



Intersecting kinematic encoding and readout of intention in autism

Noemi Montobbio^a, Andrea Cavallo^{b,c}, Dalila Albergò^{a,d}, Caterina Ansuini^b, Francesca Battaglia^{b,e}, Jessica Podda^{b,f}, Lino Nobili^{e,g}, Stefano Panzeri^{a,h,1}, and Cristina Becchio^{b,i,1}

^aNeural Computational Laboratory, Center for Human Technologies, Istituto Italiano di Tecnologia, Genoa 16152, Italy; ^bCognition, Motion and Neuroscience Laboratory, Center for Human Technologies, Istituto Italiano di Tecnologia, Genoa 16152, Italy; ^cDepartment of Psychology, University of Turin, Turin 10124, Italy; ^dCenter for Mind/Brain Sciences, University of Trento, Rovereto 38068, Italy; ^eChild Neuropsychiatric Unit, G. Gaslini Institute, Genoa 16147, Italy; ^fScientific Research Area, Italian Multiple Sclerosis Foundation, Genoa 16149, Italy; ^gDepartment of Neurosciences, Rehabilitation, Ophthalmology, Genetics, Maternal and Child Health (DiNOGMI), University of Genoa, Genoa 16132, Italy; ^hDepartment of Excellence for Neural Information Processing, Center for Molecular Neurobiology, University Medical Center Hamburg-Eppendorf (UKE), Hamburg 20251, Germany; and ⁱDepartment of Neurology, University Medical Center Hamburg-Eppendorf (UKE), Hamburg 20246, Germany

Edited by Marlene Behrmann, Psychology and Neuroscience Institute, Carnegie Mellon University, Pittsburgh, PA; received August 12, 2021; accepted December 17, 2021

Observers with autism spectrum disorders (ASDs) find it difficult to read intentions from movements. However, the computational bases of these difficulties are unknown. Do these difficulties reflect an intention readout deficit, or are they more likely rooted in kinematic (dis-)similarities between typical and ASD kinematics? We combined motion tracking, psychophysics, and computational analyses to uncover single-trial intention readout computations in typically developing (TD) children ($n = 35$) and children with ASD ($n = 35$) who observed actions performed by TD children and children with ASD. Average intention discrimination performance was above chance for TD observers but not for ASD observers. However, single-trial analysis showed that both TD and ASD observers read single-trial variations in movement kinematics. TD readers were better able to identify intention-informative kinematic features during observation of TD actions; conversely, ASD readers were better able to identify intention-informative features during observation of ASD actions. Crucially, while TD observers were generally able to extract the intention information encoded in movement kinematics, those with autism were unable to do so. These results extend existing conceptions of mind reading in ASD by suggesting that intention reading difficulties reflect both an interaction failure, rooted in kinematic dissimilarity between TD and ASD kinematics (at the level of feature identification), and an individual readout deficit (at the level of information extraction), accompanied by an overall reduced sensitivity of intention readout to single-trial variations in movement kinematics.

autism | kinematics | action observation | intention readout | intersection

The ability to intuit what others are thinking or wanting from observing their behavior—mind reading—is key to social interaction. Much like print reading, mind reading involves the derivation of meaning from signs (1). In print reading, the signs are marks on paper. In mind reading, the signs are movement traces (2). Individuals with autism spectrum disorders (ASDs) have difficulty inferring the mental states of others, including their intentions, from their body movements (e.g., refs. 3–5). However, the computational bases of these difficulties are unknown.

One proposal is that such difficulties reflect a specific deficit in mind reading (1). Typically developing (TD) observers read intention by extracting and processing subtle intention-related kinematic variations (about 3% of the total variance) out of trial-to-trial variations unrelated to intention (6). Individuals with ASD would have difficulty reading intention, possibly due to an overall lower sensitivity to single-trial kinematics or to a deficit in identifying or processing intention-informative variations in single-trial movement kinematics. This accords with the view that individuals with ASD have difficulty sampling relevant and irrelevant variability (7) and therefore, get lost in incidental,

trial-to-trial variations (8). This hypothesis predicts a general impairment in intention reading in autism.

