nature biomedical engineering

SUPPLEMENTARY INFORMATION

https://doi.org/10.1038/s41551-020-0591-0

In the format provided by the authors and unedited.

A low-power band of neuronal spiking activity dominated by local single units improves the performance of brain–machine interfaces

Samuel R. Nason[1], Alex K. Vaskov[2], Matthew S. Willsey[1,3], Elissa J. Welle[1], Hyochan An[4], Philip P. Vu[1], Autumn J. Bullard[1], Chrono S. Nu[1], Jonathan C. Kao[5,6], Krishna V. Shenoy[7,8,9,10,11,12], Taekwang Jang[4,13], Hun-Seok Kim[4], David Blaauw[4], Parag G. Patil[1,3,14,15] and Cynthia A. Chestek[1,2,4,15] ✉

[1]Department of Biomedical Engineering, University of Michigan, Ann Arbor, MI, USA. [2]Robotics Graduate Program, University of Michigan, Ann Arbor, MI, USA. [3]Department of Neurosurgery, University of Michigan Medical School, Ann Arbor, MI, USA. [4]Department of Electrical Engineering and Computer Science, University of Michigan, Ann Arbor, MI, USA. [5]Department of Electrical and Computer Engineering, University of California, Los Angeles, Los Angeles, CA, USA. [6]Neurosciences Program, University of California, Los Angeles, Los Angeles, CA, USA. [7]Department of Electrical Engineering, Stanford University, Stanford, CA, USA. [8]Department of Bioengineering, Stanford University, Stanford, CA, USA. [9]Department of Neurobiology, Stanford University, Stanford, CA, USA. [10]The Bio-X Program, Stanford University, Stanford, CA, USA. [11]Wu Tsai Neuroscience Institute, Stanford University, Stanford, CA, USA. [12]Howard Hughes Medical Institute, Stanford University, Stanford, CA, USA. [13]Department of Information Technology and Electrical Engineering, ETH Zürich, Zürich, Switzerland. [14]Department of Neurology, University of Michigan Medical School, Ann Arbor, MI, USA. [15]Neuroscience Graduate Program, University of Michigan, Ann Arbor, MI, USA. ✉e-mail: cchestek@umich.edu

Supplementary Information

Table of Contents

I. Supplementary Methods Figures

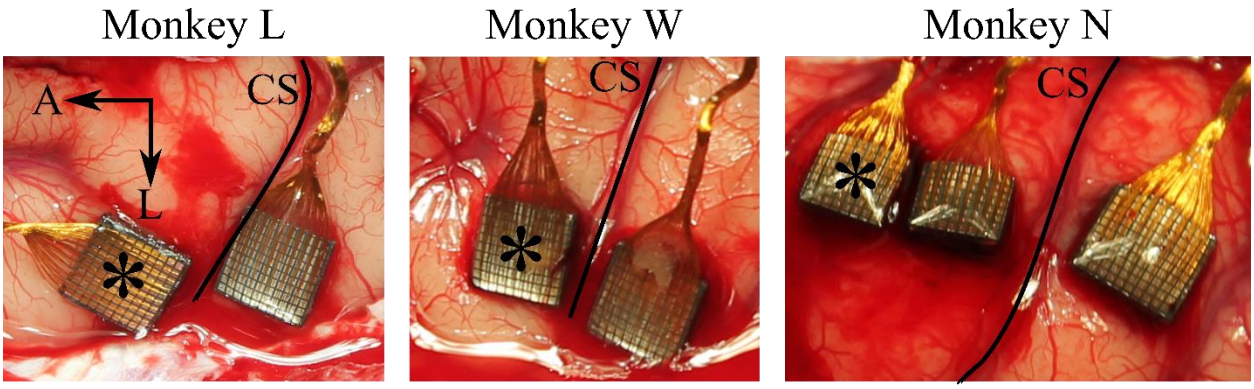

Fig. S1 | Surgical photographs of microelectrode array implants. The arrays labeled with an asterisk are the only arrays used in this study. CS indicates central sulcus, A indicates anterior direction, L indicates lateral direction.

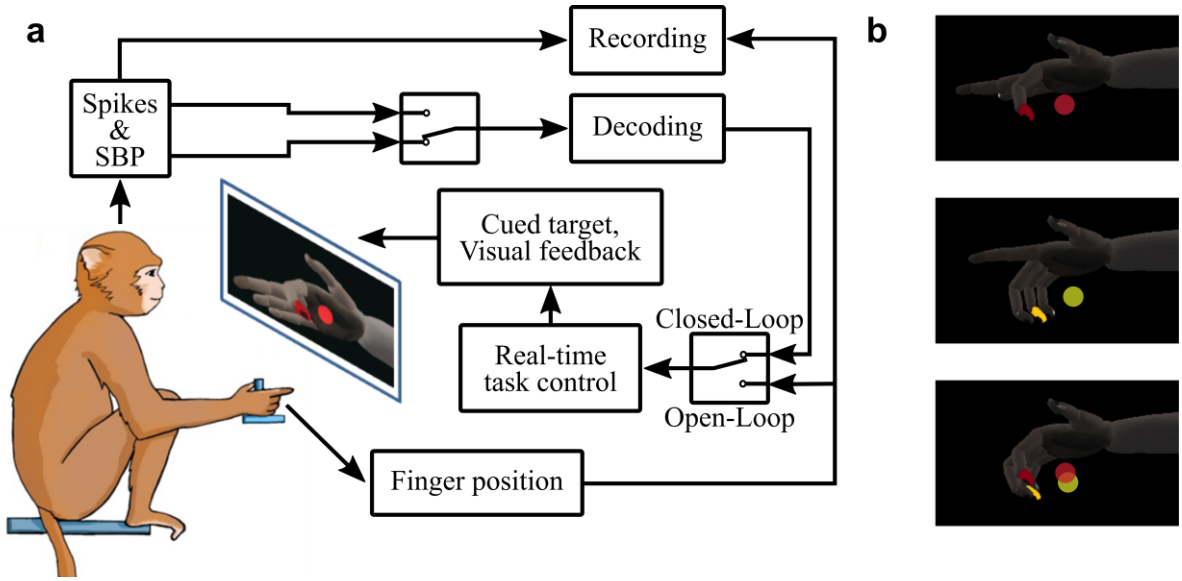

Fig. S2 | Diagram of behavioral task. (**a**) The monkey moved separate groups of fingers (either index alone, middle-ring-small (MRS) alone, or all four fingers together) to hit virtual targets presented on a computer screen. Index alone targets are presented in the diagram. The virtual fingers were controlled by the monkey's physical finger movements or the decoder's translation of the brain activity. Later we trained monkeys W and N to acquire two targets simultaneously, one in red for the index finger and one in yellow for the MRS finger group. (**b**) Example target presentations as shown to the monkeys for index only movements, MRS only movements, and two-target movements.

