Article

Task-agnostic exoskeleton control via biological joint moment estimation

https://doi.org/10.1038/s41586-024-08157-7

Received: 17 October 2023

Accepted: 4 October 2024

Published online: 13 November 2024

Check for updates

Dean D. Molinaro[1,2,5,8]✉, Keaton L. Scherpereel[1,2,6,8], Ethan B. Schonhaut[1], Georgios Evangelopoulos[3,7], Max K. Shepherd[4] & Aaron J. Young[1,2]

Lower-limb exoskeletons have the potential to transform the way we move[1–14], but current state-of-the-art controllers cannot accommodate the rich set of possible human behaviours that range from cyclic and predictable to transitory and unstructured. We introduce a task-agnostic controller that assists the user on the basis of instantaneous estimates of lower-limb biological joint moments from a deep neural network. By estimating both hip and knee moments in-the-loop, our approach provided multi-joint, coordinated assistance through our autonomous, clothing-integrated exoskeleton. When deployed during 28 activities, spanning cyclic locomotion to unstructured tasks (for example, passive meandering and high-speed lateral cutting), the network accurately estimated hip and knee moments with an average R^2 of 0.83 relative to ground truth. Further, our approach significantly outperformed a best-case task classifier-based method constructed from splines and impedance parameters. When tested on ten activities (including level walking, running, lifting a 25 lb (roughly 11 kg) weight and lunging), our controller significantly reduced user energetics (metabolic cost or lower-limb biological joint work depending on the task) relative to the zero torque condition, ranging from 5.3 to 19.7%, without any manual controller modifications among activities. Thus, this task-agnostic controller can enable exoskeletons to aid users across a broad spectrum of human activities, a necessity for real-world viability.

Lower-limb exoskeletons promise to reinvent human mobility by augmenting our capability and increasing longevity[15,16]. However, within powered exoskeleton technology lies a critical limitation: the controllers—which in many cases were optimized through extensive experimentation—only work for a single task or small set of tasks, offering little adaptability beyond passive devices[17,18]. Switching between tasks typically requires a 'high-level' task classification (for example, level walking, incline walking and stair ascent) often toggled manually or in some cases by an autonomous classifier[1,19–24]. Within each class, a 'mid-level' controller computes desired exoskeleton assistance, which is often tuned on a user- and task-specific basis[3,11,25]. For highly repetitive cyclic tasks, assistance is often designed as a function of time or phase[3,10,12,24]. For some gravity-fighting non-cyclic tasks, such as squats or sit-to-stand, impedance controllers assist, removing any dependence on time[26–28]. Although this approach has worked well for many laboratory-based experiments, this highly constrained discretization contrasts with the fluidity of natural human movement; we shuffle and side-step as we navigate a busy kitchen, stop our jog to take in a scenic view and regather our balance to again try the door that was heavier than we had anticipated. Our median walking bouts are a mere four steps[29]. Unstructured, non-cyclic and transitory tasks make up a large portion of our movements and interactions with the environment, but current exoskeleton controllers are incapable of recognizing or assisting these tasks. In fact, the expansion of traditional classification-based high-level control architectures to encompass these unstructured movements is intractable owing to the sheer number of movements that must be defined.

We have developed a task-agnostic exoskeleton controller that short-circuits the need for high-level task classification or gait phase estimation by basing assistance on a fundamental, continuous physiological state: the human's biological joint moment (Fig. 1a and Supplementary Video 1). Biological moment can be calculated using optical motion capture and high-fidelity force plates to measure interactions with the ground[30] but cannot be measured or solved for analytically through available wearable sensors, owing to sensor noise and incomplete information (particularly shear forces with the ground). Instead, biological joint moments can be estimated from wearable sensor data, often by including optimization or learning methods to account for incomplete sensor information[13,14,31–36], but very few studies have begun to explore the implications of using this technology in the control loop[13,14,35,36]. In these previous works, however, using instantaneous biological joint moment estimates in the control loop has shown substantial promise. Gasparri et al.[13] developed a joint moment-based ankle exoskeleton controller, which has shown large metabolic benefits across inclines and declines, stairs and mixed terrain in both unimpaired individuals and those with cerebral palsy[13,37–39]. Additionally,

[1]George W. Woodruff School of Mechanical Engineering, Georgia Institute of Technology, Atlanta, GA, USA. [2]Institute for Robotics and Intelligent Machines, Georgia Institute of Technology, Atlanta, GA, USA. [3]X, The Moonshot Factory, Mountain View, CA, USA. [4]College of Engineering, Bouvé College of Health Sciences, and Institute for Experiential Robotics, Northeastern University, Boston, MA, USA. [5]Present address: Boston Dynamics AI Institute, Cambridge, MA, USA. [6]Present address: Skip Innovations, San Francisco, CA, USA. [7]Present address: Google, Mountain View, CA, USA. [8]These authors contributed equally: Dean D. Molinaro, Keaton L. Scherpereel. ✉e-mail: molinarodean@gmail.com

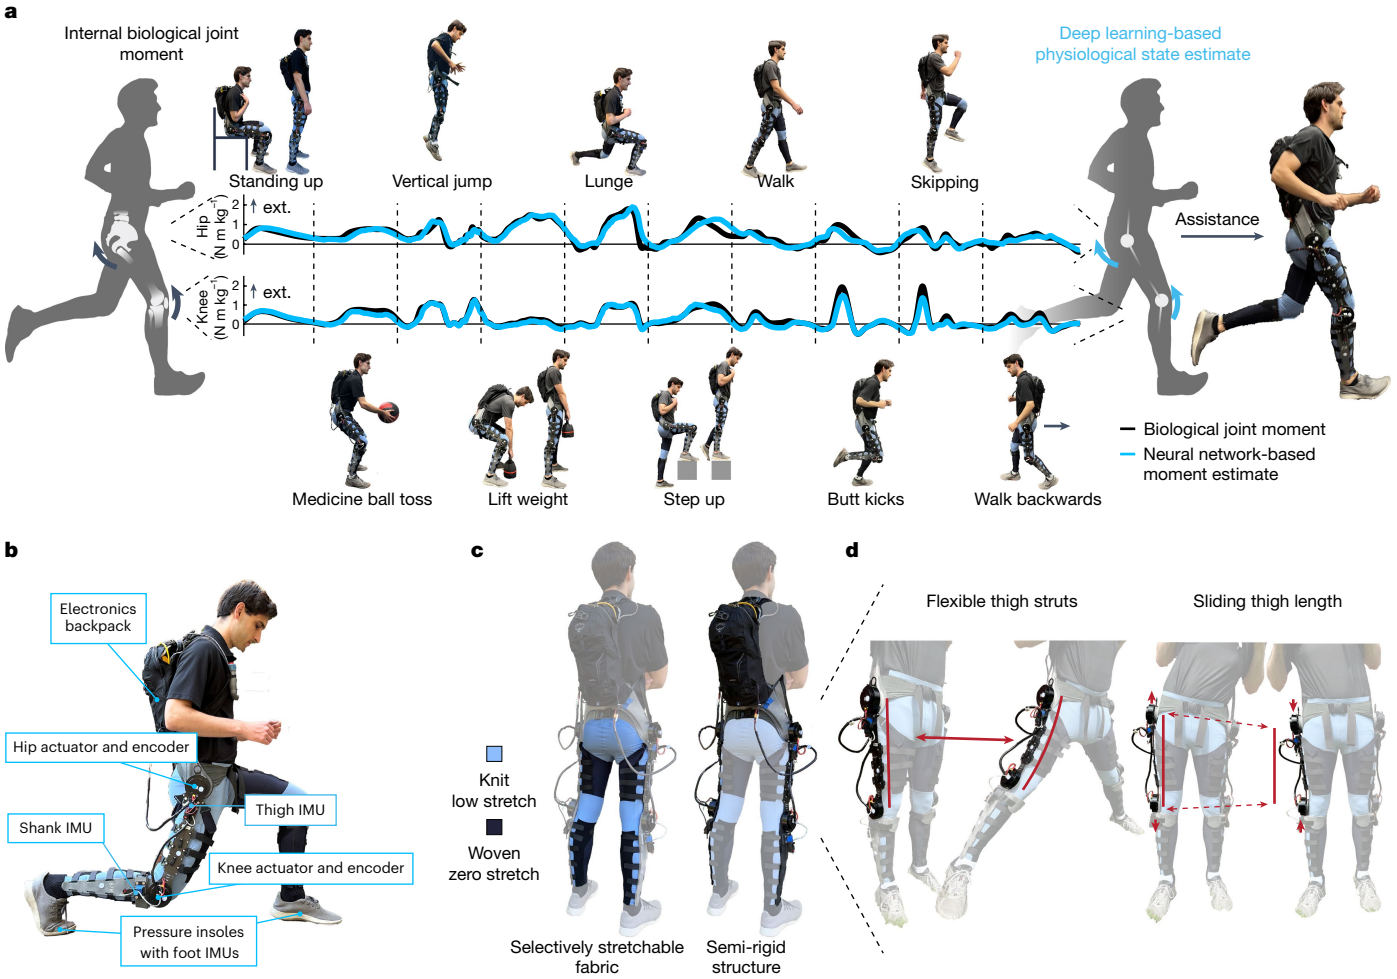

Fig. 1 | Task-agnostic exoskeleton control with a clothing-integrated exoskeleton. a, The proposed approach continuously and seamlessly provides assistive torque to the hip and knee using an estimate of the user's biological joint moments from a deep neural network. By basing assistance off a continuous physiological variable, no task classification is required; the same control law can effectively assist across the full range of human movement. The time series shown illustrates the average performance of our control approach with representative participant-averaged curves on the tasks shown. ext., extension. **b**, An autonomous hip–knee exoskeleton system was constructed to capture a rich set of sensing modalities and then assist across a wide range of mobility tasks. **c**, The hybrid design consists of both soft textiles and semirigid structural components to efficiently transmit exoskeleton assistance to both the hip and knee joints. The human interface consisted of zero-stretch woven fabric to efficiently transmit forces, whereas low-stretch knitted fabrics covering joints helped avoid restrictions in the user's range of motion. **d**, Structural compliance and a passive translational degree of freedom between the hip and knee allowed hip ab- or adduction and rotation while maintaining actuator alignment across movements.

energy shaping methods[35,36] promise joint moment-based control agnostic to a specific lower-limb joint. Further, our previous work[14] presented a deep learning-based approach that significantly reduced user metabolic cost during both level and incline walking using a hip exoskeleton; this approach demonstrated similar or even better outcomes than using previously optimized, spline-based assistance (that is, the previous standard of exoskeleton control) depending on the condition.

Although these past studies demonstrate the efficacy of exoskeleton control using real-time biological joint moment estimates, they do not yet realize the key benefit of this approach: the potential for generalizability across the broad spectrum of human movement. In fact, these previous studies are almost entirely limited to the domain of level walking, ramps and stairs (domains that have been studied in exoskeleton control for decades[15,16]) aside from the sit–stand task investigated by Lin et al.[35]. Further, this approach could autonomously coordinate assistance across many lower-limb joints, a key component of generalizing exoskeleton technology across tasks that depend on different joints; however, significant augmentation of user energetics using a multi-joint exoskeleton controlled by joint moment estimates remains to be demonstrated. In this study, we introduce a task-agnostic controller enabled by a neural network-based joint moment estimator, which runs onboard an autonomous, hip–knee exoskeleton. By training the network on a diverse dataset of time-synced exoskeleton sensor data and ground-truth joint moments, we found that it accurately estimated user joint moments during 28 cyclic and non-cyclic human activities when deployed online (corresponding dataset released with this study). Further, we found that the resulting controller significantly reduced metabolic cost (four activities tested) and lower-limb biological joint work (six activities tested) relative to a no-assistance condition in all tested activities without any manual user or experimenter intervention between activities. This work provides a path to generalizing assistance across human behaviour, a critical link for the adoption of exoskeleton technology in the real world.

Clothing-integrated robotic exoskeleton

We developed a new exoskeleton with the capacity to assist an extensive range of movements, with a particular focus on under-appreciated functional movements that are critical to independence, but can be difficult

for many populations, such as older adults. The clothing-integrated research exoskeleton presented here (Fig. 1b–d), was designed at X, The Moonshot Factory and combines the advantages of rigid exoskeletons with the comfort of soft textiles for the human–exoskeleton interface[16,40] (Fig. 1c). Compact quasi-direct drive actuators (AK80-9 T-Motor, Nanchang) mounted coaxially with the hip and knee provided up to 15 N m of assistance at each joint. The semirigid structure consisted of carbon fibre and 3D printed nylon orthotics on which the actuators and sensors were mounted. Six inertial measurement units (IMUs), joint encoders on the hips and knees, and a pair of wireless force-sensitive insoles provided real-time human movement data for the joint moment estimator (Fig. 1b) with the IMUs being the most important for joint moment estimation (Extended Data Fig. 1). Sagittal-plane actuation is provided at the hip and knee whereas passive degrees of freedom at the hip (translation and rotation) provide flexibility (Fig. 1d). This new exoskeleton architecture gave the user the flexibility and range of motion needed to perform a diverse range of structured and unstructured activities and represents a vital step forward in designing exoskeleton interfaces that are compliant, comfortable and adjustable.

Lower-limb joint moment estimation

To train the joint moment estimator, we collected exoskeleton sensor data time-synced with motion capture and ground reaction forces (GRFs) while users performed a wide range of tasks. Standard OpenSim inverse dynamics (detailed in the Supplementary Information) were used to calculate hip and knee moments[41,42], providing the ground-truth labels (Fig. 2a). To achieve both task generalizability and user-independence, our extensive dataset consisted of 15 healthy participants performing 28 different activities consisting of 66 total conditions (Extended Data Fig. 2). We categorized the 28 activities as cyclic (Supplementary Video 2), impedance-like (Supplementary Video 3) or unstructured (Supplementary Video 4) on the basis of normative joint biomechanics[43] (Fig. 2b). Using this dataset, we trained a temporal convolutional network (TCN) with optimized hyperparameters (see Extended Data Table 1) to estimate hip and knee moments from 19 of the 28 tasks, with the 19 tasks chosen by a forward selection algorithm to promote task generalization within the model (Fig. 2c and Extended Data Fig. 3a). The most helpful data for model generalization (aside from the seed task of level-ground walking whose importance cannot be assessed) was a series of static standing poses, allowing the model to learn the static characteristics of the human body (for example, standing upright requires near zero moment), which is critical for generalization (Extended Data Fig. 3b). Other critical tasks, such as jump and cut, are extremely high-effort tasks that probably helped establish the bounds of the system dynamics and thus are also important for generalization. These previously understudied activities in the exoskeleton domain, many of which are not suitable for gait phase or impedance control, were the most critical for training a model to infer joint moments across real-world tasks.