Alternatively, difficulties in attributing intentions to actions could be rooted in kinematic (dis-)similarities between typical and autistic kinematics (9). This hypothesis is based on the view that the same internal models used during action execution serve as the basis for action perception, prediction, and inference during action observation (10, 11). Because individuals with ASD move differently compared with TD individuals—in particular, they differ in the way they prospectively control their intentional actions (12–14)—this hypothesis makes the distinctive prediction that observers with ASD, with autistic internal models, should be less accurate in predicting the actions performed by TD individuals relative to those performed by individuals with ASD. Conversely, TD observers, with typical models, should be less accurate in predicting the actions performed by individuals with ASD relative to those performed by TD individuals. From this perspective, ASD difficulties would not reflect

Significance

A major challenge in studying intention reading is high motor variability. Analyses conducted across trials provide insights into what happens on average; however, they may obscure how individual observers read intention information in individual movements. We combined motion tracking, psychophysics, and computational analyses to examine intention reading in autism spectrum disorders (ASDs) with single-trial resolution. Results revealed that a sizeable fraction of ASD observers can identify intention-informative variations in ASD (but not in typically developing) movement kinematics, but they are nonetheless unable to extract the encoded intention information. This approach not only enhances our basic understanding of mind reading in ASD but also provides potential avenues for the rational design of training procedures to improve the reading of others' actions.

Author contributions: A.C., C.A., S.P., and C.B. designed research; D.A., C.A., F.B., J.P., and L.N. performed research; N.M. and S.P. contributed new analytic tools; N.M., A.C., and S.P. analyzed data; S.P. and C.B. supervised the project; S.P. and C.B. wrote the paper with contributions from N.M. and A.C., and all authors revised and approved the final manuscript.

The authors declare no competing interest.

This article is a PNAS Direct Submission.

This open access article is distributed under Creative Commons Attribution-NonCommercial-NoDerivatives License 4.0 (CC BY-NC-ND).

¹To whom correspondence may be addressed. Email: s.panzeri@uke.de or cristina.becchio@iit.it.

This article contains supporting information online at <http://www.pnas.org/lookup/suppl/doi:10.1073/pnas.2114648119/-DCSupplemental>.

Published January 31, 2022.

an individual intention reading failure but rather, would arise from reciprocal difficulties in social interaction (9, 15).

Previous work has shown a TD advantage for TD actions (3, 16) but no ASD advantage for ASD actions (3). This has been interpreted as evidence that TD observers' models are tuned to typical actions, whereas ASD observers' models are tuned (or possibly untuned) comparably with both TD and ASD actions (3). However, the advantage TD observers show for TD actions is not necessarily indicative of kinematic similarity and might instead reflect the higher informativeness of TD kinematics relative to ASD kinematics: that is, the fact that TD actions encode more intention information compared with ASD actions (12). Conversely, the lack of advantage of ASD observers for ASD actions might reflect the lower informativeness and higher variability of ASD kinematics relative to TD kinematics (3, 13). Thus, previous studies cannot rule out the possibility that group differences in intention reading relate to differences in how intention information is encoded in TD and ASD kinematics. Moreover, because intention reading was computed as the average response across individual trials with variable kinematics, these studies cannot determine the readout computations that inform intention inferences in TD and ASD observers: what information TD and ASD observers read in TD and ASD kinematics and how.

This study aimed to move beyond these limitations by combining accurate recording of movement kinematics and psychophysical measures of intention discrimination with a specifically designed analytic framework. This framework allowed us to link kinematic encoding—how intention information is encoded in TD and ASD movement kinematics during action execution—and kinematic readout—how TD and ASD observers read intention information encoded in TD and ASD visual kinematics during action observation—at the single-subject, single-trial level. In a two-by-two factorial design, TD and ASD children observed actions performed by TD and ASD children. Using a kinematic encoding model, we first quantified the intention information in TD and ASD single-movement kinematics and determined the set of kinematic features that encode this information in TD and ASD actions. Then, we developed a kinematic readout model to quantify how and how well TD and ASD observers read the intention information encoded in TD and ASD actions. Finally, adapting methods developed in refs. 6 and 17, we examined how kinematic encoding and readout intersect at the single-trial level across observer groups and observed actions. This approach allowed us to move beyond representations averaged over trials and participants and test alternative hypotheses regarding the origin of difficulties in intention reading in ASD.