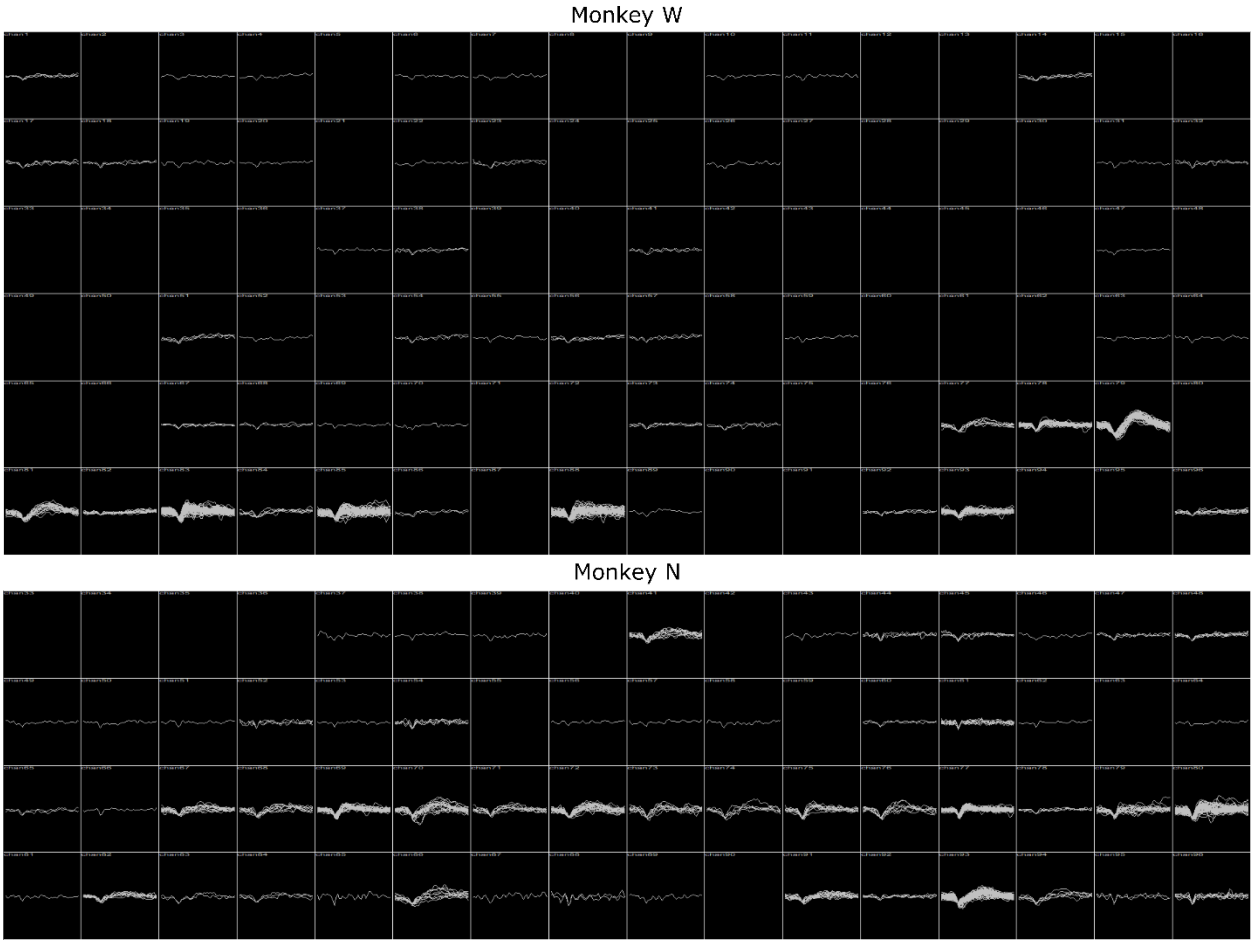

Fig. S1 | Snapshots of monkey W's (top) and monkey N's (bottom) arrays on the days of decoding presented in figure 4a with thresholds set at -4.5×root-mean-square (RMS). Monkey W's array recorded from no more than 10 potential units to achieve the presented decode performance, and monkey N's array exhibited the units as shown from premotor areas.

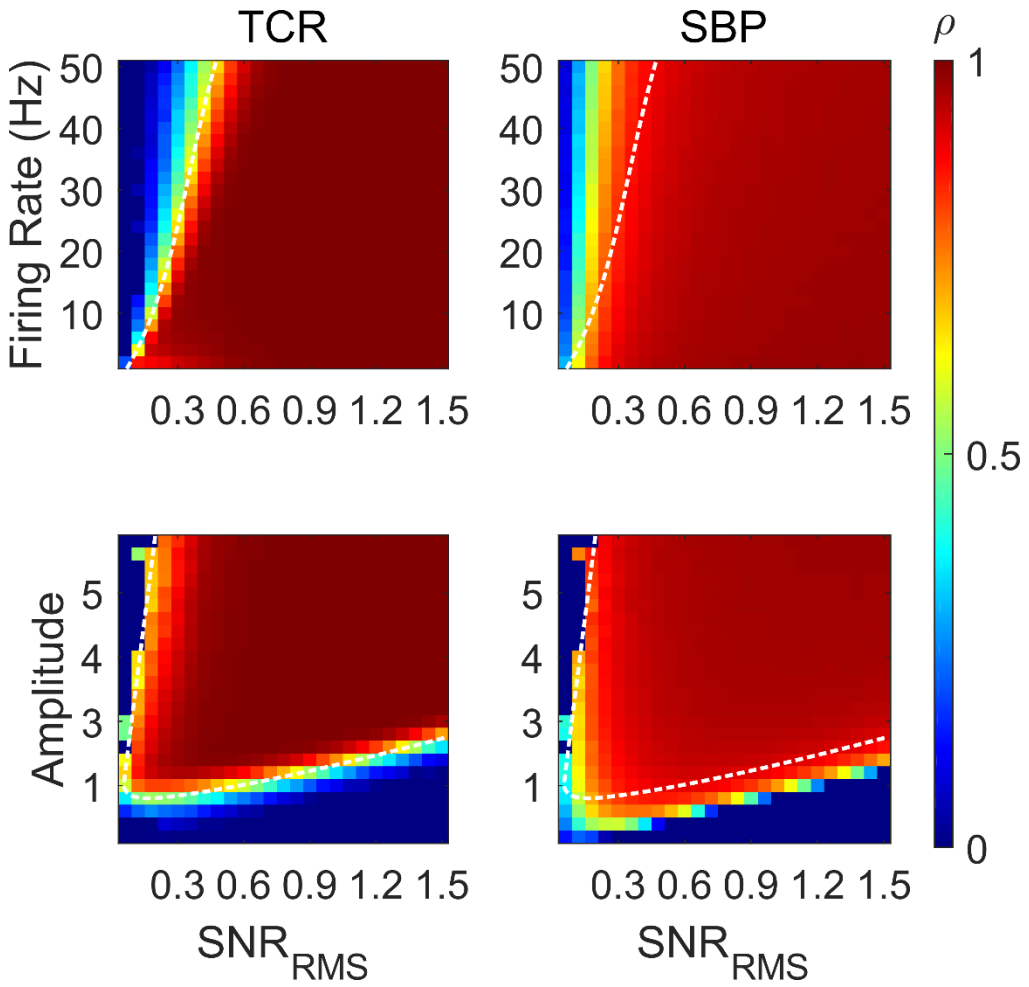

Fig. S2 | SBP and TCR prediction of true firing rate using the signal RMS-based SNR definition. Color is correlation between the true firing rate and that predicted by each feature. The white dotted lines track the yellow region of the left-most TCR plot to better compare low SNR_RMS performance. Top: Firing rate prediction performance with changing firing rate, where the spike amplitude is varied to generate the different SNR_RMS values for a given firing rate. Bottom: Firing rate prediction performance with changing spike amplitude to noise ratio, where the firing rate is varied to generate the different SNR_RMS values for a given ratio. The dark blue regions to the left of all three plots resulted from a firing rate of 0.