The joint moment estimates were mapped to applied exoskeleton torque by a continuous transformation (Fig. 2a). Hip and knee moments were scaled to 20 and 15% of the total estimated biological moments, respectively. These scaling factors were established in pilot experiments and provided comfortable assistance while preventing substantial saturation and overheating of the motors during high-torque movements. Hip moment estimates were delayed by 100 ms to maximize positive work done by the exoskeleton[44] and to potentially minimize user metabolic cost[14]. Furthermore, the delay between knee moment estimates and the resulting assistance was set to the minimum achievable by the system (a delay of 50 ms), on the basis of single-blinded pairwise preference tuning[45] during pilot testing. Finally, the delayed joint moment estimates were lowpass filtered to better match the frequency content of human movement[46] and increase user comfort.

The model was validated online with ten participants to assess its ability to accurately estimate human joint moments while providing assistance. No user-specific data were included in training to keep the tests user-independent. Furthermore, we developed a best-case baseline method based on current state-of-the-art exoskeleton control to compare against our joint moment estimator (details in the Supplementary Information); for cyclic activities, the baseline method estimated the user- and stride-averaged hip and knee moments from each activity (for example, for level walking, the baseline used the average level walking curve) and for 'impedance-like' activities (for example, jumping in place), the baseline method estimated the hip and knee moments by estimating zero moment when in swing or flight and by using a linear spring-damper model fit to each activity during stance (Fig. 3a). Unstructured activities were omitted from the baseline because of their lack of phase or impedance-like behaviour, which highlights the limitations of current exoskeleton control. We implemented the baseline method post hoc with perfect gait phase estimates and task classification (that is, a perfectly accurate classifier of 28 classes), thus representing the theoretical best possible performance achievable by this type of control architecture.

Our deep neural network estimated hip and knee moments significantly better than the baseline method for both cyclic (hip R^2 0.79, knee R^2 0.86) and impedance-like activities (hip R^2 0.81, knee R^2 0.87) without any participant-specific calibration (Fig. 3b,c). Representative time series are shown in Fig. 3d–f. Comparing within each activity, our estimator significantly outperformed the baseline method for 12 of the 19 total comparisons of R^2 at the hip and 13 of the comparisons at the knee (Extended Data Fig. 4a,b), with similar results in root mean-square error (r.m.s.e.) (cyclic hip and knee r.m.s.e. 0.15 and 0.13 N m kg^{-1}, impedance-like hip and knee r.m.s.e. 0.21 and 0.16 N m kg^{-1}) and in normalized mean absolute error (MAE) (cyclic hip and knee normalized MAEs 7.3 and 5.5%; impedance-like hip and knee normalized MAEs 7.1 and 6.0%) (Extended Data Figs. 4c–f and 5a–d). The baseline method did not significantly outperform our approach on any individual activity in R^2, r.m.s.e. or normalized MAE. In reality, the high-level state estimators required for the baseline method (that is, a task classifier, gait phase estimator and pose estimator) also have non-zero error[12,20–23,25], further detracting from this approach and highlighting the benefits of our regression-based method. We also found that during unstructured tasks that were not well-defined as cyclic or impedance-like, our approach maintained performance with an average hip R^2 of 0.80 and knee R^2 of 0.82 (Fig. 3b). Thus, our task-agnostic controller mimicked the natural behaviours of human movement, seamlessly modulating assistance throughout the transient motions common in daily life[29].

Given the black box nature of our approach it is possible that the neural network could generate large, erroneous joint moments leading to undesirable exoskeleton assistance. To analyse model under- and overestimation, we computed the normalized hip and knee estimate error at each time instance as the difference between the absolute value of the joint moment estimate and the absolute value of the ground-truth label, normalized by the peak-to-peak range of the ground-truth label. Extended Data Fig. 5e depicts the distribution of the normalized hip and knee error from the online validation trials (representing roughly 10 million instances total), in which negative and positive values correspond to under- and overestimates, respectively. As shown in the figure, large under- and overestimates were uncommon with means close to zero (hip mean −2.7%; knee mean −1.3%) and standard deviations of 7.8 and 6.6% for the hip and knee, respectively. Furthermore, time series examples of the most severe instances of under- and overestimation from the joint moment estimator are shown in Extended Data Fig. 5f.

Of the 28 evaluated activities, nine were withheld from the training set (details in Methods). The average R^2 of our estimator on these held-out tasks was 0.83 and 0.85 for the hip and knee, respectively, demonstrating the ability of the network to generalize to the hold-out tasks. To further investigate estimator generalization, three users also completed

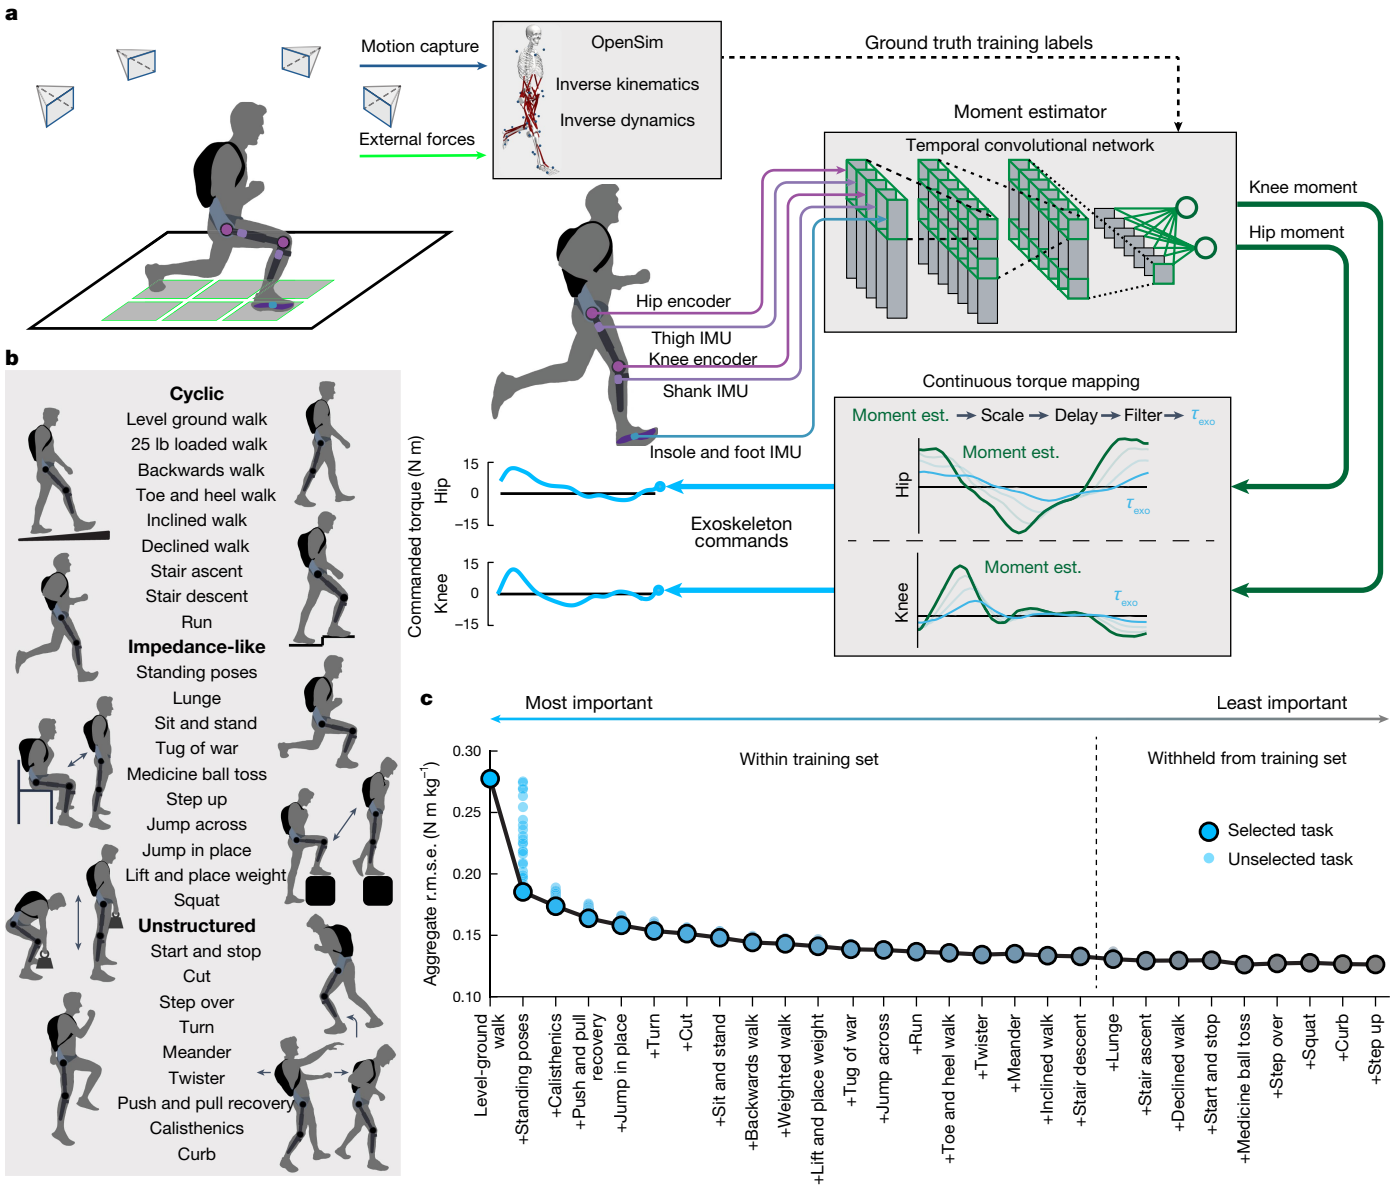

Fig. 2 | Deep neural network training and deployment for joint moment estimation. a, Lower-limb joint moment labels were calculated in OpenSim using optical motion capture, force plate data and user-specific musculoskeletal models; a TCN was trained to predict these joint moment labels from time-synchronized exoskeleton sensor data. During deployment, to improve power delivery and user comfort, the estimates were transformed into commanded exoskeleton torque through a continuous function consisting of a scale, delay and a lowpass filter. est., estimate. **b**, Users wore the exoskeleton while performing a wide range of cyclic, impedance-like and unstructured tasks. **c**, Training activities for the moment estimator were selected using a forward selection algorithm to maximize the relative improvement in model generalization across tasks. Validation r.m.s.e. decreased as the training set grew, with the first 19 tasks reducing r.m.s.e. to 0.133 N m kg^{-1}, which was within 5% of the best model accuracy with all the tasks included. This task set was used to train the real-time models used in the rest of this study. For reference, peak-to-peak hip and knee moments ranged from 2 to 4 N m kg^{-1} for most activities in the dataset. Results were computed from leave-one-participant-out cross-fold validation using a 12-participant dataset (error bars omitted for clarity).

eight completely new tasks, described in Extended Data Table 2, that had not been previously tested or analysed (Supplementary Fig. 1 and Supplementary Video 5). These tasks were intentionally designed to be highly unique from the original dataset to push the limits of our approach, including burpees, mimicking a basketball layup and walking on a split belt treadmill with differing belt speeds. Our approach generalized well to the tasks reflective of typical human movement, and when pushed to extremely dynamic behaviours outside of the training set, our approach provided directionally correct assistance, but the magnitude and shape lost accuracy (R^2 ranged from 0.24 to 0.92 at the hip and from 0.32 to 0.91 at the knee for the eight new tasks; Supplementary Fig. 1c,d). These results demonstrate the ability of the estimator to generalize to never-before-seen activities but highlights that task-specific training data is beneficial for activities with highly different dynamics (for example, when offloading bodyweight through the hands on the ground). Extra details and discussion comparing our joint moment estimator relative to previous methods and regarding its performance during new tasks are provided in the Supplementary Information.

Augmenting user energetics across tasks

To quantify the impact of our task-agnostic controller on the user, we measured user metabolic cost during four activities under three assistance conditions: wearing the exoskeleton with our task-agnostic controller (exo on), without wearing the exoskeleton (no exo) and wearing