Results

Eight- to 13-year-old ASD children ($n = 35$) and age- and intelligence quotient (IQ)-matched TD children ($n = 35$) watched a hand reaching for a bottle and judged on the intention of the observed grasp (*Materials and Methods*). To capture natural movement variability, we selected 100 representative reach-to-grasp actions (50 ASD actions and 50 TD actions) from a large action dataset obtained by tracking and simultaneously filming TD and ASD children reaching for a bottle with the intent to place or pour (12). In a within-subjects counterbalanced order, participants watched videos of reach-to-grasp actions performed by TD children and ASD children (Fig. 1A–C and *SI Appendix, Fig. S1A*). P values of all statistical comparisons are reported graphically in Figs. 1–6 and numerically in *SI Appendix, Tables S1–S5*. Effect sizes are reported in *SI Appendix, Tables S1–S5*.

Trial-Averaged Intention Discrimination in TD and ASD Observers.

We used logistic mixed effects models to test statistically whether average intention discrimination performance, computed as the

fraction of correct intention choices, differed from chance and across observer groups (TD, ASD) and observed actions (TD, ASD). We found a significant main effect of observer group, indicating that ASD observers were poorer at discriminating intention than TD observers (Fig. 1D and *SI Appendix, Table S2*). Neither the main effect of observed action nor the interaction between observer group and observed action reached significance (*SI Appendix, Table S1*). Intention discrimination performance was above chance for TD observers but not for ASD observers (Fig. 1D and *SI Appendix, Table S4*). Additional analyses conducted to explore the relationship between intention discrimination performance and autistic traits (in a subset of participants for whom autistic trait quantification was available) (*Materials and Methods*) revealed that TD observers with higher Social Responsiveness Scale (SRS) scores were poorer at discriminating intention than TD observers with lower SRS scores (*SI Appendix, Fig. S3*). For both TD and ASD observers, control analyses revealed no effect of IQ on trial-averaged performance (*SI Appendix, Table S3*).

Kinematic Encoding and Readout of Intention Information at the Single-Subject, Single-Trial Level.

The above results capture trial-averaged differences between groups. However, they do not quantify what information individual TD and ASD observers read in TD and ASD kinematics and how. To do so, we developed an analytic framework to directly model how information encoded in movement kinematics is read out with single-subject, single-trial resolution. Our formalism was inspired by recent mathematical advances in linking information encoding and readout in a neural population to inform single-trial behavior choices (17, 18). Here we adapted this formalism to investigate how information is coded in movement kinematics (rather than in a neural population).

Kinematic Encoding of Intention Information in TD and ASD Actions.

The first step was to determine kinematic encoding: that is, how intention information is encoded in trial-to-trial variations in movement kinematics of TD and ASD actions. Fig. 2A shows the temporal profile of two kinematic variables, wrist height (WH) and grip aperture (GA), under the intention “to pour” and “to place” during TD and ASD reach-to-grasp movements. Each line represents a single reach-to-grasp act. Consistent with previous reports (6, 19), individual movement traces showed a large variability across trials and individuals. To isolate the variability that conveys intention information from the trial-to-trial variability unrelated to intention, we developed a single-trial kinematic encoding model based on logistic regression. We represented single-trial kinematics as a time-dependent vector in the multidimensional space of values of 15 intention-sensitive kinematic variables (*Materials and Methods*). The kinematic encoding model computed, separately for TD and ASD actions, the probability that a reach-to-grasp movement was performed with a given intention (to pour) as a logistic regression of the time-dependent kinematic vector, with a drift term for modeling the accumulation of evidence over time (Fig. 2B and *Materials and Methods*). Across trials, model performance for TD actions, measured as the fraction of action intentions correctly predicted by the model, was above 95%. For ASD action, model performance was lower but still above 90% (Fig. 2C and *SI Appendix, Table S4*). This suggests that, despite the large variability across trials and individuals, both TD and ASD actions exhibited a consistent pattern of intention modulation.

Fig. 2D visualizes the contribution (weight) of each kinematic variable to the encoding of intention information in TD and ASD kinematics, as measured by the regression coefficient of the variable in the logistic regression. A positive (negative) encoding weight is assigned to a variable distributed across trials, with higher (lower) values for grasp-to-pour actions compared with grasp-to-place actions. For example, WH is generally higher for