II. Supplementary Data

S2.1. RMS-Based SNR Analysis

The definition of SNR used in the main-text is commonly used to represent the SNRs of single units [79, 80, 81]. However, taking a signal processing approach to the discussion of SBP's capability of extracting low-SNR single units provides a different perspective of the spiking power in a signal. To relate the RMS-based SNR definition to neural activity, we must re-evaluate the RMS formula for neurons, whose activity is not defined as a continuous time signal but rather as a relatively sparse firing rate. Thus, we propose that the RMS value of a noiseless spiking signal is the square root of the mean value of one spike's energy, multiplied by

the quantity of spikes in the signal, divided by the total number of samples in the signal. Thus, we can represent our $\mathrm{SNR}_{\mathrm{RMS}}$ definition as

$$\mathrm{SNR}_{\mathrm{RMS}} = \frac{\mathrm{RMS(signal)}}{\mathrm{RMS(noise)}} = \frac{\sqrt{n_{\mathrm{spikes}} \sum_{i=1}^{l_{\mathrm{spike}}} \left(\mathrm{A} \cdot \mathrm{spike}(i)\right)^2}}{\mathrm{RMS(noise)}}$$

where we included the term A to maintain the capability of scaling the neuron's relative amplitude. Consequently, the $\mathrm{SNR}_{\mathrm{RMS}}$ values corresponding to realistic firing rates and amplitudes seen in motor cortex, such as those presented in figure 2, will be small. As we now have a four-dimensional problem (firing rate, relative amplitude, $\mathrm{SNR}_{\mathrm{RMS}}$, and correlation coefficient), we present results in figure S4 in a similar format to figure 2 but in parameterized plots.

The results in the top of figure S4 suggest that SBP is capable of accurately predicting neural activity at lower $\mathrm{SNR}_{\mathrm{RMS}}$ values than TCR, except for small, rare motor cortical spiking frequencies in the bottom left of the plots (e.g. average firing rates below 2Hz are not informative for decoders). The bottom of figure S4 suggests that SBP can extract the firing patterns of smaller amplitude units (equivalent to the results from the amplitude-based SNR definition in figure 2a), though not at the extremely low $\mathrm{SNR}_{\mathrm{RMS}}$ values (bottom-left corner of the plots) that TCR is better able to extract.

S2.2. Open-Loop Decoding Analysis

We performed open-loop Kalman filter decodes to compare SBP to 300-6,000Hz wideband SBP (i.e. from others [32]), low-bandwidth TCR, TCR, sorted unit firing rates, and combined sorted unit firing rate with hash rates. We optimized two decoder aspects for each feature and each set of data: the number of historical neural data bins, and the offset between neural data and behavior, anywhere between 0 and 5 bins. For low-bandwidth TCR, we also optimized the threshold for each set of data. Figure S5a illustrates predictions made by each optimized decoder along with the true hand position overlaid. SBP achieved statistically equivalent or better correlation coefficients than low-bandwidth TCR, TCR, and single unit firing rate in all animals and additionally single unit firing rate with hash rate in monkeys L and N (results in table S1 with optimized parameters in table S2). We also found that wideband SBP outperformed all other neural features, which complies with previous results, and that there is a small drop in performance resulting from the bandwidth reduction from ~6kHz to ~1kHz as we found previously [32, 50].

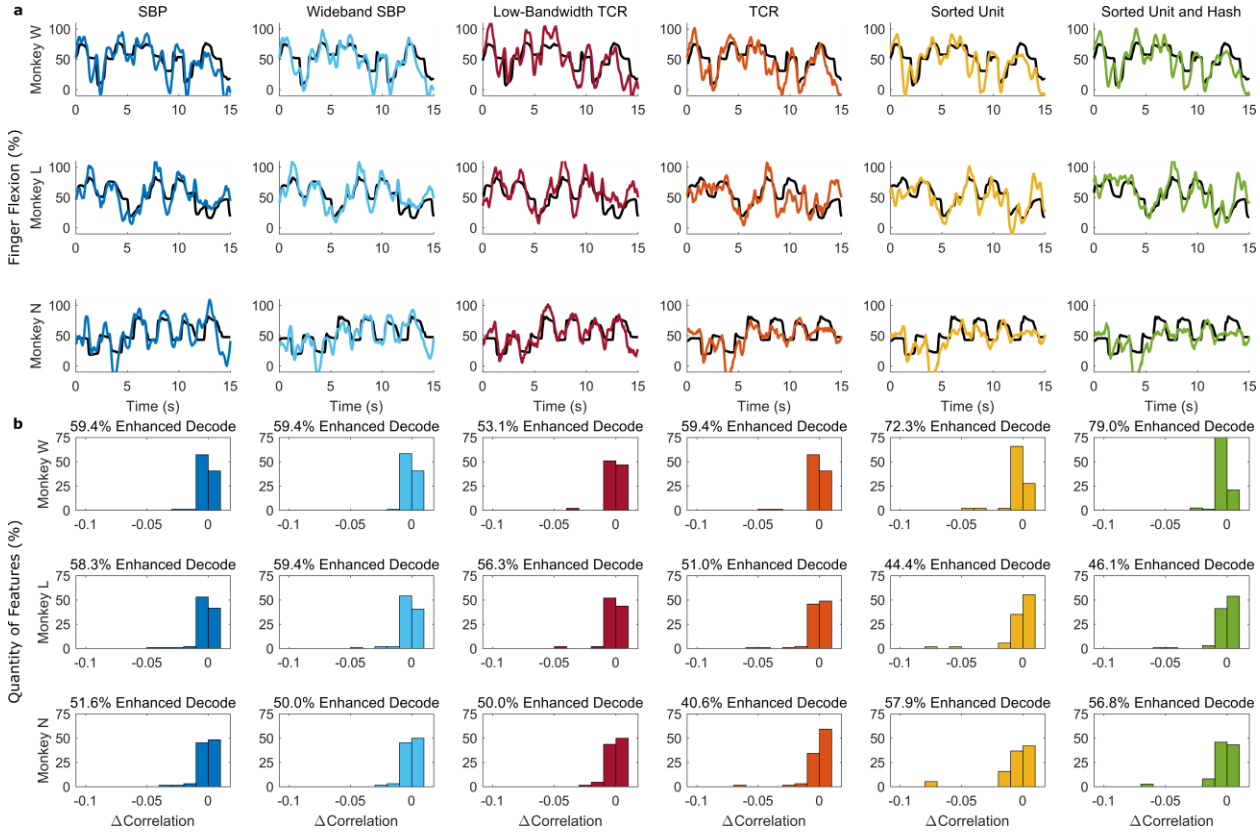

Fig. S3 | Open-loop decode results using a standard Kalman filter to decode SBP, 300-6,000Hz wideband SBP, low-bandwidth TCR, TCR, single unit firing rate, and single unit firing rate with neural hash. (**a**) Decode traces overlaid on the actual finger positions in black. (**b**) Performance lost by omitting individual channels or units during decodes. The histograms represent the percentage of total channels or units that, when left out of the decode individually, resulted in a Pearson's correlation loss within each bin. Lower numbers indicate the omitted channel was more valuable towards the overall decode performance. Above each histogram is the percentage of available features with negative correlation differences.