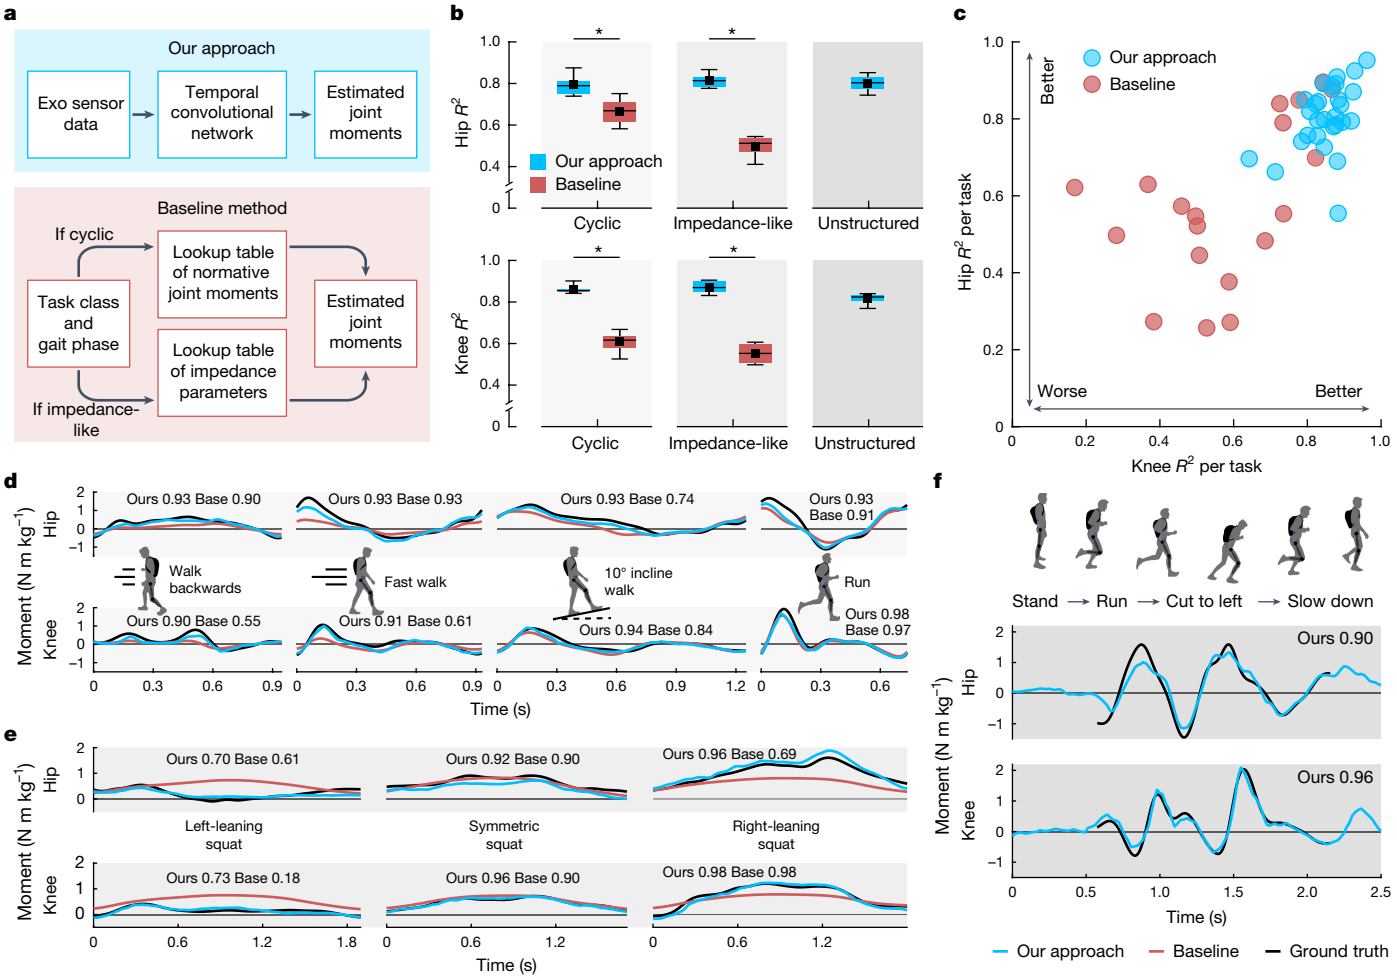

Fig. 3 | Online joint moment estimation performance. a, We compared our neural network-based joint moment estimator (deployed online) to a best-case baseline method (computed offline) that relied on perfect task classification and gait phase. **b**, Our approach significantly improved R^2 at the hip by 0.13 ± 0.04 ($19 \pm 6\%$, $P = 2 \times 10^{-7}$) and at the knee by 0.25 ± 0.04 ($40 \pm 7\%$, $P < 10^{-16}$) compared to the baseline method during cyclic activities. For impedance-like tasks, our approach improved R^2 by 0.31 ± 0.05 ($63 \pm 10\%$, $P = 4 \times 10^{-16}$) at the hip and by 0.32 ± 0.02 ($57 \pm 3\%$, $P < 10^{-16}$) at the knee compared to the baseline method. Black squares depict inter-participant mean, coloured boxes depict interquartile range, horizontal lines within boxes depict inter-participant median and error bars depict inter-participant minimum and maximum ($n = 10$). **c**, Estimator R^2 is shown per task for our approach and the baseline method. Each marker corresponds to the inter-participant average per single task ($n = 10$, except for the run condition where $n = 9$). **d**, Representative strides from various cyclic tasks are shown. The baseline method required a different task classification for each depicted ambulation mode, whereas our approach did not require any discrete switching. **e**, Representative trials are shown when squatting to the left, right and symmetrically. The impedance control-based approach failed to capture changes in joint moments by relying solely on kinematics. Instead, our approach accurately modified joint moments with the change in weight distribution across the user's legs. **f**, A representative trial during leftward cutting is shown, depicting the ability of our approach to seamlessly modulate assistance during highly unstructured behaviours. As it is unclear how to extend the baseline method to these types of activity, it was omitted. Estimator R^2 relative to ground truth is shown for our approach (ours) and the baseline method (base) above each representative trial.

the exoskeleton without assistance (zero torque). The task-agnostic controller significantly reduced metabolic cost for all four tasks compared to zero torque ($P < 0.05$) with relative reductions ranging from 8.0% during the lift weight task to 19.7% during 5° inclined walking (Fig. 4a). Relative to no exo, our approach significantly reduced user metabolic cost during the weight lifting task and during running ($P < 0.05$); however, our approach increased metabolic cost during level-ground walking ($P < 0.05$). Given the similarity in level walking estimator accuracy and metabolic cost reduction relative to zero torque in our previous work, which did reduce metabolic cost relative to no exo using a lighter weight hip-only exoskeleton[14], it is likely that the increase in metabolic cost in this study was due to the added mass penalty of the hip–knee exoskeleton, not the controller itself[47]. Nevertheless, these results demonstrate the ability of our approach to autonomously modulate assistance across tasks in a beneficial manner, a critical hurdle in developing task-agnostic exoskeleton controllers.

To further quantify the effect of the device during transient tasks, we measured metabolic cost for three participants performing a varying speed and incline circuit that ranged from walking to running with inclines ranging from 0° to 15°. During this highly transient trial, our approach reduced user metabolic cost by 12.9% relative to zero torque and by 1.6% relative to no exo (Fig. 4b). Thus, our approach seamlessly accommodated these transient behaviours so common to daily life[29] without any extra tuning or calibration.

Furthermore, positive lower-limb biological joint work of the user was also evaluated during six extra tasks described in Extended Data Table 3, which provided insight into the joint-level effects of our controller[14,48–50] and was less taxing on the participants than the metabolic trials. Total positive lower-limb biological joint work was computed as the sum of the components computed at the hip, knee and ankle. Positive biological joint work was computed by integrating the positive biological power at each joint, which involved subtracting the

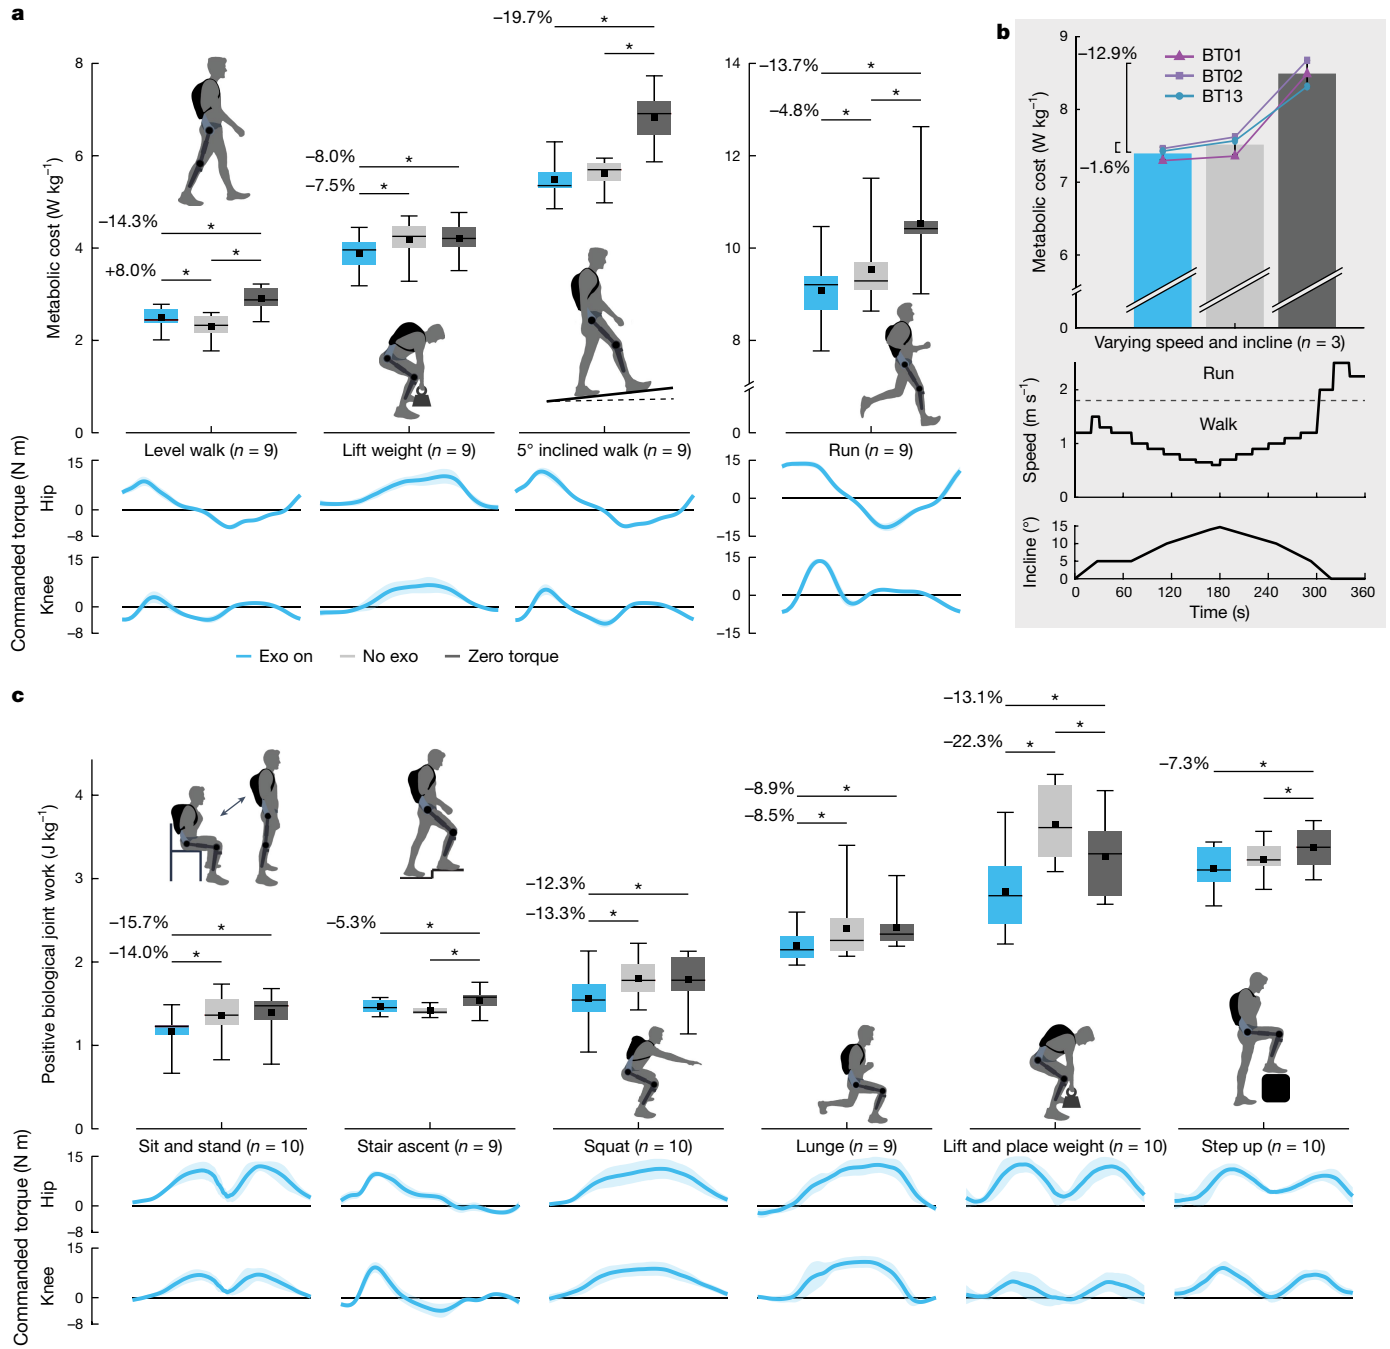

Fig. 4 | Human outcome performance. a, Participant metabolic cost was measured during four activities while wearing the exoskeleton using the task-agnostic controller (exo on), without wearing the exoskeleton (no exo) and while wearing the exoskeleton without assistance (zero torque). Tasks other than running were conducted using 6 min trials and a counter-balanced design (ABCCBA). Owing to the strenuous nature of the running trials, conditions were only completed once (ABC) and each trial only lasted 3.5 min. Below each activity is the cycle-averaged commanded torque during exo on as a function of movement percentage. The shaded region around each curve depicts ±1 standard deviation about the mean. **b**, Three participants returned and performed a varying speed and incline protocol while metabolic cost was measured. Individual traces are provided for each participant as well as the average. The detailed variations in speed and incline are shown below the activity. **c**, The average positive biological joint work per movement cycle summed across each participant's hip, knee and ankle are shown. Participants completed the six activities under the same three assistance conditions as the metabolic trials. Again, the cycle-average commanded torque during exo on for each activity is shown as a function of movement percentage whereas the shaded region around each curve depicts ±1 standard deviation about the mean. Each black square depicts the inter-participant mean, each coloured box depicts the interquartile range, each horizontal line within the boxes depicts the inter-participant median and each error bar depicts the inter-participant minimum and maximum. Asterisks indicate statistical significance ($P < 0.05$; exact P values are provided in Supplementary Data 1).

exoskeleton torque from the ground-truth total joint moment from inverse dynamics (further details are provided in the Methods section). Our controller significantly reduced positive lower-limb biological joint work of the user during all six tasks compared to zero torque ($P < 0.05$) with decreases ranging from 5.3 to 15.7% (Fig. 4c). During four of the six tasks, our controller also reduced positive joint work compared to no exo with significant decreases ranging from 8.5 to 22.3%. The other two tasks, stair ascent and step up, showed no significant difference.

Furthermore, participant lower-limb kinematics showed little variation across exoskeleton conditions (Supplementary Fig. 2), suggesting that reductions in biological joint work were achieved primarily through reductions in biological joint moments as opposed to modified kinematics (details in the section 'Exoskeleton effects on user kinematics' in the Supplementary Information).

With further improvements to the exoskeleton, such as reducing exoskeleton mass, relocating knee actuators more proximal to the body[47], and increasing the maximum magnitude of assistance, we anticipate substantially greater capacity for our approach to enhance human performance across activities compared to this first evaluation. Furthermore, it is likely that our mid-level controller was not optimal for all tasks and similarly may not be optimal for populations outside of young, able-bodied individuals. Further optimization of the mid-level controller[3,8,11], consideration of how this approach could generalize to extra populations and investigation of the physiological mechanisms that drive the relationship between exoskeleton assistance and user outcomes could result in further improvements in user outcomes and expand the scope of this approach. Nevertheless, these comparisons demonstrate the efficacy of our task-agnostic controller to dynamically and beneficially modulate assistance with changing user behaviour without user- or task-specific tuning, which is a critical component for exoskeleton controllers deployed in the real world.