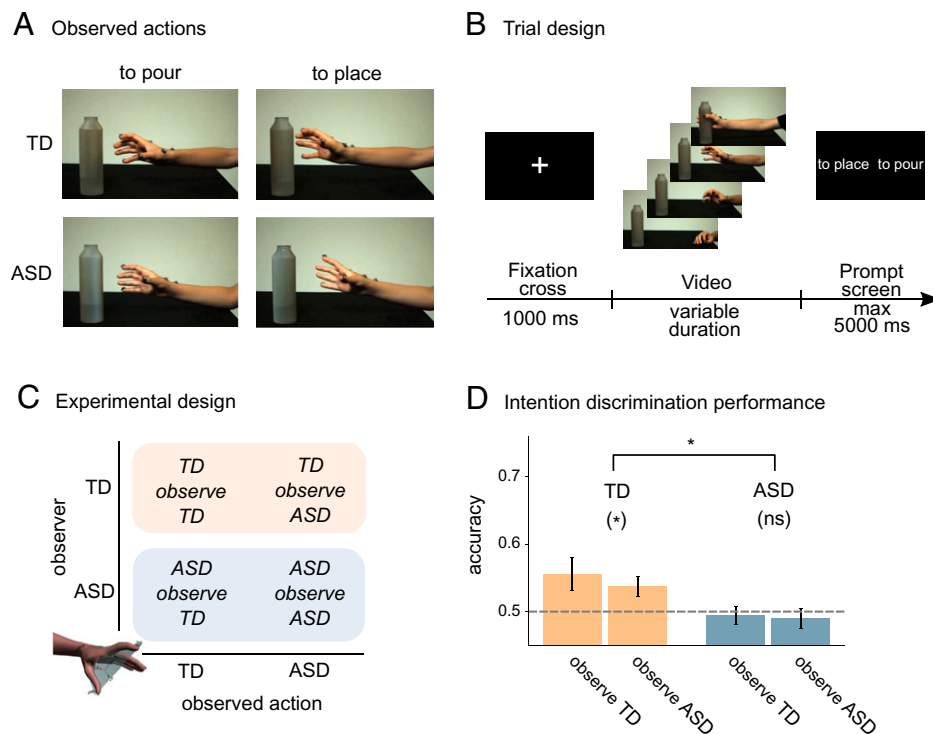


Fig. 1. Experimental design and results of intention discrimination. (A) Example video frames of grasp-to-pour and grasp-to-place actions produced by TD and ASD children. Each video began with reach onset and ended with the contact between the hand and the bottle. (B) Trial design of the intention discrimination task. (C) Schematic representation of the experimental design. (D) Trial-averaged intention discrimination performance (fraction correct) for each observer group and observed action. Histograms represent mean \pm SEM across participants. ns indicates $P > 0.05$. * $P < 0.05$.

the grasp-to-place action and is, therefore, negatively weighted for both TD and ASD actions (Fig. 2A). TD and ASD actions exhibited partially different patterns of intention encoding. For TD actions, intention information was encoded in WH and in the displacement of the thumb (z thumb [TZ]) and index finger (z index [IZ]) along the z axis (SI Appendix, Fig. S2). For ASD actions, intention information was distributed across a larger set of variables. Some variables carrying intention information in ASD kinematics also carried intention information in TD kinematics (WH, TZ). Other variables informative in ASD kinematics (wrist velocity [WV], GA, y thumb [TY], and x dorsum plane [DPX]) did not carry intention information in TD kinematics. Consistent with previous work demonstrating differences in the way that TD and ASD prospectively control their actions (13), these results demonstrate differences in the kinematic encoding of to pour and to place intentions in TD and ASD actions.

Kinematic Readout of Intention Information in TD and ASD Observers.

Having determined how intention information is encoded in the single-trial kinematics of TD and ASD actions, we next fitted single-trial intention choices to a kinematic readout model to investigate how TD and ASD observers read such information from observing TD and ASD actions. The kinematic readout model computed the probability of intention choice (to pour) in each trial as a logistic regression of the time-dependent kinematic vector for that trial (Fig. 3A).

Across trials and conditions, kinematic readout model performance, measured as the fraction of intention choices correctly predicted by the model, was significantly above chance (Fig. 3B and SI Appendix, Table S4). The strong correlation between observed intention discrimination performance and performance predicted by the readout model confirmed that our kinematic readout model was able to capture intention discrimination performance at the individual level (Fig. 3C). Although

reaction times were not used to fit the model parameters, we also found weak, but significant, negative trial-to-trial relationship between reaction time and model prediction confidence (Fig. 3D). This suggests that observers were slightly faster to judge intention on trials that were classified with greater confidence by the model. Taken together, these analyses suggest that our kinematic readout model provided a plausible description of how well observers discriminated intention from single-trial kinematics.