Since SBP can extract unit-specific neural activity from SNRs below typical threshold levels, we also investigated the quantity of recorded channels valuable for decoding. We re-processed the optimized decodes, except trained and tested on neural data lacking one feature at a time. Then, we calculated the new correlation between the finger position predicted by the decoder lacking a feature and the actual finger position and found the difference between it and the correlation resulting from all features included. Thus, a more negative difference means the excluded feature was more valuable to the decoder. The histograms in figure S5b represent these differences in correlation. As follows from figure 2, there were generally more SBP features that resulted in negative correlation differences than threshold crossings, suggesting SBP may extract decodable neural information from channels that threshold crossings may find invaluable. Sorted unit features expectedly had some higher values, as manual feature extraction excludes lots of invaluable information.

Subject ρ RMSE	**SBP**	**WSBP**	**LbTCR**	**TCR**	**SU**	**SUH**
Monkey W	0.730^{a}	0.770^{b}	0.694^{c}	0.727^{a}	0.722^{a}	0.750^{d}
n = 45,124	0.176^{a}	0.161^{b}	0.190^{c}	0.168^{b}	0.178^{a}	0.167^{b}
Monkey L	0.760^{a}	0.782^{b}	0.715^{c}	0.694^{d}	0.665^{e}	$0.695^{c,d}$
n = 37,203	0.149^{a}	0.143^{a}	0.163^{b}	0.164^{b}	0.188^{c}	$0.177^{b,c}$
Monkey N	$0.708^{a,b}$	0.732^{a}	0.684^{b}	0.689^{b}	0.676^{b}	0.686^{b}
n = 7,547	0.194^{a}	0.185^{a}	0.173^{a}	0.171^{a}	0.174^{a}	0.172^{a}

Table S1 | Optimized open-loop decode results. Correlation ρ is Pearson's correlation coefficient between the actual finger trace and that predicted by each feature. Root-mean-squared error (RMSE) was also calculated between the actual and predicted finger traces. *n* is the number of samples 50ms in size used for each calculation. Given that the optimal amount of historical bins and offset between the neural data and behavior differed slightly between features, the minimum number of bins per dataset is shown. The letters noted by each statistic indicate statistically similar numbers within one subject (i.e. numbers with the same letter are not statistically different, $p < 1 \times 10^{-4}$, two-tailed two-sample *z*-tests on Fisher's *z*-scores for the correlation coefficients comparisons, one-tailed two-sample Wilcoxon rank-sum tests for the RMSE comparisons). WSBP: wideband SBP, LbTCR: low-bandwidth TCR, SU: sorted unit firing rate, SUH: sorted unit firing rate with neural hash rate.

Optimized Parameter	**Monkey**	**SBP**	**WSBP**	**LbTCR**	**TCR**	**SU**	**SUH**
No. Additional Bins	W	5	5	4	1	5	5
	L	1	1	1	1	4	4
	N_1*	4	5	5	0	0	0
	N_2	5	5	0	0	0	0
Bin Offset	W	0	0	0	0	0	0
	L	0	0	0	0	0	0
	N_1*	0	0	0	3	2	3
	N_2	0	0	3	2	2	2
Threshold	W	-	-	2.40	-	-	-
	L	-	-	2.25	-	-	-
	N_1*	-	-	2.20	-	-	-
	N_2	-	-	2.15	-	-	-

Table S2 | Parameters optimized per dataset for open-loop decoding. Monkey N's dataset 1 (marked by *) was not used for any other analyses, but included here to demonstrate variability in optimal parameters within one animal. “No. Additional Bins” is the quantity of additional historical neural bins, between 0 and 5, that achieved the highest correlation coefficient. “Bin Offset” is the delay, in bins from 0 to 5, between the neural activity and the behavior that achieved the highest correlation coefficient. “Threshold” is the RMS threshold that achieved the highest correlation coefficient.