Conclusion

By relying on internal physiological state estimates rather than human-engineered gait parameterizations, control is given back to the user and the exoskeleton can respond to that user's specific, real-time joint moments without any user-specific calibration or hand-engineered state machine criteria. Whereas recent work has demonstrated how to leverage instantaneous joint moment estimates as a promising alternative for exoskeleton control during walking[13,14,35,36], here we present the missing piece: task generalization. Specifically, our deep neural network approach to exoskeleton control provides a task-agnostic framework, capable of seamlessly augmenting human effort by coordinating exoskeleton assistance across joints during both structured and unstructured tasks. Building from these advances, we believe this technology will become paramount in many industries, allowing researchers to test new hypotheses during natural, less-structured behaviours, relieving worker exhaustion in manual labour industries and enabling disaster relief teams to act with extra endurance in time-sensitive missions.

Online content

Any methods, additional references, Nature Portfolio reporting summaries, source data, extended data, supplementary information, acknowledgements, peer review information; details of author contributions and competing interests; and statements of data and code availability are available at https://doi.org/10.1038/s41586-024-08157-7.

1. Kim, J. et al. Reducing the metabolic rate of walking and running with a versatile, portable exosuit. *Science* **365**, 668–672 (2019).
2. Witte, K. A., Fiers, P., Sheets-Singer, A. L. & Collins, S. H. Improving the energy economy of human running with powered and unpowered ankle exoskeleton assistance. *Sci. Robot.* **5**, eaay9108 (2020).
3. Slade, P., Kochenderfer, M. J., Delp, S. L. & Collins, S. H. Personalizing exoskeleton assistance while walking in the real world. *Nature* **610**, 277–282 (2022).
4. Awad, L. N. et al. A soft robotic exosuit improves walking in patients after stroke. *Sci. Transl. Med.* **9**, eaai9084 (2017).
5. Malcolm, P., Derave, W., Galle, S. & Clercq, D. D. A simple exoskeleton that assists plantarflexion can reduce the metabolic cost of human walking. *PLoS ONE* **8**, e56137 (2013).
6. Mooney, L. M., Rouse, E. J. & Herr, H. M. Autonomous exoskeleton reduces metabolic cost of human walking during load carriage. *J. NeuroEngineering Rehabil.* **11**, 80 (2014).
7. Ishmael, M. K., Archangeli, D. & Lenzi, T. Powered hip exoskeleton improves walking economy in individuals with above-knee amputation. *Nat. Med.* **27**, 1783–1788 (2021).
8. Franks, P. W. et al. Comparing optimized exoskeleton assistance of the hip, knee, and ankle in single and multi-joint configurations. *Wearable Technol.* **2**, e16 (2021).
9. Baltrusch, S. J. et al. The effect of a passive trunk exoskeleton on metabolic costs during lifting and walking. *Ergonomics* **62**, 903–916 (2019).
10. Lim, B. et al. Delayed output feedback control for gait assistance with a robotic hip exoskeleton. *IEEE Trans. Robot.* **35**, 1055–1062 (2019).
11. Zhang, J. et al. Human-in-the-loop optimization of exoskeleton assistance during walking. *Science* **356**, 1280–1284 (2017).
12. Shepherd, M. K., Molinaro, D. D., Sawicki, G. S. & Young, A. J. Deep learning enables exoboot control to augment variable-speed walking. *IEEE Robot. Autom. Lett.* **7**, 3571–3577 (2022).
13. Gasparri, G. M., Luque, J. & Lerner, Z. F. Proportional joint-moment control for instantaneously adaptive ankle exoskeleton assistance. *IEEE Trans. Neural Syst. Rehabil. Eng.* **27**, 751–759 (2019).
14. Molinaro, D. D., Kang, I. & Young, A. J. Estimating human joint moments unifies exoskeleton control, reducing user effort. *Sci. Robot.* **9**, eadi8852 (2024).
15. Siviy, C. et al. Opportunities and challenges in the development of exoskeletons for locomotor assistance. *Nat. Biomed. Eng.* **7**, 456–472 (2023).
16. Sawicki, G. S., Beck, O. N., Kang, I. & Young, A. J. The exoskeleton expansion: improving walking and running economy. *J. NeuroEngineering Rehabil.* **17**, 25 (2020).
17. Collins, S. H., Wiggin, M. B. & Sawicki, G. S. Reducing the energy cost of human walking using an unpowered exoskeleton. *Nature* **522**, 212–215 (2015).
18. Yang, J., Park, J., Kim, J., Park, S. & Lee, G. Reducing the energy cost of running using a lightweight, low-profile elastic exosuit. *J. NeuroEngineering Rehabil.* **18**, 129 (2021).
19. Li, Y. D. & Hsiao-Wecksler, E. T. Gait mode recognition and control for a portable-powered ankle-foot orthosis. In *Proc. 2013 IEEE 13th International Conference on Rehabilitation Robotics (ICORR)* 1–8 (IEEE, 2013).
20. Laschowski, B., McNally, W., Wong, A. & McPhee, J. Environment classification for robotic leg prostheses and exoskeletons using deep convolutional neural networks. *Front. Neurorobotics* **15**, 730965 (2022).
21. Kang, I., Molinaro, D. D., Choi, G., Camargo, J. & Young, A. J. Subject-independent continuous locomotion mode classification for robotic hip exoskeleton applications. *IEEE Trans. Biomed. Eng.* **69**, 3234–3242 (2022).
22. Camargo, J., Flanagan, W., Csomay-Shanklin, N., Kanwar, B. & Young, A. A machine learning strategy for locomotion classification and parameter estimation using fusion of wearable sensors. *IEEE Trans. Biomed. Eng.* **68**, 1569–1578 (2021).
23. Qian, Y. et al. Predictive locomotion mode recognition and accurate gait phase estimation for hip exoskeleton on various terrains. *IEEE Robot. Autom. Lett.* **7**, 6439–6446 (2022).
24. Medrano, R. L., Thomas, G. C., Keais, C. G., Rouse, E. J. & Gregg, R. D. Real-time gait phase and task estimation for controlling a powered ankle exoskeleton on extremely uneven terrain. *IEEE Trans. Robot.* **39**, 2170–2182 (2023).
25. Kang, I. et al. Real-time gait phase estimation for robotic hip exoskeleton control during multimodal locomotion. *IEEE Robot. Autom. Lett.* **6**, 3491–3497 (2021).
26. Huo, W. et al. Impedance modulation control of a lower-limb exoskeleton to assist sit-to-stand movements. *IEEE Trans. Robot.* **38**, 1230–1249 (2022).
27. Yang, X. et al. Spine-inspired continuum soft exoskeleton for stoop lifting assistance. *IEEE Robot. Autom. Lett.* **4**, 4547–4554 (2019).
28. Shepherd, M. K. & Rouse, E. J. Design and validation of a torque-controllable knee exoskeleton for sit-to-stand assistance. *IEEE ASME Trans. Mechatron.* **22**, 1695–1704 (2017).
29. Orendurff, M. S., Schoen, J. A., Bernatz, G. C., Segal, A. D. & Klute, G. K. How humans walk: bout duration, steps per bout, and rest duration. *J. Rehabil. Res. Dev.* **45**, 1077–1089 (2008).
30. Winter, D. in *Biomechanics and Motor Control of Human Movement* Ch. 5, 107–138 (John Wiley & Sons, Ltd, 2009).
31. Dorschky, E. et al. CNN-based estimation of sagittal plane walking and running biomechanics from measured and simulated inertial sensor data. *Front. Bioeng. Biotechnol.* **8**, 604 (2020).
32. Molinaro, D. D., Kang, I., Camargo, J., Gombolay, M. C. & Young, A. J. Subject-independent, biological hip moment estimation during multimodal overground ambulation using deep learning. *IEEE Trans. Med. Robot. Bionics* **4**, 219–229 (2022).
33. Camargo, J., Molinaro, D. & Young, A. Predicting biological joint moment during multiple ambulation tasks. *J. Biomech.* **134**, 111020 (2022).
34. Hossain, M. S. B., Guo, Z. & Choi, H. Estimation of lower extremity joint moments and 3D ground reaction forces using IMU sensors in multiple walking conditions: a deep learning approach. *IEEE J. Biomed. Health Inform.* https://doi.org/10.1109/JBHI.2023.3262164 (2023).
35. Lin, J., Divekar, N. V., Thomas, G. C. & Gregg, R. D. Optimally biomimetic passivity-based control of a lower-limb exoskeleton over the primary activities of daily life. *IEEE Open J. Control Syst.* **1**, 15–28 (2022).
36. Zhang, J., Lin, J., Peddinti, V. & Gregg, R. D. Optimal energy shaping control for a backdrivable hip exoskeleton. In *Proc. 2023 American Control Conference (ACC)* 2065–2070 (IEEE, 2023).
37. Fang, Y., Orekhov, G. & Lerner, Z. F. Improving the energy cost of incline walking and stair ascent with ankle exoskeleton assistance in cerebral palsy. *IEEE Trans. Biomed. Eng.* **69**, 2143–2152 (2022).
38. Bishe, S. S. P. A., Nguyen, T., Fang, Y. & Lerner, Z. F. Adaptive ankle exoskeleton control: validation across diverse walking conditions. *IEEE Trans. Med. Robot. Bionics* **3**, 801–812 (2021).
39. Tagoe, E. A., Fang, Y., Williams, J. R. & Lerner, Z. F. Walking on real-world terrain with an ankle exoskeleton in cerebral palsy. *IEEE Trans. Med. Robot. Bionics* **6**, 202–212 (2024).
40. Young, A. J. & Ferris, D. P. State of the art and future directions for lower limb robotic exoskeletons. *IEEE Trans. Neural Syst. Rehabil. Eng.* **25**, 171–182 (2017).
41. Delp, S. L. et al. OpenSim: open-source software to create and analyze dynamic simulations of movement. *IEEE Trans. Biomed. Eng.* **54**, 1940–1950 (2007).

42. Seth, A. et al. OpenSim: simulating musculoskeletal dynamics and neuromuscular control to study human and animal movement. *PLoS Comput. Biol.* **14**, e1006223 (2018).
43. Scherpereel, K., Molinaro, D., Inan, O., Shepherd, M. & Young, A. A human lower-limb biomechanics and wearable sensors dataset during cyclic and non-cyclic activities. *Sci. Data* **10**, 924 (2023).
44. Ding, Y. et al. Effect of timing of hip extension assistance during loaded walking with a soft exosuit. *J. NeuroEngineering Rehabil.* **13**, 87 (2016).
45. Ingraham, K. A., Tucker, M., Ames, A. D., Rouse, E. J. & Shepherd, M. K. Leveraging user preference in the design and evaluation of lower-limb exoskeletons and prostheses. *Curr. Opin. Biomed. Eng.* **28**, 100487 (2023).
46. Winter, D. A. Biomechanical motor patterns in normal walking. *J. Mot. Behav.* **15**, 302–330 (1983).
47. Browning, R. C., Modica, J. R., Kram, R. & Goswami, A. The effects of adding mass to the legs on the energetics and biomechanics of walking. *Med. Sci. Sports Exerc.* **39**, 515 (2007).
48. Farris, D. J. & Sawicki, G. S. The mechanics and energetics of human walking and running: a joint level perspective. *J. R. Soc. Interface* **9**, 110–118 (2012).
49. Farris, D. J., Hampton, A., Lewek, M. D. & Sawicki, G. S. Revisiting the mechanics and energetics of walking in individuals with chronic hemiparesis following stroke: from individual limbs to lower limb joints. *J. NeuroEngineering Rehabil.* **12**, 24 (2015).
50. Farris, D. J. & Sawicki, G. S. Linking the mechanics and energetics of hopping with elastic ankle exoskeletons. *J. Appl. Physiol.* **113**, 1862–1872 (2012).

Methods

Autonomous robotic hip–knee exoskeleton

In this study, we used the clothing-integrated robotic exoskeleton developed at X, which was designed to enhance human mobility by providing powered sagittal-plane assistance to the hips and knees (design rationale provided in the supplementary section 'Why the hip and knee and not the ankle?'). A Raspberry Pi 4B (RPi) (Raspberry Pi) served as the primary onboard computer that ran the exoskeleton control loop at 55 Hz. The RPi managed Controller Area Network (CAN) Bus and Bluetooth communication with peripherals, saved experimental data locally and provided all functions other than joint moment estimates. Joint moment estimates were generated on a machine learning coprocessor (NVIDIA Jetson Nano) also mounted onboard the device; thus, all computation was fully onboard the exoskeleton. The Jetson Nano provided a low power consumption (5 V, 2 A) coprocessor, easily integrated into the exoskeleton by means of ethernet connection and commonly available portable charging banks. Actuated hip flexion and extension and knee flexion and extension was provided by quasi-direct drive actuators (T-Motor AK80-9s, Nanchang), with a peak intermittent torque constrained to 15 N m. Open-loop torque commands were sent to the actuators from the RPi over CAN and encoder measurements were returned to the RPi. Encoder velocity was lowpass filtered using a second-order Butterworth filter with a 10 Hz cut-off frequency. Six-axis IMUs (OpenIMUA) were mounted to the shank and thigh struts and communicated with the RPi by means of CAN. Pressure-sensitive insoles (Moticon) measured vertical GRF and centre of pressure (COP) and had an embedded six-axis IMU. They communicated with the RPi by Bluetooth and were powered by coin-cell batteries. The RPi and actuators were powered by two 20 V, 3 Ah drill batteries (DeWalt) connected in parallel, providing power for roughly 2 h of continuous walking at a minimum. The RPi interfaced through WiFi with a laptop for data visualization using a custom user interface.

The human–robot interface was designed to be as compliant, comfortable and adjustable as possible, while maintaining the minimal structure required for effective transfer of actuator torques to the body. The semirigid structure included thin, waterjet-cut carbon fibre plates in the thigh struts that were compliant in ab- or adduction and internal or external rotation of the hip, but supported flexion and extension torques. A passive translational degree of freedom allowed the shortening of the thigh strut required for hip ab- or adduction. The 3D printed nylon shin struts, shin cuffs and pelvic orthosis were designed to apply the assistance to comfortable sections on the shin and pelvis. The custom trousers provided tight integration of the semirigid structure with the body, preventing buckling of the thin thigh struts. Quick release snaps, adjustable with Velcro, connected the pelvic orthosis and thigh or shank struts to the trousers. Woven zero-stretch fabric on the thigh and shank allowed slop-free transfer of the applied exoskeleton torque to the leg through a 'hammocking' effect. Knit fabric, with some stretch and elasticity, around the knee and hip allowed free range of motion and was used on the lateral sides of the trousers to allow variation in sizing between users. Zippers on the shank allowed tight integration of the fabric with the lower shank while allowing don and doff over the heel. Overall, the entire exoskeleton added roughly 7 kg to the user depending on the exoskeleton size worn during the experiment. Furthermore, although comfort was a key consideration of the exoskeleton design, some participants voiced discomfort regarding the shape of the backplate and the load-bearing seams of the soft textile trousers.