Sensitivity of Intention Readout to Movement Kinematics. Having verified the ability of the kinematic readout model to capture statistical dependencies between intention choices and single-trial variations in movement kinematics, we used it to test the hypothesis that poor intention discrimination in ASD (Fig. 1D) reflects an overall reduced sensitivity of intention readout to single-trial variations in visual kinematics. One concrete way to assess this is to measure how well the kinematic readout model predicts single-trial intention choices (regardless of whether the predicted intention choices are correct or incorrect). If intention readout in ASD is not sensitive to single-trial variations in movement kinematics, the kinematic readout model should be at chance in predicting ASD intention choices. As shown in Fig. 3B, this was not the case. Sensitivity of intention readout as measured by kinematic readout model performance was lower in ASD compared with TD (SI Appendix, Table S2) but still significantly above chance for both observer groups and observed actions (SI Appendix, Table S4). This suggests that lower sensitivity of the intention readout to the single-trial kinematics cannot fully account for ASD failure to discriminate intention apparent in Fig. 1D.

Identifying Readers. Readout patterns are variable across observers. We next used the kinematic readout model to parse this heterogeneity and identify, at the individual level, observers whose intention readout was sensitive to single-trial variations in movement

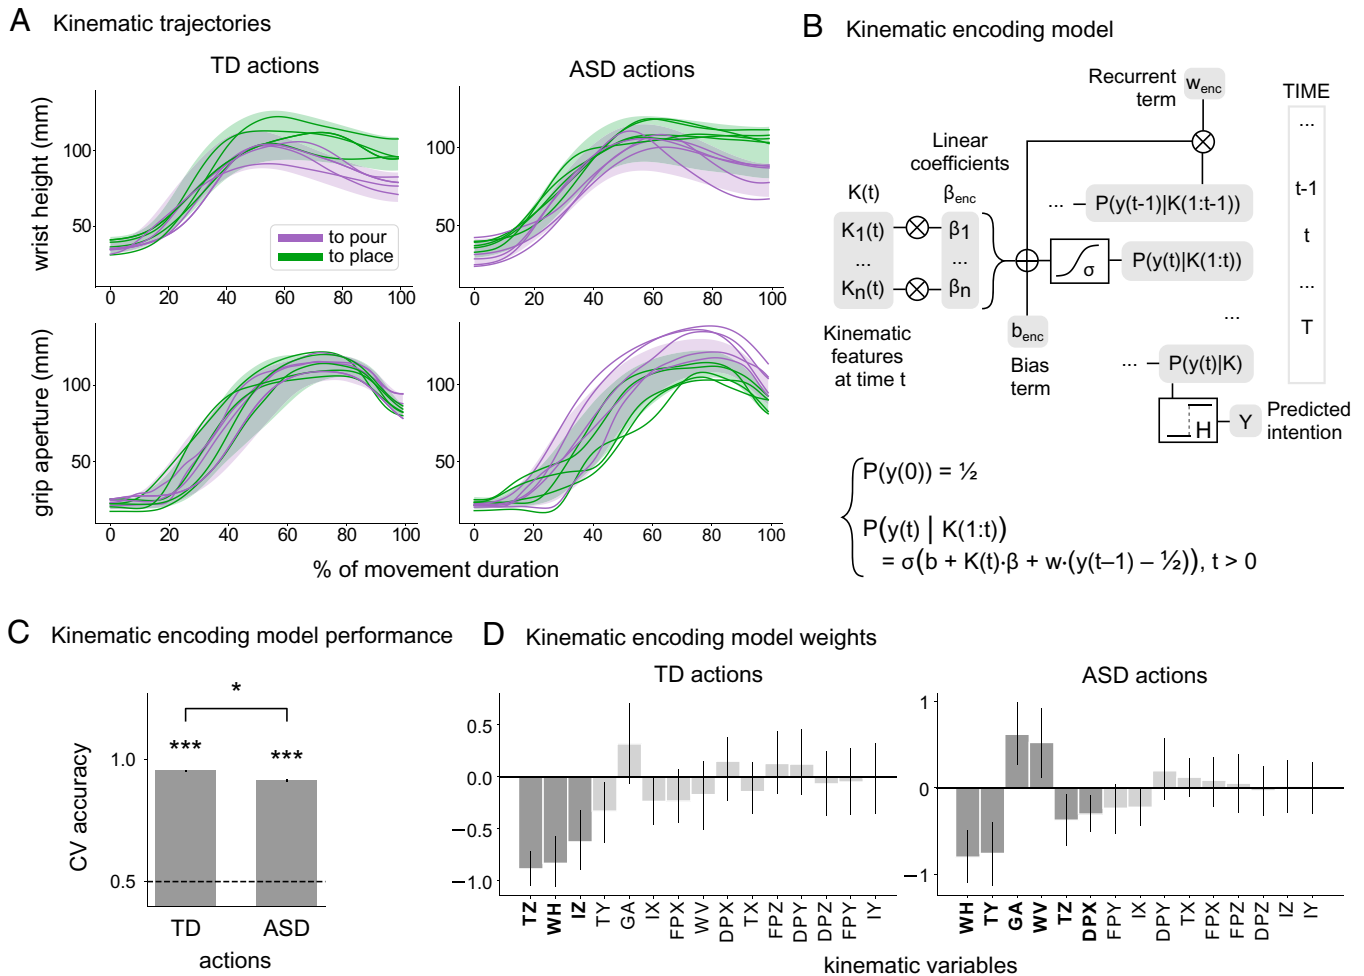


Fig. 2. Encoding of intention information in movement kinematics. (A) Time course of WH and GA for reach-to-grasp actions performed by TD and ASD children with the intention to pour or to place. Colored curves show representative trajectories for each intention, and colored areas show one SD across executed trials. (B) Block diagram and equation of the kinematic encoding model used to quantify intention information in movement kinematics. σ is the sigmoid function, β is the vector containing the values of the regression coefficients of each kinematic variable, and w is a coefficient weighting the accumulation of information over time. (C) Cross-validated (CV) performance of kinematic encoding models trained on TD actions and ASD actions quantified as the fraction of trial correctly predicted. Histograms represent mean \pm SEM across folds. $*P < 0.05$; $***P < 0.001$. (D) Contribution (weight) of each kinematic variable to the kinematic encoding of intention in TD and ASD movement kinematics as measured by the regression coefficient of the variable in the logistic regression. A positive (negative) weight is assigned to a variable distributed across trials with higher (lower) values for grasp-to-pour actions compared with grasp-to-place actions. Dark bars indicate variables carrying intention information. Error bars indicate 95% CIs computed by bootstrapping. DPY, y dorsum plane; DPZ, z dorsum plane; FPX, x finger plane; FPY, y finger plane; FPZ, z finger plane; IX, x index; IY, y index; TX, x thumb.