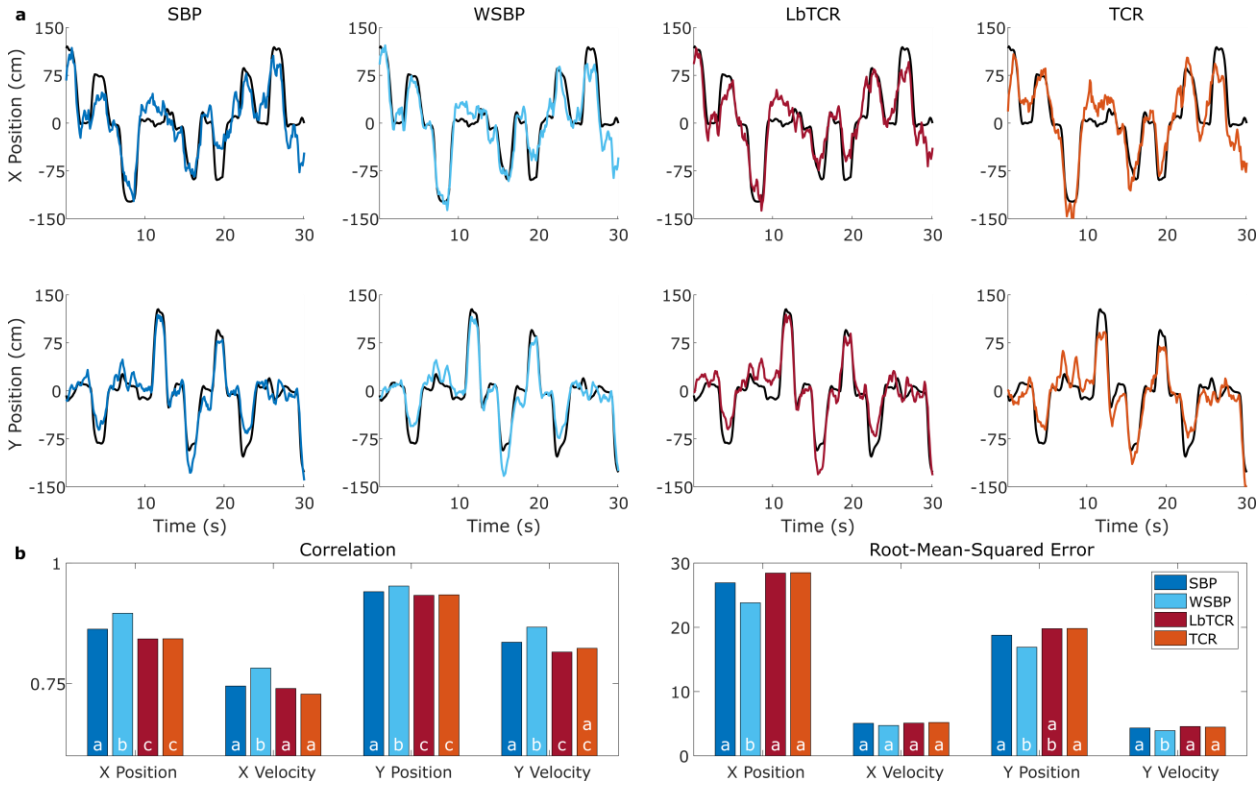

Fig. S4 | Open-loop two-dimensional cursor control task for monkey J over 506 trials, or *n* = 9,295 samples 50ms in size. (**a**) Example decodes for SBP, wideband SBP, low-bandwidth TCR, and TCR. (**b**) Decode statistics. Bars in a group labelled with different letters are statistically different. $p < 1 \times 10^{-3}$, two-sided two-tailed *z*-test for Pearson's correlation coefficients. $p < 1 \times 10^{-3}$, two-tailed two-sample Wilcoxon rank-sum test for root-mean-squared errors.

S2.3. Two-Dimensional Cursor Control

To validate that SBP maintains performance in a clinically-relevant state-of-the-art task, we compared SBP, wideband SBP, TCR (threshold at -4.5RMS), and low-bandwidth TCR (thresholded at 2.25RMS, the averaged optimal value from the finger offline decodes) at predicting two-dimensional cursor control. Specifically, we analyzed some offline datasets from monkey L from a prior study (L120502) and monkey J (J121009), which is the same animal with the same implants as described previously [57]. Briefly, the monkey was trained to perform center-out-and-back arm reaches along 8 directions while 30kSps neural activity was recorded. We used a standard Kalman filter to decode 506 trials within one day for monkey J using all 96 channels and 117 trials within one day for monkey L using 95 channels. We excluded channels without units from all decodes of all features, and all features were binned in 50ms bins. To avoid over-fitting, we 10-fold cross-validated each dataset. Statistical tests followed the procedures defined in the open-loop decode methods, section 3.5.

We found that SBP achieved statistically as good or better decoding performance than TCR and low-bandwidth TCR, which is consistent with the results of our offline one-dimensional finger tasks. Wideband SBP achieved the highest performance of any feature, as

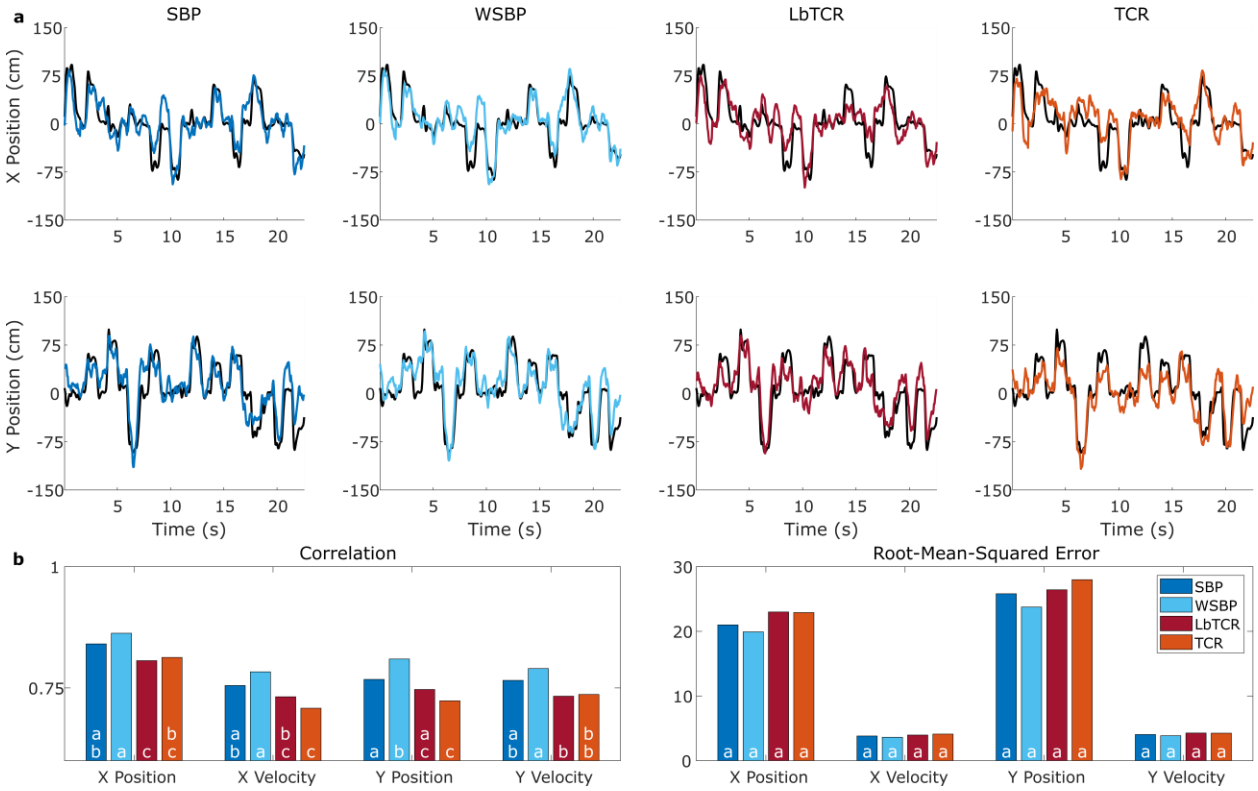

Fig. S5 | Open-loop two-dimensional cursor control task for monkey L over 117 trials, or *n* = 2,283 samples 50ms in size. (**a**) Example decodes for SBP, wideband SBP, low-bandwidth TCR, and TCR. (**b**) Decode statistics. Bars in a group labelled with different letters are statistically different. $p < 1 \times 10^{-3}$, two-sided two-tailed *z*-test for Pearson's correlation coefficients. $p < 1 \times 10^{-3}$, two-tailed two-sample Wilcoxon rank-sum test for root-mean-squared errors.

was expected from others' results and our one-dimensional results [32] Figures S6 and S7 compare the decodes for a few movements along with statistics. Monkey J's SBP had significantly higher positional correlation coefficients and as good or better velocity correlation coefficients than TCR features, and monkey L's SBP correlation coefficients were as good or better than TCR features ($p < 1 \times 10^{-3}$, two-sided two-sample *z*-test). Root-mean-squared errors were all statistically similar between SBP and TCR features, though numerically, SBP root-mean-squared errors were lower than the TCR features for both monkeys ($p < 1 \times 10^{-3}$, two-tailed two-sample Wilcoxon rank-sum test). Extrapolating these results to the outcomes of the one-dimensional open-loop and closed-loop decodes suggests that SBP may maintain performance at least as good as TCR in multi-dimensional closed-loop tasks.