Real-time joint moment estimation

To deploy the joint moment estimator within the exoskeleton controller, we integrated the machine learning coprocessor (that is, the Jetson Nano) into the device using an asynchronous TCP/IP connection over wired ethernet with the RPi. With each control loop, sensor data were measured from the actuators, IMUs and pressure insoles. The sensor data were sent from the RPi to the coprocessor, which returned estimates of the total hip and knee flexion or extension moments from the neural network. Biological joint moment estimates were then computed by subtracting the measured actuator torques from the previous loop from each corresponding joint moment estimate.

Desired torque assistance was computed from the resulting biological joint moment estimates using three steps. (1) The biological joint moments at the hip and knee were scaled by 20 and 15%, respectively, to maximize assistance while maintaining safe operating regions for the device hardware (15 N m at each joint). (2) The scaled hip and knee moments were then delayed by 100 and 50 ms, respectively. This delay was chosen at the knee because 50 ms was the minimum possible delay to guarantee a consistent relationship between biological joint moment estimates and exoskeleton assistance owing to limitations in loop rate reliability of the exoskeleton. Hip assistance was further delayed by an extra 50 ms because this approach can maximize the positive work done by the exoskeleton during walking, which can lead to further benefits for the user[14,44]. (3) The exoskeleton assistance was lowpass filtered using a second-order 10 Hz Butterworth filter to preserve the frequency content of human motion[43,51–53] while removing jitter from the estimator[14]. This filter added an extra delay of 25 ms, resulting in a total hip delay of 125 ms and knee delay of 75 ms. In the case in which the resulting assistance was larger than the peak exoskeleton torque, commanded torque was clamped to the peak exoskeleton torque.

Because the exoskeleton controller intentionally delayed the hip moment estimates by an extra 50 ms relative to the minimum achievable system delay with our system, we chose to train the network to estimate hip moments delayed by 50 ms relative to the input sequence. Our previous work found that delaying joint moment estimates relative to the input sequence can further improve model accuracy[54], and we found that this approach resulted in an extra 5% improvement in hip moment validation MAE for this study.

Neural network architecture optimization

In our previous work, we achieved state-of-the-art accuracy in estimating user joint moments with a TCN[32,54]. In this study, we implemented the TCN as originally introduced by Bai et al.[55] and modified for joint moment estimation in our previous work[14,32,54]. The TCN input consisted of a sequence of unilateral hip and knee encoder data, thigh, shank and foot IMU data, and pressure insole data (vertical GRF and COP). Owing to the exoskeleton loop rate of 55 Hz, the data were upsampled in real-time to 200 Hz to match the sampling frequency previously used for the TCN. Each of the model inputs were also normalized using their corresponding mean and standard deviation computed from the training set. The TCN was designed with two output heads for the instantaneous estimates of the total hip and knee flexion or extension moments (that is, the sum of exoskeleton torque and human biological moment). Joint moment labels in the training set were scaled by participant body mass during training, such that the model was trained to estimate joint moments in units of N m kg^{-1} (refs. 32,54). Furthermore, we trained the TCN to estimate the total joint moments to maintain the relationship between exoskeleton sensor data and TCN joint moment outputs, regardless of the specific parameters of the exoskeleton controller, such as assistance magnitude. We then computed biological joint moments later in the control framework by subtracting the exoskeleton torque from the total estimated moment.

In our previous work, we conducted a thorough hyperparameter optimization of the TCN for estimating sagittal-plane hip moments[32], however, this approach did not consider model generalizability (that is, during the hyperparameter optimization the model training set and validation set consisted of the same ambulation modes). Furthermore, this optimization was conducted under different conditions (only sagittal-plane hip moments, cyclic ambulatory activities and kinematic sensors). Thus, we conducted a rigorous hyperparameter optimization using a multi-stage approach, specifically targeting model

generalizability under conditions consistent with this study (Extended Data Table 1). In stage 1 of the optimization, we used our previous dataset of human lower-limb biomechanics during cyclic and non-cyclic activities[43] to optimize a large, 11-dimensional hyperparameter space using Bayesian optimization implemented in Vizier[56] (training and testing more than 10,000 models). In stage 2, we finetuned the network hyperparameters using actual exoskeleton sensor data from phases 1 and 2 of the experimental protocol (details below) over a smaller, six-dimensional search space that could be rigorously optimized using grid search (training and testing 1,440 models). The six-dimensional space was constructed to finetune the most sensitive network hyperparameters determined from the marginal and conditional results of the stage 1 optimization. The resulting network hyperparameters (shown in Extended Data Table 1) resulted in an 8% improvement in MAE using leave-one-participant-out cross-fold validation compared to using the original hyperparameters from Molinaro et al.[32].

Task optimization for generalizability

Because the collection of actuated, motion capture-labelled data is difficult and costly, we first sought to discover a subset of tasks that could allow a user-independent lower-limb joint moment estimator to generalize to the rest of human activities. To determine the subset of training activities that best promoted generalization, we used the same dataset used for stage 1 of the hyperparameter optimization to conduct a forward activity selection optimization (Extended Data Fig. 3a). During each optimization step, the TCN was trained and tested using leave-one-participant-out validation to compute the expected model performance when evaluated on a new participant. First, model performance was computed by training the model using only the level-ground walking data but tested on all activities. The TCN was then iteratively trained from random initialization, including one activity from the candidate task set into the training set at a time. On the ith optimization step, the relative improvement in generalizability $s_g[i]$ associated with including a candidate activity (g) into the training set (that is, the relative improvement in model performance on all activities beside g) was computed as

$$s_g[i] = \frac{1}{n-1}\sum_{j=0}^{n-1} e_{g^*,j}[i-1] - e_{g,j}[i], \quad j \neq g, \tag{1}$$

$$g^*[i] = \underset{g}{\operatorname{argmax}}\, s_g[i] \quad \forall\, g \in G[i], \tag{2}$$

where n is the total number of tasks and $e_{g,j}[i]$ is the MAE of estimating joint moments during the jth task when trained using the updated training set over the first $i-1$ optimization steps and the extra candidate task g. Thus, $s_g[i]$ evaluated the overall improvement in model performance across all activities, excluding the relative improvement of the task at hand. Furthermore, $g^*[i]$ was the activity associated with the largest improvement score, which was then added to the training set for all further optimization steps and removed from the set of candidate tasks G to be selected in the next optimization step. This process was repeated until all tasks were selected.

We defined the activity set that saturated model generalization as the minimum set of selected activities that contributed more than 95% of the sum total of relative improvement in generalizability across the complete optimization. After the first seven tasks, the model had reached this threshold, indicating that further tasks failed to substantially improve the model's ability to estimate joint moments on other tasks (Extended Data Fig. 3b). Aside from level-ground walking, which was used to seed the optimizer, none of the selected tasks was cyclic and four were unstructured.

Although model generalization saturated rapidly, task-specific data continued to improve user-independent, task-specific validation error (Fig. 2c). The overall model performance required 19 selected activities before validation MAE fell within 5% of the validation MAE when trained on all activities. This demonstrates that there was added benefit to be gained from training on task-specific data even after generalization saturated. For all further analyses, the TCN was trained using the data from these 19 selected tasks unless otherwise stated.

Sensor contribution to moment estimation

Previous joint moment estimation studies often use data from IMUs and joint encoders as model inputs[14,32,34,54,57,58]; however, the relative importance of each of these sensors (and others) on model performance is less explored. Furthermore, the sensitivity of model performance relative to sensor dropout (for example, from sensor disconnection) is also a critical real-world consideration. To investigate these two topics, we trained several extra models under two different conditions: first, using different subsets of available sensors and second, simulating sensor disconnection during model deployment. We tested the performance of these models offline on the data from the ten participants used to test our approach online. For the first condition, two of the sensor sets were inspired by common exoskeleton design choices: (1) removing the GRF and COP contributed by the insoles (that is, -insoles) and (2) removing all foot mounted sensing (that is, -insoles, -foot IMU). The other three sensor sets were chosen to demonstrate the contribution from each unique sensor modality: (1) only using the thigh and shank IMUs (IMU only), (2) only using the hip and knee encoders (encoder only) and (3) only using the GRF and COP from the foot insole (insole only). For the second condition, a single sensor was effectively dropped out during model testing, by zeroing that respective sensor's inputs to simulate a sensor losing connection during device deployment. These comparisons are presented in Extended Data Fig. 1.

In the first condition, removing the GRF and COP resulted in a 0.03 and a 0.07 reduction in R^2 at the hip and knee, respectively, demonstrating the moderate benefit of adding kinetic based sensing, with an additional penalty when removing the foot IMU. The 'IMU only' condition resulted in a further drop in performance; the benefit gained from the encoders indicated that our six-axis IMUs did not fully capture the relevant kinematic input information. Overall, the IMUs contributed the most to the accuracy of the model, followed by the encoders and, last, the insole. Previous work has shown that kinematic sensors can be effectively used to estimate GRFs, indicating a potential reason the IMUs contributed the most to the model[34,59]. Likewise, because the insoles initially measure pressure, from which the GRF is calculated, the amount of information provided to the model may be less than other more accurate sensors, such as in-ground force plates.

In the second condition, dropping out different sensors showed the reliance of our network on each sensor. In general, loss of a sensor resulted in a significant drop in accuracy, indicating that the model generally used all available sensing inputs in the neural network weights. The loss of the knee encoder, however, was much worse for estimating knee joint moments and, similarly, the loss of the hip encoder was much worse for estimating hip joint moments. The losses of the thigh IMU and foot insole were likewise much more consequential to model accuracy when compared to the loss of the shank or foot IMU. These results indicate that our trained network learned to rely on each sensor for generating joint moment estimates, underscoring the importance of the need for high-quality hardware and sensor integration; however, further analyses could explore training the network with synthesized sensor dropout to potentially improve robustness.

Experimental data collection

The data used in this study were collected over three phases to facilitate model training and online testing of the joint moment estimator. Over the three phases, a total of 22 able-bodied participants participated in the study protocol. Participants provided written informed consent to participate in the study under Georgia Institute of Technology Institutional Review Board protocol H21184. Consent to publish participant

images was also obtained. In each phase, participants completed a set of 28 cyclic and non-cyclic activities, consisting of 66 total conditions as detailed in Scherpereel et al.[43] and outlined in Extended Data Fig. 2. During these activities, retroreflective markers on both the body and exoskeleton were tracked using a motion capture system at 200 Hz (Vicon Motion Systems). Furthermore, overground force plates and an instrumented treadmill (Bertec) were used to measure GRFs at 1,000 Hz. Owing to software limitations early in this study, motion capture data collected for two participants during phase 1 were collected at 120 Hz and subsequently upsampled to 200 Hz.

To sync the exoskeleton data with the participants' ground-truth joint biomechanics, the exoskeleton data were first upsampled to 200 Hz to match the frequency of the motion capture data. At the start of each exoskeleton trial when motion capture and GRF data were collected, participants kicked three times with their right leg while exoskeleton actuation was off. On the basis of this movement, the exoskeleton data were time shifted to maximize the R^2 between the right knee encoder and the resulting right knee joint kinematics computed from the biomechanical model.

Phase 1: initial data collection

In phase 1, ten participants (six males, four females, age of 23.7 ± 2.0 years, height of 176.3 ± 8.7 cm and body mass of 76.5 ± 13.3 kg) participated in a single day protocol collecting exoskeleton data and ground-truth human joint moments during the 28 cyclic and non-cyclic activities. As we could not deploy the task-agnostic controller until initial training data were collected, participants completed all activities with exoskeleton assistance turned off while sensor data were collected. To further increase the richness of our training data, we also hand-designed activity-specific controllers to collect actuated data for 46 out of the 66 total experimental conditions. Spline-based assistance was implemented as a function of gait phase when possible using eight- and six-node piecewise cubic Hermite interpolating polynomials at the hip and knee, respectively. Gait phase was estimated in real-time using the duration of the previous two strides measured from the pressure insole data[12]. The splines were shaped on the basis of the biological hip and knee moments reported in our previous work[43]. Furthermore, an impedance-based stance-swing state machine was used during several non-cyclic, impedance-like activities to control the device. Stance and swing phase were determined using the vertical GRF data measured from the pressure insoles. During leg swing, assistance was set to zero. During stance, the exoskeleton commanded torques using a spring-mode with a stiffness of 5 N m rad^{-1} and an equilibrium angle of 0°, which allowed all tasks to be completed without exceeding the actuator capabilities. These controllers were not intended to provide optimal assistance, but instead were used to collect a rich dataset for model training that included actuated exoskeleton data as well as the unactuated data.

Phase 2: training data with pilot model

In phase 2, five new participants (four males, one female, age of 23.4 ± 4.9 years, height of 171.4 ± 7.4 cm and body mass of 68.0 ± 12.7 kg) participated in a single day protocol similar to phase 1. However, during phase 2 we deployed the task-agnostic controller on the exoskeleton using a preliminary joint moment estimator trained using the phase 1 dataset. Each participant completed all 28 tasks while the exoskeleton provided assistance, which generated a dataset closely representing the controller to be deployed in the final phase.

Phase 3: online accuracy and user outcomes

In phase 3, ten participants (seven males, three females, age of 21.8 ± 2.5 years, height of 174.8 ± 8.5 cm and body mass of 71.7 ± 10.2 kg) participated in a multi-day protocol consisting of three sessions each centred on a specific outcome metric: model accuracy, metabolic cost and lower-limb biological joint work. Model accuracy was assessed with R^2, r.m.s.e. and normalized MAE with respect to ground-truth total hip and knee moments[31,32,60]. Normalized MAE was normalized to the peak-to-peak range of the corresponding ground-truth joint moments per task. Metabolic cost and lower-limb biological joint work provided an indication of the exoskeleton's impact on user effort[3,16,61–63], with metabolic cost more directly related to user energetics but only possible to measure and fairly compare during long bouts of repetitive tasks with consistent mechanical work requirements. During each session, the exoskeleton was controlled using the task-agnostic controller informed by the model trained on phase 1 and 2 data. Three of the participants were already enrolled in previous phases of the protocol. To ensure that the controller was evaluated on a participant-independent basis, we retrained separate models from random initialization for each of these participants while withholding the participant-specific data from the training set, ensuring that each participant was truly new to the network.

