (Life Technologies, Gaithersburg, MD), and then purified by PAGE. The final concentration of this 71-bp duplex when added to the nanopore was $0.05\ \mu\text{M}$.

Formation of α -hemolysin pores in horizontal bilayers. Each experiment was conducted using one α -hemolysin channel inserted into a diphytanoyl-phosphatidylcholine/hexadecane bilayer across a $20\ \mu\text{m}$ diameter horizontal Teflon aperture, as described^{5,7}. Chambers of $70\ \mu\text{l}$ volume on either side of the bilayer contained $1.0\ \text{M}$ KCl buffered at pH 8.0 ($10\ \text{mM}$ HEPES/KOH). Voltage was applied across the bilayer between Ag-AgCl electrodes that were replaced each day. DNA was added to the *cis* chamber, giving a $10\ \mu\text{M}$ final concentration unless otherwise noted. All experiments were conducted at room temperature ($22 \pm 1^{\circ}\text{C}$).

Data acquisition and analysis. Ionic current was filtered at 100 kHz bandwidth using a lowpass Bessel filter and recorded at $5\ \mu\text{s}$ intervals (200 kHz) using an Axopatch 200B amplifier (Axon Instruments, Foster City, CA) coupled to an Axon Digidata 1320A digitizer. Applied potential was 120 mV (*trans* side positive) unless otherwise noted. Manual analysis of shoulder blockade current and duration was performed using Fetchan 6.0.6.01 on data filtered at 10 kHz using a digital lowpass Gaussian filter (Axon Instruments). Ionic current blockade levels are reported as percentage of open channel current, I/I_0 , where I is the average event current and I_0 is the average open channel current. In unusual cases ($<10\%$ of channels), the toggle rate and amplitude difference between conductance states caused by 8- and 9-bp hairpins were measurably different from the representative events shown in Figure 5. Data from those channels were not included in the SVM analysis comparing 7-, 8-, and 9-bp hairpin blockade events.

Molecular modeling. The molecular model of α -hemolysin combined with the 6-bp hairpin shown in Figure 1 was prepared using the crystal structure of α -hemolysin⁹ and Chem-Site¹¹.

Support vector machine learning. Blockade events from each data file were recognized and cut out using a customized FSA, and then passed to the SVM. Classification by SVM requires a training phase and a test phase. In its training phase, an SVM learns to distinguish one class of blockade events (the query class) from all other blockade events using separate data files. SVM scores cover a range (from +3 to -3 in our experiments), with positive scores given to events that fit the query class criteria and negative scores given to events that do not fit the query class criteria. Scores near zero are ambiguous. After training, the SVM is tested on independent mixed data. In the analysis of the DNA molecule signals, the vectors fed to the SVM were based on three criteria: (1) blockade shoulder duration, average, minimum, maximum, and standard deviation of I/I_0 for the entire blockade shoulder; (2) minimum, maximum, and standard deviation of I/I_0 over each of 10 time-domain bins; and (3) a wavelet profile based on averaging the fifth-order wavelet coefficients derived from the signal for the whole event. The statistical measures for SVM are sensitivity and specificity. Sensitivity is defined as true positives/(true positives + false negatives), and specificity is defined as true positives/(true positives + false positives). A true positive is an event in the test data that comes from the positive class and is assigned a positive value; a false positive occurs when the SVM assigns a positive score to an event in the test data when that event actually comes from the negative class. A false negative is an event that is assigned a negative value, but actually comes from the positive class. These values provide a measure of how well the SVM was able to discriminate among classes of data.

Acknowledgments

This work was supported by NRHGI grant HG01826-01. We wish to thank Laura Steinmann for thoughtful comments on the manuscript.

Received 2 August 2000; accepted 1 December 2000



1. Rief, M. Clausen-Schaumann, H. & Gaub, H.E. Sequence-dependent mechanics of single DNA molecules. *Nat. Struct. Biol.* **6**, 346–349 (1999).
2. Perkins, T.T., Quake, S.R., Smith, D.E. & Chu, S. Relaxation of a single DNA molecule observed by optical microscopy. *Science* **264**, 822–826 (1994).
3. Smith, S.B., Cui, Y. & Bustamante, C. Overstretching B-DNA: the elastic response of individual double-stranded and single-stranded DNA molecules. *Science* **271**, 795–799 (1996).
4. Voss, D. Gene express. *New Sci.* **164**, 40–43 (1999).
5. Akeson, M., Branton, D., Kasianowicz, J.J., Brandin, E. & Deamer, D.W. Microsecond time-scale discrimination among polycytidylic acid, polyadenylic acid, and polyuridylic acid as homopolymers or as segments within single RNA molecules. *Biophys. J.* **77**, 3227–3233 (1999).
6. Meller, A., Nivon, L., Brandin, E., Golovchenko, J. & Branton, D. Rapid nanopore discrimination between single polynucleotide molecules. *Proc. Natl. Acad. Sci. USA* **97**, 1079–1084 (2000).
7. Kasianowicz, J.J., Brandin, E., Branton, D. & Deamer, D.W. Characterization of individual polynucleotide molecules using a membrane channel. *Proc. Natl. Acad. Sci. USA* **93**, 13770–13773 (1996).
8. Bayley, H., Braha, O. & Gu, L.Q. Stochastic sensing with protein pores. *Advan. Mater.* **12**, 139–142 (2000).
9. Song, L. *et al.* Structure of staphylococcal alpha-hemolysin, a heptameric transmembrane pore [see comments]. *Science* **274**, 1859–1866 (1996).
10. Senior, M.M., Jones, R.A. & Breslauer, K.J. Influence of loop residues on the relative stabilities of DNA hairpin structures. *Proc. Natl. Acad. Sci. USA* **85**, 6242–6246 (1988).
11. Michael, D. *Chem-Site 3.01*. (Pyramid Learning LLC, Hudson, OH; 1999).
12. Vapnik, V. *The nature of statistical learning theory*, Edn. 2. (Springer-Verlag, New York; 1999).
13. Burges, C.J.C. A tutorial on Support Vector Machines for pattern recognition. *Data Mining and Knowledge Discovery* **2**, 121–167 (1998).
14. Cormen, T.H., Leiserson, C.E. & Rivest, R.L. *Introduction to algorithms*. (McGraw-Hill, New York; 1989).
15. Nievergelt, Y. *Wavelets made easy*. (Birkhauser, Boston; 1999).
16. Jaakkola, T.S. & Hausler, D. Exploiting generative models in discriminative classifiers. In *Advances in neural information processing systems 11*. (eds Kearnes, M.S., Solla, S.A. & Cohn, D.A.) (MIT Press, Cambridge, MA; 1999).
17. SantaLucia, J., Jr. A unified view of polymer, dumbbell, and oligonucleotide DNA nearest-neighbor thermodynamics. *Proc. Natl. Acad. Sci. USA* **95**, 1460–1465 (1998).