Using these two-dimensional datasets, we can also investigate how robust each neural feature is to various amounts of noise. The results of section 1.1 suggest that low-bandwidth features are more robust to noise than higher bandwidth threshold crossings when predicting unit firing rates, so we superimposed white noise at RMS levels from 0 to 30μV in increments of 1μV before extracting features. We selected a maximum of 30μV of noise because, at that noise level, decode correlations were around 0.5, a number low enough to suggest poor performance.

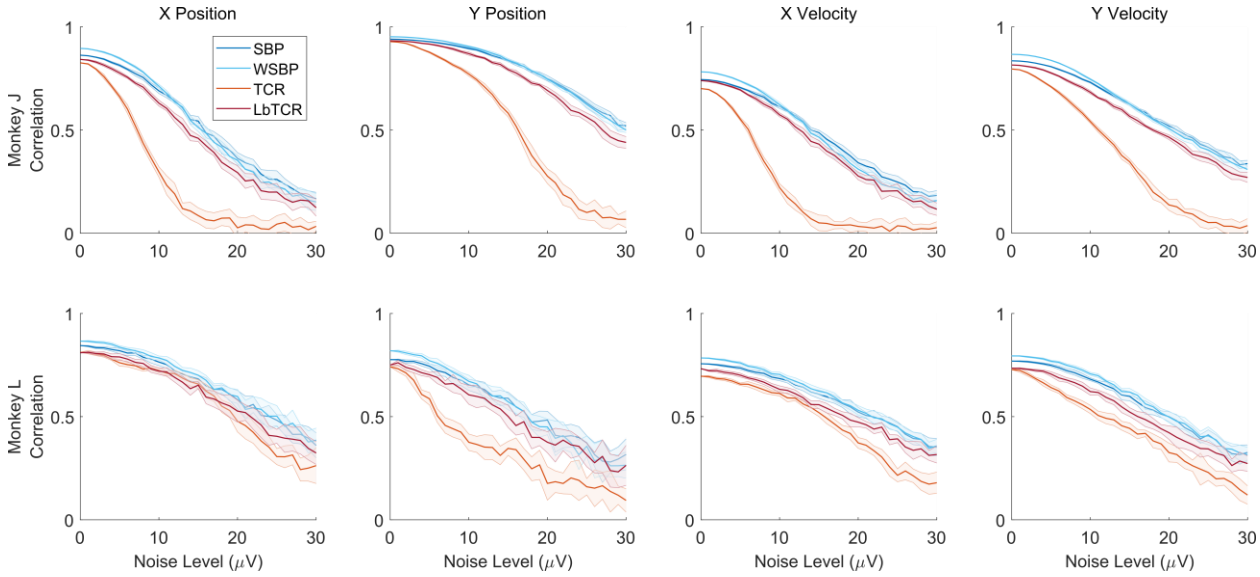

Fig. S6 | Open-loop two-dimensional cursor control correlation coefficients for monkeys J (top, 506 trials or *n* = 9,295 samples 50ms in size) and L (bottom, 117 trials or *n* = 2,283 samples 50ms in size) with manually injected white noise at various levels from 0 to 30μV. Solid lines represent the mean Pearson's correlation coefficient of each feature's decode at each injected noise level. Shaded area surrounding each solid line represents standard deviation across the 10 trials of randomly generated noise at each level.

Noise Added to Generate 20% Performance Loss $\mu V, \frac{nV}{\sqrt{Hz}}$								
Neural	Monkey J				Monkey L			
Feature	X Pos.	Y Pos.	X Vel.	Y Vel.	X Pos.	Y Pos.	X Vel.	Y Vel.
SBP	9, 340	19, 718	10, 378	12, 454	15, 567	12, 454	15, 567	13, 491
WSBP	9, 119	19, 252	9, 119	12, 159	15, 199	11, 146	14, 185	13, 172
LbTCR	8, 302	17, 643	9, 340	11, 416	15, 567	10, 378	13, 491	12, 454
TCR	5, 66.2	11, 146	4, 53.0	7, 92.7	15, 199	4, 53.0	13, 172	7, 92.7

Table S3 | The amount of noise added to each channel before extracting features that resulted in a 20% loss in performance in $\mu V, nV/\sqrt{Hz}$. Higher numbers imply more noise could be tolerated to maintain 80% of the maximum performance. Results averaged from 10 random generations of noise for each noise level.

We generated random noise ten times at each noise level to obtain the average effects of noise on decoding performance.

The averaged results are plotted in figure S8. We found that broadband TCR showed the most rapid dropoff of decode performance with increasing noise level, where the other features remained in the same neighborhood of performance. Additionally, SBP always maintained a higher decoding performance than low-bandwidth TCR at any level of injected noise, suggesting the performance enhancements of SBP are maintained even in lower-quality recordings. Lastly,

Circuit Type	NEF	Noise V_{RMS}	U_T	k	T	Bandwidth
Low-Bandwidth	4.0	2μV	26.7mV	1.38×10^{-23}	310K	700Hz
High-Bandwidth						5,750Hz

Table S4 | Summary of analog parameters for integrated circuit simulations. NEF is the noise-efficiency factor, U_T is the thermal voltage, k is the Boltzmann Constant, T is temperature in Kelvin.

table S3 demonstrates that the low-bandwidth features can accommodate almost three times the amount of noise as the higher bandwidth features to maintain 80% of the maximum performance, with SBP able to withstand at least as much noise as low-bandwidth TCR. Overall, this suggests that eliminating noise by restricting the passband to the spiking band increases low-bandwidth TCR and SBP's tolerance of noisy recordings, with the optimal low-bandwidth feature being SBP as it persistently achieved higher performance.

S2.4. Integrated Circuit Simulations

Many groups have attempted to reduce circuit power by using lower power process nodes, optimized digital hardware, and data compression [7, 19, 5, 6, 64]. To better realize the applicable benefits of low-bandwidth neural recording, we simulated fully customized low- and high-bandwidth-recording integrated neural interfaces using established equations to compare their estimated power consumptions. We defined the following requirements for our simulated circuits: acquire data, extract neural features, decode the features, and transmit the decoded values off-chip. The necessary components to accomplish those goals included analog amplifiers and filters, analog to digital converters (ADCs), feature extraction circuitry at a 50ms bin size, a Kalman filter hardware accelerator to decode, and a controller area network (CAN) interface operating at 100kbaud to communicate the predictions.