The network accuracy session was performed using the same protocol as phase 2 in which the exoskeleton was powered on for all 28 activities. During each condition, motion capture and GRF data were collected to compute ground-truth joint moments, which were then used to evaluate model performance. Each participant completed this session first, which served as a training session for the participant to adapt to the exoskeleton assistance, a critical component to evaluating human–exoskeleton interactions[64]. Owing to an error when exporting the segmented trials, data for the run task from one participant was not included in this study. Additionally, three participants returned to perform eight truly new tasks while collecting the same data to compare estimates to ground-truth moments (Extended Data Table 2).

During the lower-limb biological joint work session, participants completed six tasks detailed in Extended Data Table 3. These tasks were selected for our joint work analysis as measuring their resulting metabolic cost was infeasible owing to participant fatigue and time constraints. Each activity was repeated under three conditions: (1) wearing the exoskeleton with the task-agnostic controller providing assistance (exo on), (2) without wearing the exoskeleton (no exo) and (3) wearing the exoskeleton without actuation (zero torque). The order of these conditions was randomized; however, to minimize the time taken to complete the experimental protocol, the no exo condition was always placed either at the beginning or at the end, and all of the no exo activities were completed in succession, whereas the zero torque and exo on were alternated between each activity. Lower-limb positive biological joint work of the participant was computed by summing the average positive biological joint work from the hip, knee and ankle joints. Positive biological joint work for each joint was computed as the integral of the positive biological joint power for each joint, which was calculated as the product of the biological joint velocity and the biological joint moment (that is, the joint moment computed from ground-truth inverse dynamics after subtracting exoskeleton assistance torque). Results are presented as the average right leg positive joint work per repetition or stride. Owing to a bug in the exoskeleton data logger during the experimental protocol, stair ascent results for one participant and lunging results for one participant were not included in this study. Further, although this analysis quantified changes in the energetics of the biological joints, it did not account for muscle-level changes, such as muscle cocontraction.

At the beginning of the metabolic cost session, each participant completed a habituation protocol to reacclimate to the device. The habituation protocol consisted of level-ground walking on the treadmill at 1.25 m s^{-1} while wearing the exoskeleton controlled with the task-agnostic controller. Exoskeleton assistance was sequentially ramped up in four evenly spaced increments every 2 min until the participant reached full torque assistance. The participant then walked at full assistance (20 and 15% of the estimated biological hip and knee moments, respectively) for 5 min to complete the habituation protocol.

During each metabolic trial, oxygen intake ($\dot{V}O_2$) and carbon dioxide exhaust ($\dot{V}CO_2$) from each breath were measured using a metabolic

measurement system (TrueOne 2400, ParvoMedics). User metabolic cost was computed from the $\dot{V}O_2$ and $\dot{V}CO_2$ measurements using the modified Brockway equation[11,65]. These measurements were taken during four activities: level-ground walking at 1.25 m s^{-1}, lifting a 25 lb weight, 5° incline walking at 1.25 m s^{-1} and running on level ground at 2.5 m s^{-1}. For the lift weight trials, a metronome played a tone at 10 bpm and participants were instructed on each tone to use both hands to lift a 11 kg (25 lb) kettle bell off a shelf at waist height, touch the weight to the ground between their feet and then replace the weight on the shelf before returning to neutral standing and waiting for the next tone. For each activity, we tested the same three conditions as those of the joint work protocol: no exo, exo on and zero torque. Furthermore, the basal metabolic rate for each participant was measured from a 6 min standing trial while not wearing the exoskeleton. Owing to experimental time constraints, motion capture data were not collected during the metabolic experiments.

The level-ground walking, lift weight and incline walking trials were each completed using a within-participant counter-balanced design (ABCCBA)[12]. Each condition lasted 6 min, and steady-state metabolic cost for each trial was computed as the average of the last 3 min of data[11,66]. Owing to the strenuous nature of running, especially during zero torque, the running conditions were only completed once (ABC) and lasted 3.5 min. Steady-state metabolic cost was then computed from the output of a first-order model fit to the running data[11,66]. The basal metabolic rate measured for each participant was subtracted from the steady-state metabolic cost of each condition to compute the user's net metabolic rate required to complete each activity. Within each activity, the order of conditions was pseudorandomized, with the no exo condition either in the A or C position to minimize don and doff time. Owing to improper calibration of the metabolic system, the metabolic data for one participant were omitted from this study.

To further understand the effect of our approach on user metabolic cost during transient activities, we developed an extra protocol in which we recorded metabolic cost during a varying speed and incline treadmill circuit for three of the participants. Each 6 min trial involved walking at speeds varying from 0.6 to 1.5 m s^{-1} and running speeds of 2.0, 2.25 and 2.5 m s^{-1}. The speed of the treadmill was changed according to the profile shown in Fig. 4b. Similarly, the treadmill incline varied between 0° and 15° throughout the trial (Fig. 4b). The speed and incline profile of the treadmill was designed to decrease walking speeds at higher inclines, to keep the participant within an acceptable aerobic respiration range for valid metabolic cost measurements. Similar to the outcomes testing described above, the participants completed a randomized ladder protocol consisting of no exo, exo on and zero torque conditions (ABCCBA[12] as above) while wearing a metabolic measurement system (TrueOne 2400, ParvoMedics). Owing to an exoskeleton malfunction, one trial of data from one participant had to be removed; however, this trial was part of the C condition within the protocol, which should mitigate any adverse ordering effects from the removed datapoint. As described above, metabolic cost was computed using the Brockway equation and averaged across the last three minutes of the treadmill trial. This represents a consistent snapshot of metabolic cost across varying conditions.

Whereas qualitative feedback from participants about exoskeleton assistance was not formally collected, it is important to note that some participants expressed that the knee extension assistance did not feel helpful during level-ground walking and occasionally felt uncomfortable. Otherwise, the overall response from participants about the exoskeleton assistance was quite positive across the trials used to evaluate metabolic cost and lower-limb biological joint work.

Data availability

All of the exoskeleton sensor data and time-synced ground-truth biomechanics used for training and evaluating our joint moment estimator have been released with this publication. The data and details are available at https://doi.org/10.35090/gatech/75759. The training data includes 15 users performing all 66 conditions from the 28 task groups (Extended Data Fig. 2). The first 10 users include both unactuated data for all tasks and actuated data for those that could be mimicked with a heuristic controller. The following five users include actuated data for all tasks using a preliminary joint moment estimation model. The validation data includes ten users performing the same 66 conditions from the 28 task groups. The exoskeleton was actuated for all activities using the final model. Overall, this dataset includes more than 22 million labels of ground-truth moments across the left and right legs per lower-limb joint. Furthermore, the *P* values and corresponding participant count for all statistical tests are available in the Supplementary Data and a detailed description of our statistical tests is provided in the Supplementary Information.

Code availability

A Python package containing the code to deploy and test the models trained in this study using our corresponding dataset is available at https://doi.org/10.24433/CO.7641031.v2.

51. Camargo, J., Ramanathan, A., Flanagan, W. & Young, A. A comprehensive, open-source dataset of lower limb biomechanics in multiple conditions of stairs, ramps, and level-ground ambulation and transitions. *J. Biomech.* **119**, 110320 (2021).
52. Reznick, E. et al. Lower-limb kinematics and kinetics during continuously varying human locomotion. *Sci. Data* **8**, 282 (2021).
53. Winter, D. A., Sidwall, H. G. & Hobson, D. A. Measurement and reduction of noise in kinematics of locomotion. *J. Biomech.* **7**, 157–159 (1974).
54. Molinaro, D. D., Park, E. O. & Young, A. J. Anticipation and delayed estimation of sagittal plane human hip moments using deep learning and a robotic hip exoskeleton. In *Proc. 2023 IEEE International Conference on Robotics and Automation (ICRA)* 12679–12685 (IEEE, 2023).
55. Bai, S., Kolter, J. Z. & Koltun, V. An empirical evaluation of generic convolutional and recurrent networks for sequence modeling. Preprint at https://doi.org/10.48550/arXiv.1803.01271 (2018).
56. Golovin, D. et al. Google Vizier: a service for black-box optimization. In *Proc. 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* 1487–1495 (Association for Computing Machinery, 2017).
57. Lim, H., Kim, B. & Park, S. Prediction of lower limb kinetics and kinematics during walking by a single IMU on the lower back using machine learning. *Sensors* **20**, 130 (2020).
58. Mundt, M. et al. A comparison of three neural network approaches for estimating joint angles and moments from inertial measurement units. *Sensors* **21**, 4535 (2021).
59. Ancillao, A., Tedesco, S., Barton, J. & O'Flynn, B. Indirect measurement of ground reaction forces and moments by means of wearable inertial sensors: a systematic review. *Sensors* **18**, 2564 (2018).
60. Forner-Cordero, A., Koopman, H. J. F. M. & van der Helm, F. C. T. Inverse dynamics calculations during gait with restricted ground reaction force information from pressure insoles. *Gait Posture* **23**, 189–199 (2006).
61. Nuckols, R. W. et al. Mechanics of walking and running up and downhill: a joint-level perspective to guide design of lower-limb exoskeletons. *PLoS ONE* **15**, e0231996 (2020).
62. Alexander, N., Strutzenberger, G., Ameshofer, L. M. & Schwameder, H. Lower limb joint work and joint work contribution during downhill and uphill walking at different inclinations. *J. Biomech.* **61**, 75–80 (2017).
63. Lenton, G. K. et al. Lower-limb joint work and power are modulated during load carriage based on load configuration and walking speed. *J. Biomech.* **83**, 174–180 (2019).
64. Poggensee, K. L. & Collins, S. H. How adaptation, training, and customization contribute to benefits from exoskeleton assistance. *Sci. Robot.* **6**, eabf1078 (2021).
65. Brockway, J. M. Derivation of formulae used to calculate energy expenditure in man. *Hum. Nutr. Clin. Nutr.* **41**, 463–471 (1987).
66. Selinger, J. C. & Donelan, J. M. Estimating instantaneous energetic cost during non-steady-state gait. *J. Appl. Physiol.* **117**, 1406–1415 (2014).

Acknowledgements We acknowledge E. Rouse and K. Zealand for their contribution to the overarching technical direction; and A. Azocar, A. Memo and R. Jackson for their contribution to experimental design and analysis. In addition, E. Lamers, T. Malko, A. Metzger, N. Hite, C. Muntz, K. Chen, B. Piercy, P. Franks, X. Qin, D. Tachibana and J. Cogan contributed to developing and maintaining the exoskeleton hardware, soft goods and software used in this study. This research was also supported in part through research cyberinfrastructure resources and services provided by the Partnership for an Advanced Computing Environment (PACE) at the Georgia Institute of Technology, Atlanta, GA, USA. This project was funded in part by X, The Moonshot Factory (A.J.Y.), in part by the Nation Science Foundation NRI grant no. 1830215 (A.J.Y.), in part by the National Science Foundation FRR grant no. 2233164 (A.J.Y.), and in part by the National Science Foundation Graduate Research Fellowship Program grant no. DGE-2039655 (D.D.M. and K.L.S.).

Author contributions D.D.M. and K.L.S. contributed equally to the conceptualization, data acquisition, analysis, interpretation and writing of this work. E.B.S. contributed to the data acquisition, analysis and writing of this work. G.E. contributed to the conceptualization, analysis and writing of this work. M.K.S. contributed to the conceptualization, data acquisition, interpretation and writing of this work. A.J.Y. contributed to the conceptualization, interpretation and writing of this work.

Competing interests D.D.M. and A.J.Y. are inventors on a patent filed with the US Patent Office (US 18/340,981) by the Georgia Institute of Technology that covers some of the methods for state estimation described in this work. X, The Moonshot Factory, funded this work, contributed to its conceptualization and methodological design, and contributed much of the exoskeleton hardware and software used in this study.

Additional information

Supplementary information The online version contains supplementary material available at https://doi.org/10.1038/s41586-024-08157-7.

Correspondence and requests for materials should be addressed to Dean D. Molinaro.

Peer review information *Nature* thanks the anonymous reviewers for their contribution to the peer review of this work.

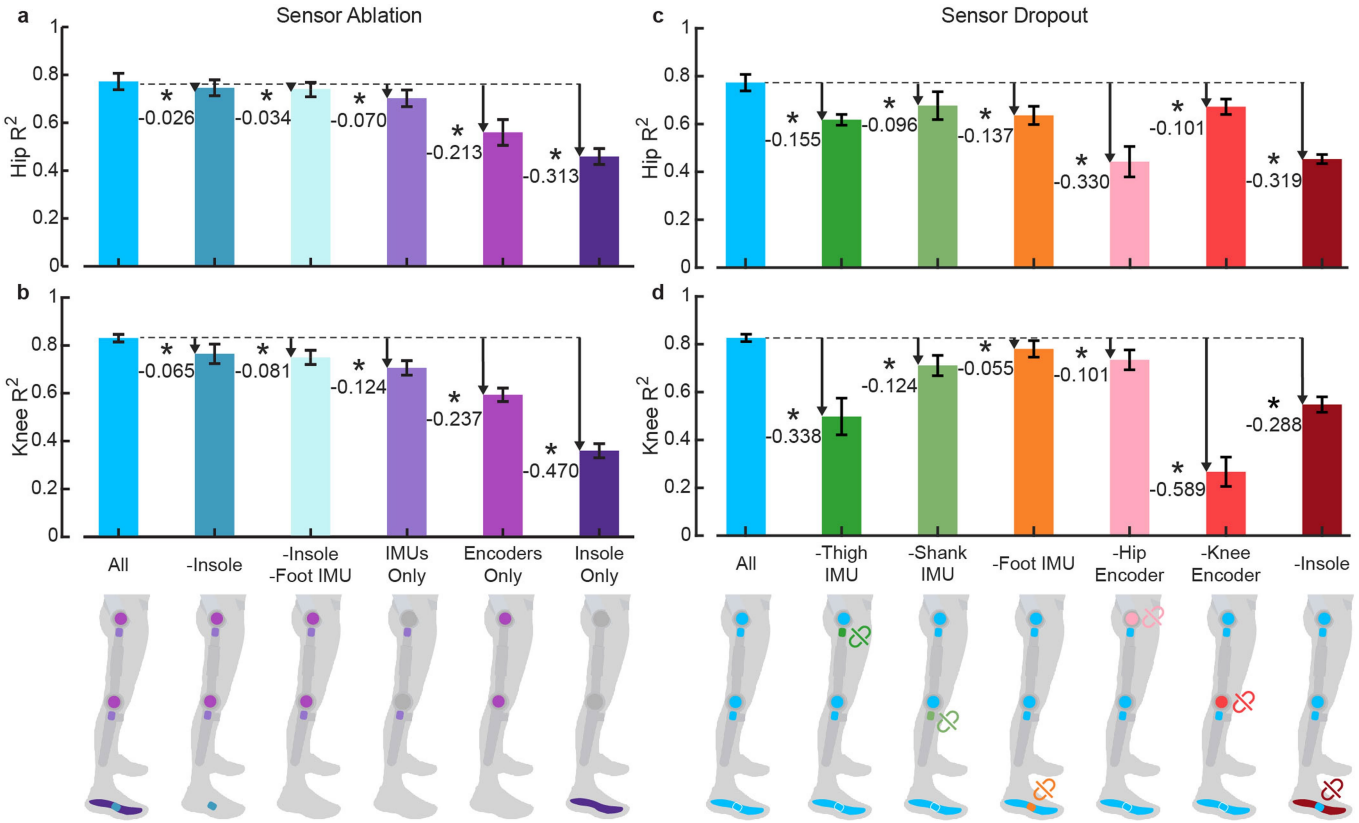

Extended Data Fig. 1 | Sensor ablations and dropouts for joint moment estimation. R^2 for (a) hip and (b) knee joint moment estimation was compared across five conditions where different combinations of sensors were removed. R^2 for (c) hip and (d) knee joint moment estimation was compared across six conditions where different individual sensor signals were set to zero to simulate sensor dropout. A one-way ANOVA was used to test for a main effect at the hip and knee, and for each joint, a *post-hoc* multiple comparisons test with a Bonferroni correction was conducted to evaluate comparisons between the All condition and each ablation or dropout condition. The bar height depicts the average performance across all 28 tasks and 10 participants. Error bars depict ±1 standard deviation across participants. The relative decrease in performance is depicted relative to the model trained on all sensors and without artificial sensor dropout. Asterisks indicate statistical significance ($P < 0.05$). Exact P values are provided in Supplementary Data 1.