kinematics (hereinafter readers). To this end, for each observer, we computed the intention choices predicted by the kinematic readout model (regardless of whether the predicted choices were correct or incorrect) and compared the obtained value with a null distribution of randomly permuted choices. The distribution of readers and nonreaders in each group is shown in Fig. 4 as a function of intention discrimination performance and readout strength (defined as individual model performance, z scored with the null hypothesis accuracy on randomly permuted choices). For both TD and ASD observed actions, the proportion of readers was higher in the TD group (20 of 35) than in the ASD group (11 of 35). In both groups, the proportion of readers exceeded the proportion expected by chance for both TD actions and ASD actions, with no significant difference between observed actions (SI Appendix, Table S5).

Readers Good (and Bad) at Reading TD and ASD Actions. The notion of “reader” is agnostic with respect to intention discrimination performance—observers might read movement kinematics (as measured by kinematic readout model performance) and still

perform at chance or even below chance. For example, readers would perform at chance if they read variations that do not encode intention information; they would perform below chance if they read variations that encode intention information but do not read the encoded information correctly: for instance, they interpret a decrease in WH, encoding the intention to pour, as indicative of to place. To look at the relationship between readout and intention discrimination performance, we used a binomial test to stratify readers on the basis of their ability to discriminate intention (SI Appendix, Table S5).

As shown in Fig. 4, readers with intention discrimination above chance (“good readers”) in the TD group outnumbered good readers in the ASD group. In the TD group, the proportion of good readers was significantly higher for TD actions (10 observers of 14) compared with ASD actions (4 observers of 17). This increase was partially offset by the presence of three TD bad readers for TD actions. For both TD and ASD actions, the proportion of good readers was higher than expected by chance. In the ASD group, the proportion of good readers for both TD (one observer of five) and ASD actions (two observers