To estimate the power consumption of the analog front-ends, we first calculated the power of the amplifier front-end chain using the noise efficiency factor (NEF) formula [82]:

$$I = \left(\frac{NEF}{V_{RMS}}\right)^2 \frac{\pi \cdot U_T \cdot 4kT \cdot BW}{2}$$

with a thermal voltage U_T of 26.7mV, Boltzmann constant k of 1.38×10^{-23}, temperature T of 310K, and signal bandwidth BW. These parameters are summarized in table S4. Even though some amplifiers [83] have approached the ideal NEF of 1, the NEF of a single bipolar junction transistor, neural recording amplifiers typically show NEF values larger than 4 because of other stringent specifications for their robust operation, such as high input impedance, high common mode ratio, and high power supply rejection ratio [84, 85, 86, 87]. Therefore, we assumed an NEF of 4.0. For the robust detection of neural spikes whose amplitude ranges around 100μV, 2μV input referred noise is assumed. The signal bandwidth was set to 300-1,000Hz for the low-bandwidth amplifier and 250-6,000Hz for the high-bandwidth amplifier for accurate spike detection (which is in the range of similar devices presented previously) [7, 19, 6, 64]. We chose a primary supply voltage of 3.3V for a high-resolution ADC (described later) and a matched last-stage amplifier to support high amplifier linearity, with an optional 1.2V supply for the amplifier chain given multiple supply domains. We computed the power consumption by multiplying the

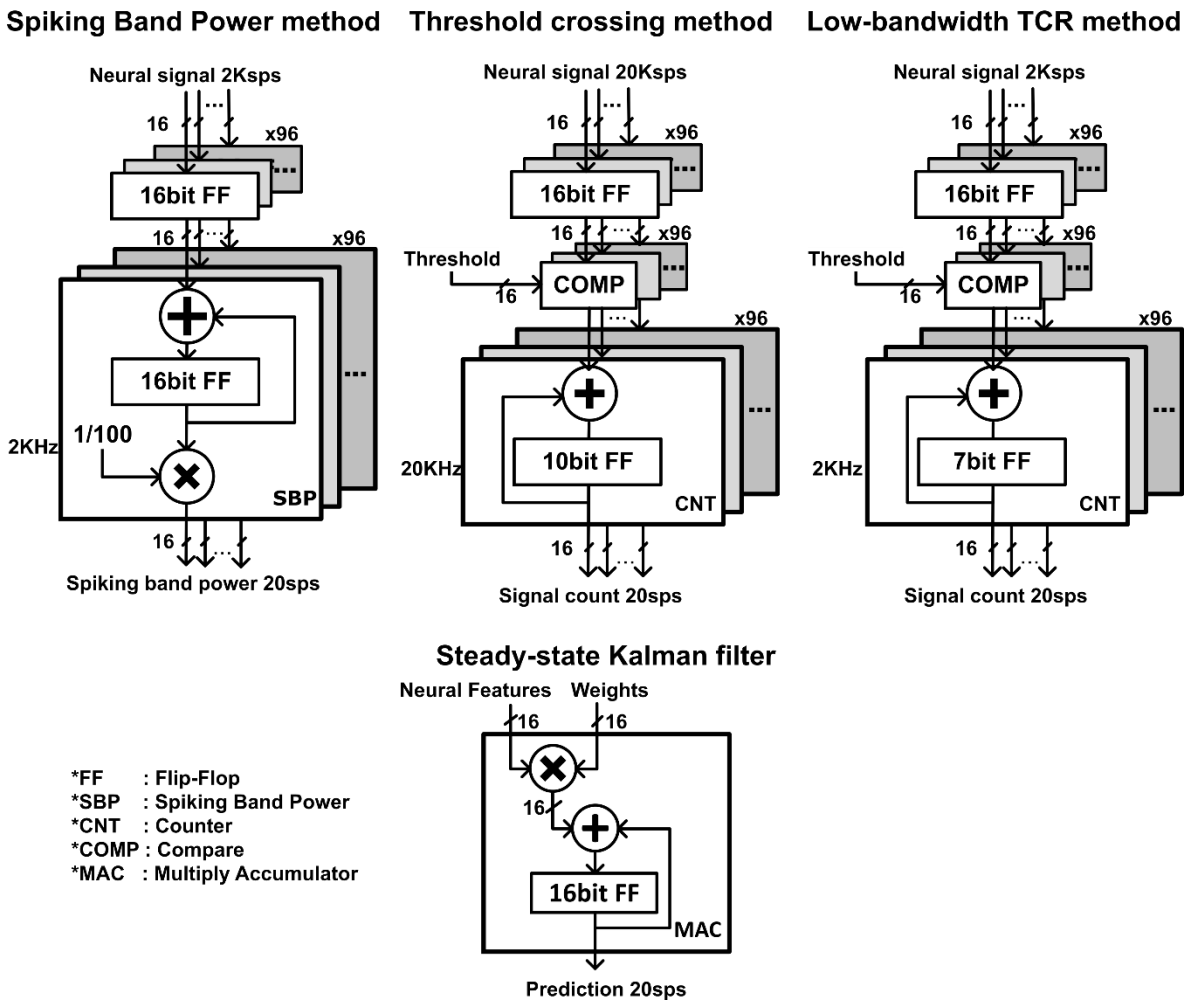

Fig. S7 | Digital circuit designs. (Left) Digital circuit designs for SBP feature extraction with 2kSps neural recording. 16-bit data of 96 channels are stored in flip-flops to be accumulated and averaged for SBP. This logic executes at 2kHz to match the sampling rate of SBP. (Middle) Digital circuit design for threshold crossing rate with 20kSps neural recording. 16-bit data of 96 channels are stored in flip-flops to count the number of events where the sampled data has crossed the threshold value. This logic must execute at 20kHz to match the sampling rate of the broadband recordings. (Right) Digital circuit design for low-bandwidth TCR with 2kSps neural recording. 16-bit data of 96 channels are stored in flip-flops to count the number of events where the 2kSps data has crossed the threshold value. This logic executes at 2kHz to match the sampling rate of the low-bandwidth recordings. (Bottom) Computational unit for the steady-state Kalman filter using a 16-bit multiply-accumulator unit (MAC).

estimated current consumption by the voltage supply level and by 96 for the total quantity of channels. We estimated the power consumption of each 16-bit analog to digital converter using the Schreier Figure-of-Merit (FoM_S) formula, whose state-of-the-art envelope at a sampling rate less than 5MHz is 178dB [88]:

$$FoM_S = SNDR(dB) + 10\log_{10}\left(\frac{BW}{P}\right)$$

We assumed a signal-to-noise-and-distortion-ratio (SNDR) of 96dB in anticipation of potential movement or electrical stimulation artifacts (which can be as large as 20mV [89]), resulting in an effective number of bits of 15.6 to accommodate artifacts and neural recordings with the more typical low SNR units or the occasional high SNR units. Additionally, we assumed sampling frequencies of 2kHz for the low-bandwidth circuit or 20kHz for the high-bandwidth circuit and one converter for each of the 96 channels. The analog to digital converters' outputs were assumed to be transmitted in parallel to the digital circuitry.