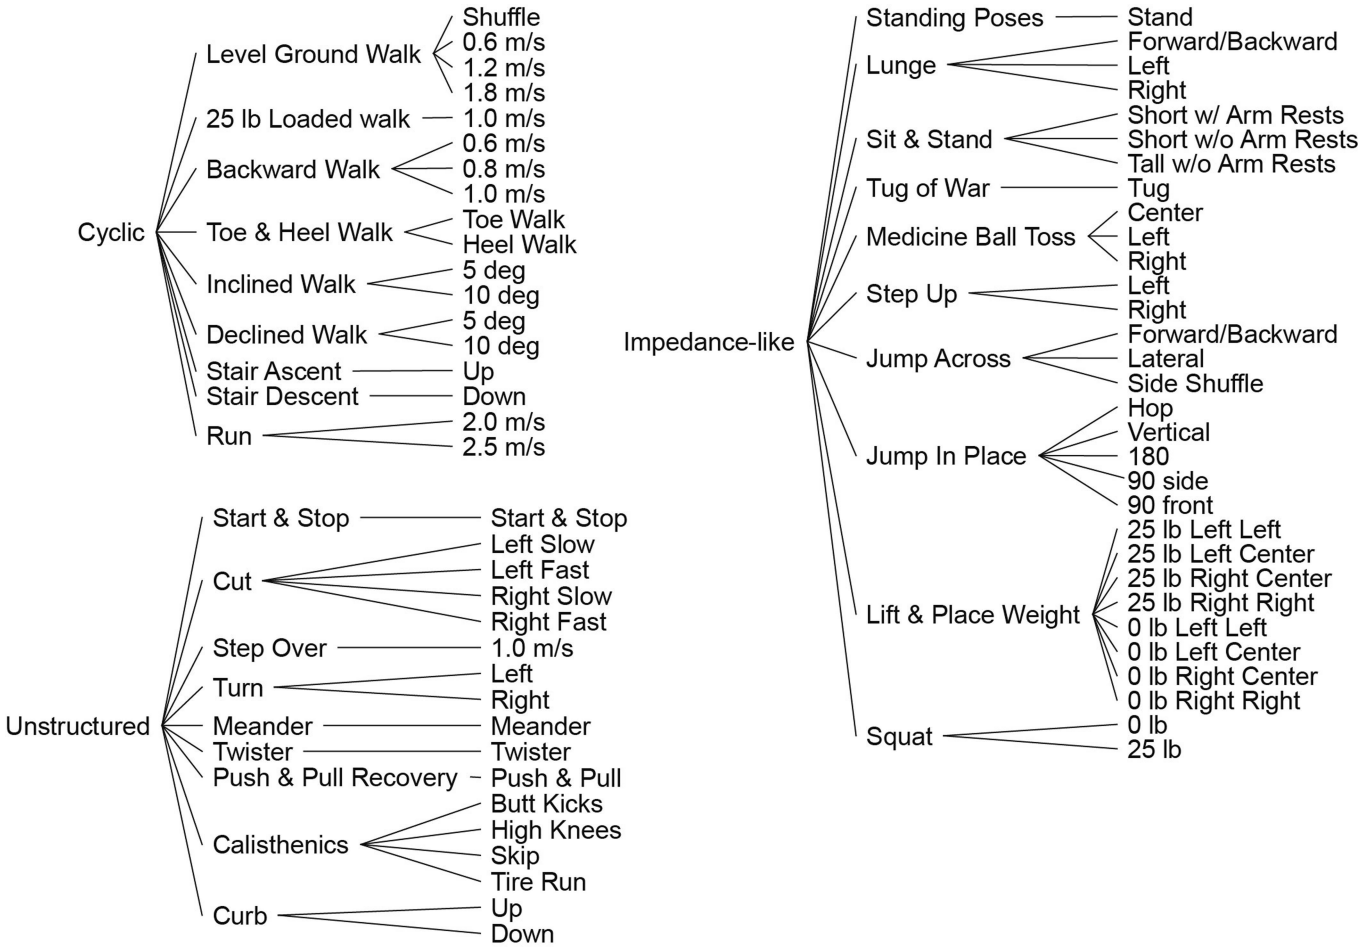

Extended Data Fig. 2 | Task breakdown of our exoskeleton dataset labeled with ground-truth biomechanics. The 28 tasks were binned into three categories: cyclic, impedance-like, and unstructured based on their normative biomechanics. Additionally, many tasks contained multiple experimental conditions (66 conditions total).

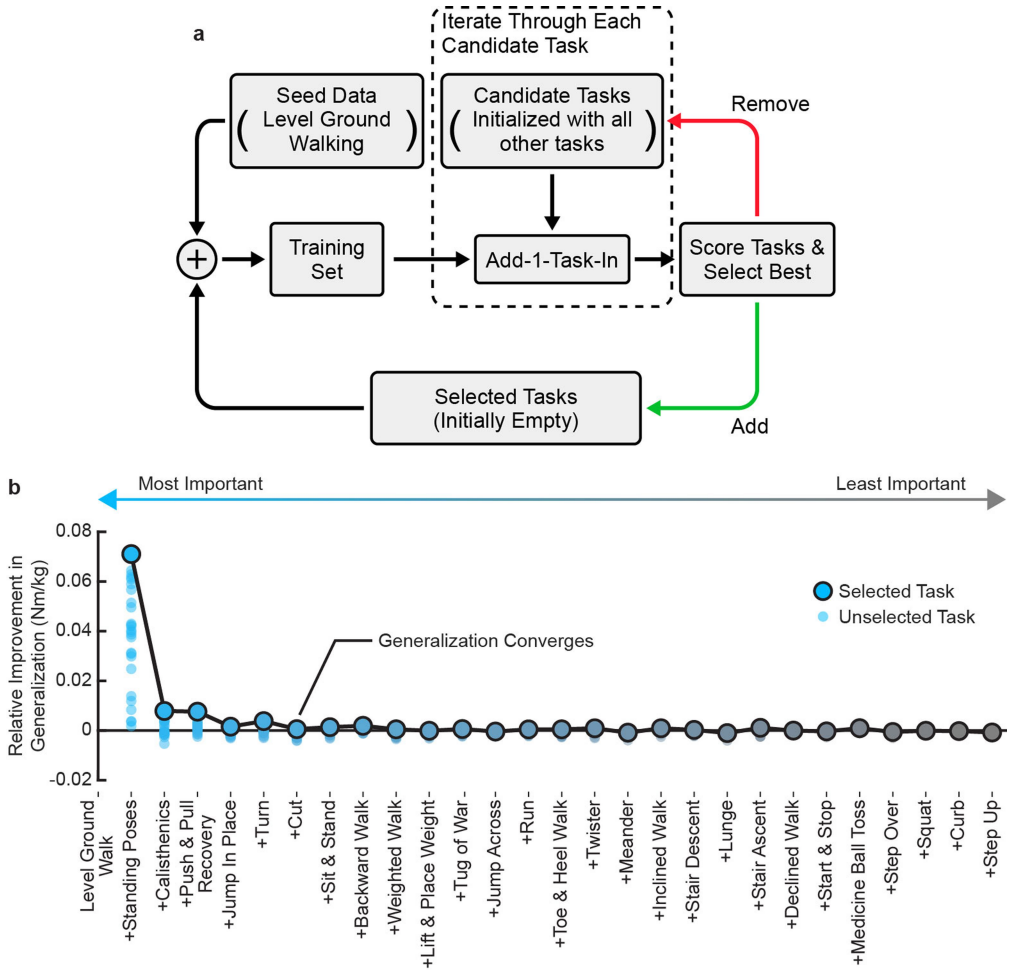

Extended Data Fig. 3 | Forward task selection algorithm and results. (a) To determine the set of activities most important for model generalizability, tasks were iteratively added into the model training set, scored, and subsequently selected based on their contribution to overall model generalizability on the validation data. (b) The relative improvement in generalization for each selected task is shown. With each new task, the model was retrained and the relative improvement in its ability to predict biological hip and knee moments across all other (27) tasks was computed (average improvement shown). Saturated generalization was reached when the sum of the relative improvement reached 95% of the total sum from adding in all of the activities. Results were computed from leave-one-subject-out cross-fold validation using a 12-participant dataset; however, error bars were omitted for clarity.