Fig. 3. Readout of intention from single-trial kinematics. (A) Block diagram and equation of the model used to quantify kinematic readout of intention. (B) Cross-validated performance of the kinematic readout models quantified as the fraction of correctly predicted intention choices. Histograms represent mean \pm SEM across participants for each observer group and observed action. The light subbars represent chance-level performance quantified as the mean of the null hypothesis distribution of cross-validated model performance. $**P < 0.01$; $***P < 0.001$. (C) Scatterplots of the relationship between the observed intention discrimination performance and the performance predicted by the kinematic readout model across individual participants for TD and ASD observers separately. Pearson's correlation coefficients (r) and their significance values (P) are reported. (D) Scatterplots of the relationship between trial-level reaction times and model prediction confidence computed as the deviation of the estimated probability of to pour from chance. Spearman's correlation coefficients (ρ) and their significance values (P) are reported.

of nine) did not differ from that expected by chance, with no difference between observed actions (*SI Appendix, Table S5*). Although the small sample of subgroups urges caution in interpretation, these results suggest a predominance of good readers among TD observing TD actions.

Intersecting Kinematic Encoding and Readout. Our results so far reveal differences in the ability to read out intention information across observers and observed actions. However, these analyses do not identify the specific features that are read out, whether readers read informative variables or noninformative variables, and how well they read the encoded information. To address this issue, we examined how specific features were read by TD and ASD readers during observation of TD and ASD actions.

We computed the contribution (weight) of each kinematic variable to the intention readout as the variable regression coefficient in the readout logistic regression. A positive (negative) weight is assigned to a variable distributed across trials with higher (lower) values for the intention choice to pour compared with to place. We examined, separately for each observer group and observed action, the overlap in the distribution of readout weights relative to encoding weights: whether readout weights were assigned to intention-informative variables. For variables carrying intention information, we also examined whether the signs of the readout weights correctly aligned with the signs of the encoding weights. A positive (negative) readout weight assigned to a variable with a positive (negative) encoding

weight would indicate correct alignment; a positive (negative) readout weight assigned to a variable with a negative (positive) encoding weight would indicate incorrect alignment. For example, an increase in WH encodes to place, and thus, WH is assigned a negative encoding weight (Fig. 2A and D). An incorrectly aligned positive readout weight would incorrectly interpret an increase in WH as signaling to pour.

Fig. 5A visualizes the overlap and alignment of readout weights relative to encoding weights across kinematic variables (averaged across observers). To provide a complementary visualization of the interindividual reproducibility of readout, Fig. 5B shows the number of readers who read a given variable in each condition. For TDs observing TD actions, comparison of the distribution of readout weights relative to encoding weights revealed a near-perfect overlap—the three variables that are read out more and by more observers (WH, TZ, and IZ) are also the three variables that encode intention information in TD kinematics (Fig. 5A). Filled bars indicate that the readout weights mostly aligned to the encoding weight correctly. Also, as shown in Fig. 5B, most observers correctly interpreted the intention information encoded in these variables. Although IZ does not carry intention information in ASD kinematics (Fig. 2D), WH, TZ, and IZ were also the three variables most frequently read by TD observers in ASD actions. Variables such as GA and WV, which encode intention information in ASD kinematics but not in TD kinematics, were only read out—mostly incorrectly—by a limited fraction of TD observers who observed ASD actions.