We designed the digital circuits in Verilog HDL [90], synthesized using Synopsys Design Compiler Version L-2016.03-SP2 in TSMC 0.18 UM MIXED SIGNAL GENERAL PURPOSE II 1P6M/1P5M SALICIDE 1.8V/3.3V technology, and simulated using Cadence NCverilog simulator Version 15.20-s005, with power estimated using Synopsys PrimeTime PX Version M-2017.06. To represent the digital circuits' power consumptions in a more advanced process node, we have additionally performed all digital simulations in 40nm technology. It should be noted that the analog components' noise constraints result in larger than minimum size analog transistors, even using the older 180nm technology. This means that the analog domain will not gain any power advantage by using a smaller process size. Figure S9 contains diagrams illustrating the components of our circuit designs. The input vectors of the gate-level simulations were randomized 2kSps neural signals for the low-bandwidth circuit and randomized 20kSps neural signals for the high-bandwidth circuit. For relevance to this manuscript and the devices presented in literature, we simulated two low-bandwidth circuits, one that extracted SBP as the neural feature and the other that extracted low-bandwidth TCR, while the high-bandwidth circuit extracted TCR as the neural feature. The low-bandwidth SBP circuit was composed of 96 copies of 16-bit registers, a 16-bit adder, and 96 copies of average generators. The TCR circuits were composed of 96 copies of 16-bit registers, a threshold comparator, and 96 copies of 10-bit counters for the high-bandwidth circuit or 7-bit counters for the low-bandwidth circuit. We estimated the power consumptions for the Kalman filter and CAN blocks using the same simulation software. For all digital logic simulated, we used a near-threshold supply voltage of 0.6V, which we found to be low power while adequate for the performance requirement and logic depth.

Our simulations and calculations demonstrated that the overall power consumption of the high-bandwidth circuit is dominated by the front-end, needing 11mW for a 3.3V amplifier supply or 8.0mW for a 1.2V amplifier supply and 1.6µW for the digital back-end components (in 180nm, 0.55µW in 40nm). When we reduced the front-end bandwidth to the spiking band, the simulation predicted a drastic cut in the power consumption to 1.1mW for a 3.3V amplifier supply or 0.65mW for a 1.2V amplifier supply, while the consumption of the digital components remained low at 1.3µW (in 180nm, 0.47µW in 40nm). The power consumptions of the individual circuit components are listed in table S5. The results of these simulations suggest that, regardless of the process node, the power consumption of integrated neural devices is dominated by the analog front-end, which can be cut by at least 90% when restricted to the 300-1,000Hz band, despite maintaining a power efficiency factor (PEF) near the state-of-the-art at 19.2 with a 1.2V supply [91]. However, TCR can accurately predict firing patterns at much narrower bandwidths, as demonstrated in figure 2. Simulating the low-bandwidth TCR circuit revealed that it can

Circuit Type	Analog (mW)		Digital (μW)			Total (mW)
	Amplifier	ADC	Feature Extractor	3-DoF SSKF	CAN	
	(3.3V, 1.2V)		(180nm, 40nm)			(3.3V, 1.2V)
Low-Bandwidth (SBP)	(0.64, 0.23)	0.42	(0.67, 0.22)	(0.64, 0.22)	0.032	(1.1, 0.65)
Low-Bandwidth (TCR)	(0.64, 0.23)	0.42	(0.14, 0.05)	(0.64, 0.22)	0.032	(1.1, 0.65)
High-Bandwidth (TCR)	(5.2, 1.9)	6.1	(0.89, 0.30)	(0.64, 0.22)	0.032	(11, 8.0)

Table S5 | Breakdown of power estimation for each integrated circuit component. (3.3V, 1.2V) indicates the power consumptions of the amplifiers calculated at a 3.3V supply or a 1.2V supply. (180nm, 40nm) indicates the digital circuit power consumption when simulated at a 180nm or a 40nm process node. TCR is threshold crossing rate, ADC is analog-to-digital converter, DoF is degrees of freedom, SSKF is steady-state Kalman filter, CAN is controller area network.

further reduce the power consumption of the digital components to a total of 0.81μW (in 180nm, 0.30μW in 40nm). While extracting low-bandwidth TCR using the spiking band would result in an equivalent noise spectral density as SBP and save equivalent amounts of amplifier power and more digital power than the low-bandwidth SBP circuit, the analog front-end remains the dominating power-hungry component. Thus, since the decoding results in figures S5 – S8 suggest that SBP persistently achieves higher decoding performances than low-bandwidth TCR and low-bandwidth TCR enables a small 0.05% power savings over SBP, we conclude that SBP would remain the highest performing neural feature for low-power decoding in all cases.

S3. Supplementary Bibliography

The bibliography below is a continuation of that in the main manuscript containing items cited exclusively in this Supplementary Information document.

79. Wark, H. A. C. et al. A new high-density (25 electrodes/mm2) penetrating microelectrode array for recording and stimulating sub-millimeter neuroanatomical structures. *J. Neural Eng.* **10**, 045003 (2013).

80. Xu, H. et al. Acute in vivo testing of a conformal polymer microelectrode array for multi-region hippocampal recordings. *J. Neural Eng.* **15**, 15 (2018).

81. Byun, D. et al. Recording nerve signals in canine sciatic nerves with a flexible penetrating microelectrode array. *J. Neural Eng.* **14**, 14 (2017).

82. Steyaert, M. S., Sansen, W. M. & Zhongyuan, C. A micropower low-noise monolithic instrumentation amplifier for medical purposes. *IEEE J. Solid State Circuits* **22**, 1163–1168 (1987).

83. Shen, L., Lu, N. & Sun, N. A 1V 0.25 μW inverter-stacking amplifier with 1.07 noise efficiency factor. *IEEE J. Solid-State Circuits* **53**, 3, 896-905 (2018).

84. Chandrakumar, H. & Markovic, D. An 80-mVpp linear-input range, 1.6-G Ω input impedance, low-power chopper amplifier for closed-loop neural recording that is tolerant to 650-mVpp common-mode interference. *IEEE J. Solid State Circuits* **52**, 2811–2828 (2017).

85. Chandrakumar, H. & Markovic, D. A high dynamic-range neural recording chopper amplifier for simultaneous neural recording and stimulation. *IEEE J. Solid State Circuits* **52**, 645–656 (2017).

86. Muller, R. et al. A miniaturized 64-channel 225 μW wireless electrocorticographic neural sensor. *IEEE International Solid-State Circuits Conference Digest of Technical Papers (ISSCC)*, **2014**, 412-413 (2014).

87. Mahajan, A. et al. A 64-channel ultra-low power bioelectric signal acquisition system for brain–computer interface. *IEEE Biomedical Circuits and Systems Conference (BioCAS)*, **2015**, 1-4 (2015).

88. Murmann, B. *ADC Performance Survey 1997–2020*; http://web.stanford.edu/~murmann/adcsurvey.html

89. Young, D. et al. Signal processing methods for reducing artifacts in microelectrode brain recordings caused by functional electrical stimulation title signal processing methods for reducing artifacts in microelectrode brain recordings caused by functional electrical stimulation. *J. Neural Eng.* **15**, 026014 (2018).

90. *IEEE Standard for Verilog Hardware Description Language* (IEEE, 2006).

91. Ng, K. A. & Xu, Y. P. A low-power, high CMRR neural amplifier system employing CMOS inverter-based OTAs with CMFB through supply rails. *IEEE J. Solid State Circuits* **51**, 724–737 (2016).