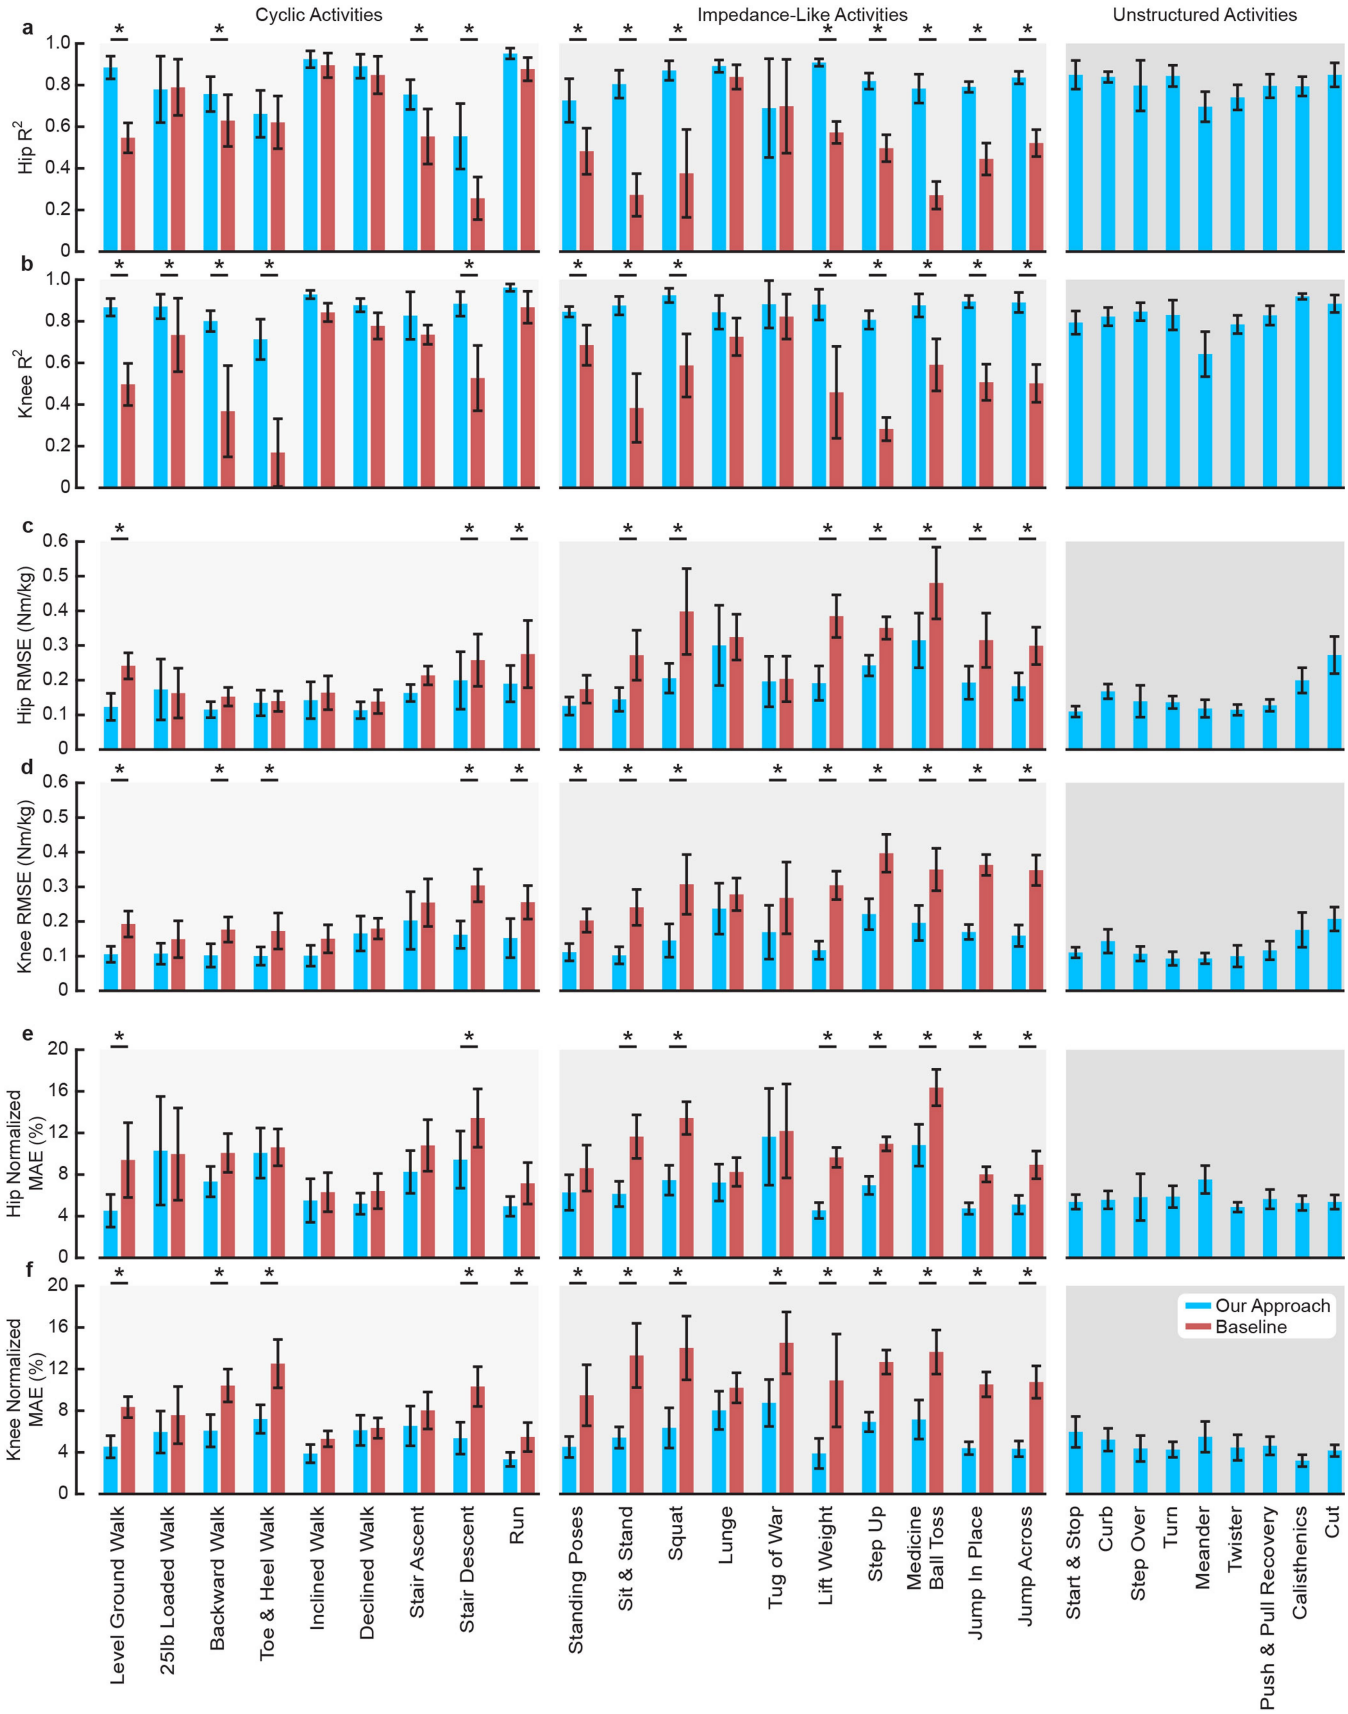

Extended Data Fig. 4 | Detailed online joint moment estimation performance. The resulting (a) hip R^2, (b) knee R^2, (c) hip RMSE, (d) knee RMSE, (e) hip normalized MAE, and (f) knee normalized MAE of our joint moment estimator is shown for each activity and is compared to the baseline method. The MAE results are normalized to their corresponding peak-to-peak range of the ground-truth joint moments. The bars depict the inter-subject mean across 10 subjects except for the Run condition where n = 9, and the error bars depict ±1 standard deviation. Asterisks indicate statistical significance ($P < 0.05$; exact P values are provided in Supplementary Data 1).

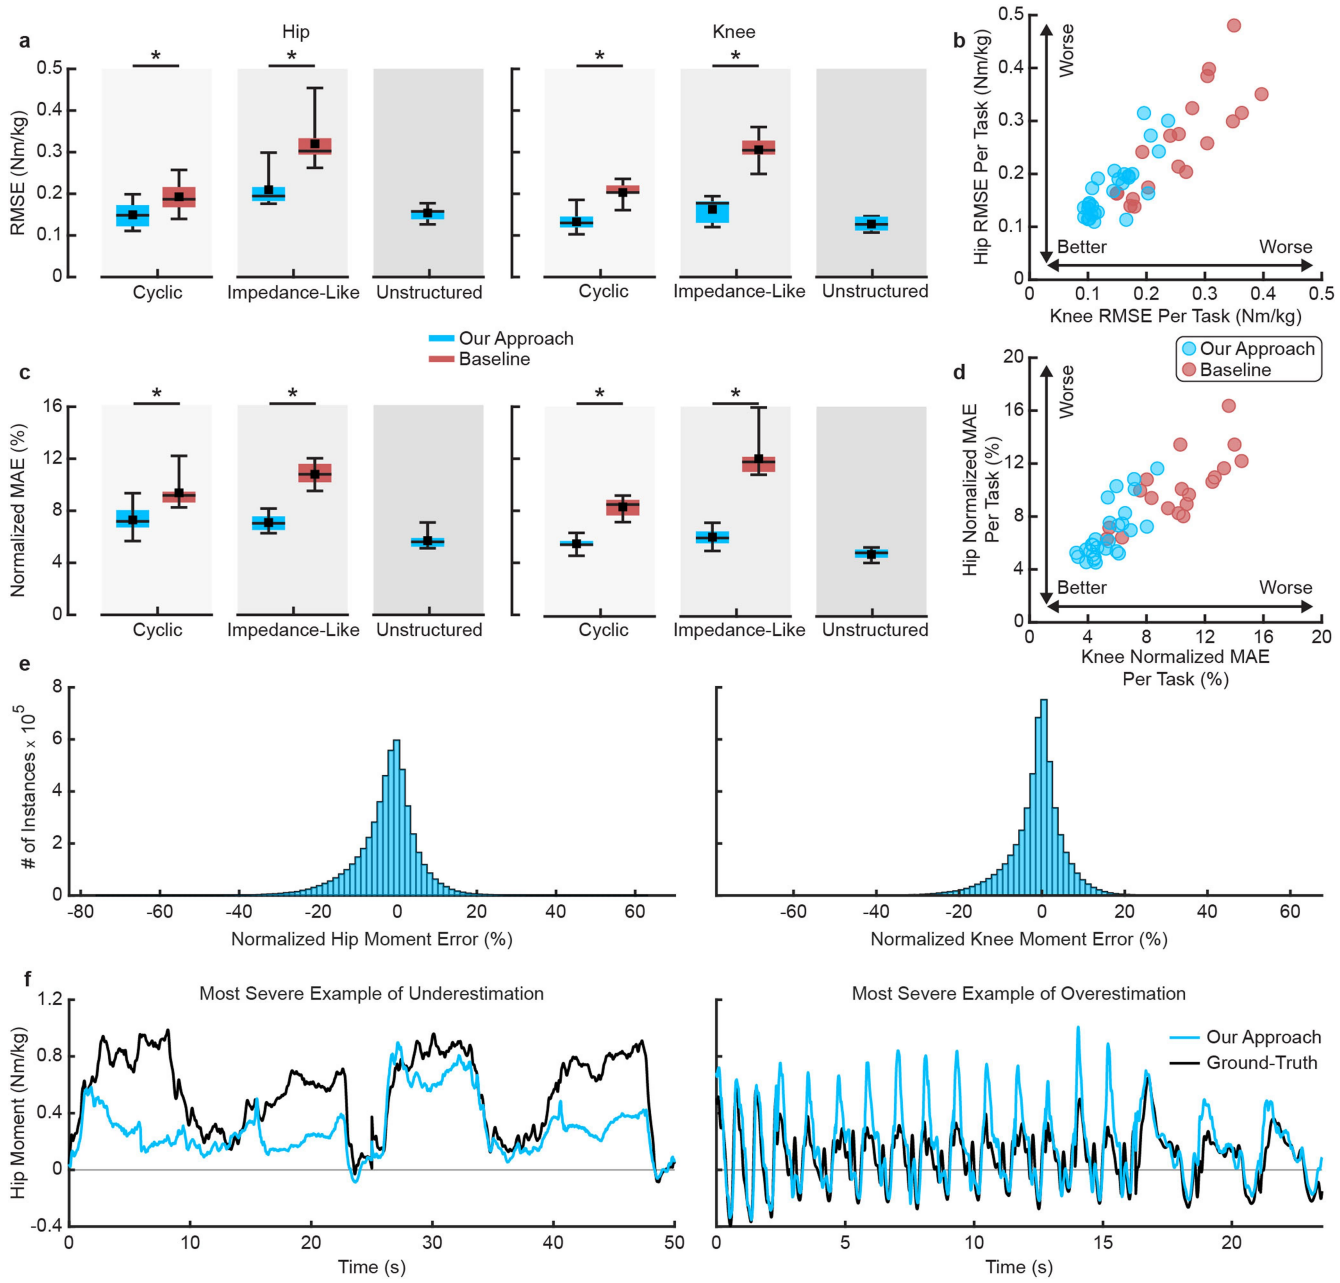

Extended Data Fig. 5 | Online joint moment estimation RMSE, normalized MAE, and analysis of under- and overestimation. (a) Our approach significantly reduced RMSE at the hip by 0.04 ± 0.02 Nm/kg (22 ± 10%, P = 0.03) and at the knee by 0.07 ± 0.02 Nm/kg (35 ± 8%, P = 3×10^{-8}) compared to the baseline method during cyclic activities. For impedance-like tasks, our approach reduced RMSE by 0.11 ± 0.03 Nm/kg (35 ± 9%, P = 8×10^{-7}) at the hip and by 0.14 ± 0.03 Nm/kg (47 ± 9%, P = 3×10^{-15}) at the knee compared to the baseline method. (b) Estimator RMSE is shown per task for our approach and the baseline method. (c) Our approach significantly reduced normalized MAE at the hip by 22.2% ± 11.8% (P = 6×10^{-5}) and by 34.1% ± 8.9% (P = 6×10^{-7}) at the knee compared to the baseline method during cyclic activities. For impedance-like tasks, our approach reduced normalized MAE by 34.4% ± 8.0% (P = 3×10^{-9}) at the hip and by 50.2% ± 14.2% (P = 8×10^{-14}) at the knee compared to the baseline method. (d) Estimator normalized MAE is shown per task for our approach and the baseline method. (e) Histograms depicting the distribution of under- and overestimation of the hip and knee moments relative to ground-truth are shown. (f) Time series depicting the most severe example of underestimation and overestimation across all validation trials are shown (both of which occurred at the hip). In the box plots, black squares depict inter-subject mean; colored boxes depict the interquartile range; horizontal lines within boxes depict the inter-subject median; error bars depict the inter-subject minimum and maximum (n = 10). Asterisks indicate statistical significance (P < 0.05; exact P values are provided in Supplementary Data 1). In the scatter plots, each marker corresponds to the inter-subject average result for a single task (n = 10 expect for the Run condition where n = 9).

Extended Data Table 1 | TCN hyperparameter optimization

Hyperparameter	Stage 1 Search Space	Stage 2 Search Space	Selected Value
# of Filters per Layer	8, 16, 24, 32, 48, 56, 64, 72, 80	56, 64, 72, 80	80
# of Residual Blocks (2 Convolutional Layers per Block)	2, 3, 4, 5, 6	5, 6	5
Kernel Size	2, 3, 4, 5, 6, 7, 8, 9, 10	2, 4, 5	5
Activation Function	ELU, GELU, ReLU, Swish	ELU, ReLU, Swish	ReLU
Learning Rate	[1e-6, 1e-3]	5e-6, 1e-5, 5e-5	5e-5
Block Normalization	Batch Normalization, Layer Normalization, Weight Normalization	Batch Normalization, Weight Normalization	Weight Normalization
Dropout Type	Element-Wise, Spatial	Spatial	Spatial
Dropout Probability	[0, 0.3]	0.15	0.15
L1 Kernel Regularization	[1e-5, 1e-3]	0	0
L2 Kernel Regularization	[1e-5, 1e-3]	0	0
L2 Bias Regularization	[1e-5, 1e-3]	0	0

Values shown without brackets were optimized categorically. Arrays shown with closed brackets were optimized on the closed interval. Hyperparameter combinations resulting in an input sequence greater than 250 were omitted due to limitations in the simulated data at the bounds of each trial. In Stage 2, kernel and bias regularization were turned off since Stage 1 optimization generally minimized these values across multiple optimization instances.

Extended Data Table 2 | Activities for testing exoskeleton performance on entirely novel tasks

Task	Description
Burpees	Participants were instructed to place their hands on the ground and step back into a push up posture. Then they stepped forward to return to standing. This was performed 3 times leading with the right leg and 3 with the left.
Cart Walk	Participants were instructed to push a stationary rod (simulating a heavy cart) while walking on the treadmill at 0.8 m/s for 20 seconds. They then performed the same task while pulling on the rod and walking backward.
Crouch Walk	Participants were instructed to walk on the treadmill at 0.8 m/s while in a crouched posture for 20 seconds.
Dizzy Walk	Participants spun around for 10 seconds and then walked back and forth across the force plates until no longer showing signs of dizziness.
Steep Incline/Decline Walk (15°)	Participants were instructed to walk up a 15° incline at 1.2 m/s for 20 seconds. They then repeated this while walking down the same decline.
Layup	Participants were instructed to run up to the force plates, plant a single foot, and jump to maximum vertical height while landing with both feet on the force plates. This was performed 3 times for both right and left legs.
Mountain Climbers	Participants were instructed to assume a push-up posture and then alternate bringing each knee up toward their chest as quickly as possible. This was performed for 20 seconds.
Split Walk	Participants walked on a treadmill with the right leg belt at 1.6 m/s and the left leg belt at 0.8 m/s for 20 seconds. They then repeated this task with the left belt at 1.6 m/s and the right at 0.8 m/s.

Extended Data Table 3 | Activities for testing exoskeleton effects on lower-limb biological joint work

Task	Speed	Description
Lift & Place Weight	12 lifts per min	Participants were instructed to lift a 25lb weighted bag with both hands. On the first tone, participants started from standing and lifted the bag to waist height. They then paused for the second tone, at which they set the bag back on the floor and then returned to standing.
Squat	15 squats per minute	Participants were instructed to squat with the 25lb weight until they barely touched a pillow resting on a low stool. Then they returned to standing.
Lunge	15 lunges per minute	Participants were instructed to lunge with their right leg forward until their left knee barely touched a pillow laid on the ground. Then they returned to standing.
Sit & Stand	12 sit & stands per minute	Participants were instructed to alternate sitting down and standing up on the tone without using their arms for assistance.
Right Leg Step Up	12 step ups per min	Participants were instructed to step up onto a 46 cm box with their right leg. On the first tone, participants stepped up with their right leg and paused standing on one foot. On the second tone they stepped back down to have both feet on the ground.
Stair Ascent	Self-selected speed	Participants were instructed to complete four bouts of walking up a six-step staircase at their natural walking pace.