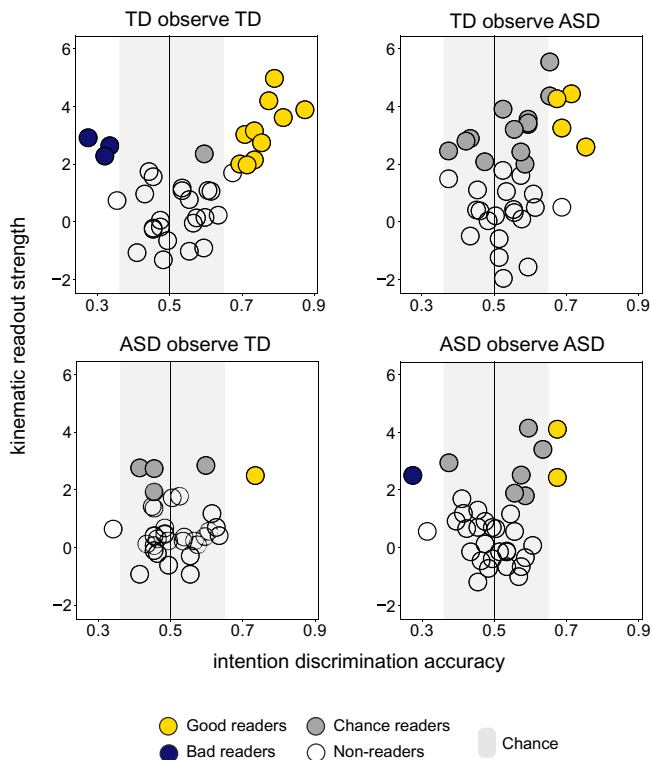


Fig. 4. Distribution of readers and nonreaders as a function of intention discrimination accuracy and readout strength. Readout strength is quantified for each observer as the z-scored readout model performance. The larger the readout strength, the stronger the sensitivity of intention readout to movement kinematics. Readers (observers whose intention readout is sensitive to single-trial movement kinematics) are represented as filled circles; good readers (intention discrimination accuracy above chance level) are represented as yellow circles, bad readers (intention discrimination accuracy below chance level) are represented as blue circles, and chance readers (intention discrimination accuracy at chance level) are represented as gray circles. Nonreaders are shown as open circles.

For ASD readers, the readout weights showed greater (although not perfect) overlap with the encoding weights during observation of ASD actions compared with TD actions. Specifically, ASD observers consistently read out two variables—WH and TZ—of the six variables encoding intention information in ASD actions. While WH and TZ also carry intention information in TD kinematics, ASD observers assigned little readout weight to these variables (or other informative variables) when observing TD actions. Diagonal striped bars indicate that, regardless of the observed actions (TD vs. ASD), information was misread in most variables.

We computed two indices that quantitatively summarized the above results across all variables. The first index quantified the overlap in the distribution of readout and encoding weights as the normalized scalar product between the absolute values of the encoding and readout vectors. The second index quantified the alignment of the readout weights relative to the encoding weights as the normalized scalar product between the encoding and readout vectors (Fig. 6A). A reader good who is good at both identifying informative features and interpreting their information would have both high overlap and high alignment. A reader good who is good at identifying informative features but not good at interpreting their information would have high overlap but low alignment. Consistent with the intuition conveyed by Fig. 5, the results showed a significant overlap in the distribution of readout and encoding weights for TD readers observing TD actions (but not ASD actions) and for ASD readers observing

ASD actions (but not TD actions). TD readers showed significant alignment across both TD and ASD actions. In contrast, alignment was not significant for ASD readers for either action (Fig. 6B and *SI Appendix*, Table S4).

Fig. 6C illustrates the correlation of overlap and alignment with individual intention discrimination performance separately for TD and ASD readers. Overlap did not correlate with individual intention discrimination. However, we found a significant positive correlation between overlap and deviation of individual intention discrimination performance from chance (defined as the absolute value of the difference between task performance and the 0.5-chance level). Alignment correlated positively with individual intention discrimination for both TD and ASD readers. This indicates that, in readers, individual intention discrimination depended not only on the selection of informative features but also, on their correct interpretation. In other words, the (in-)ability of readers to discriminate intentions was related to their (in-)ability to correctly interpret the intention information extracted from informative features.

Discussion

Many current perspectives on action reading in autism are based on the quantification of average intention discrimination across repeats of observed actions (3, 16). However, kinematics are variable across trials and individuals (20). Trial-averaged analyses may obscure how intention information is encoded in and read out in single-trial kinematics. Here, we have developed an analytic approach that enabled us to reveal intention readout computations with single-trial resolution.

By applying this approach, we were able to uncover that single-trial intention choices in ASD systematically reflected trial-to-trial variations in visual kinematics. This is demonstrated by the finding of a lower but still significant sensitivity of intention readout to single-trial kinematics (as measured by kinematic readout model performance) in ASD compared with TD. Corroborating this finding, the proportion of ASD observers who read trial-to-trial variations in movement kinematics (ASD readers, about one-third of ASD observers), although lower than the proportion of TD readers (about two-thirds of TD observers), exceeded the proportion of readers expected by chance for both TD and ASD actions. These findings indicate that while the average intention discrimination in ASD was at chance, single-trial intention choices by a sizeable proportion of individual observers were not random.

A second implication of our results is that for both TD and ASD readers kinematic similarity was important for identifying variations that carry intention-related information. Unlike in print reading, where all marks on paper encode meaning, in mind reading, readers must first extract, from trial-to-trial variations, those variations that encode intention information. Our single-trial results show that TD readers were able to extract such variations during observation of TD actions but not ASD actions. Conversely, ASD readers were able to extract intention-informative variations during the observation of ASD actions but not TD actions. This “same group” advantage is consistent with the principle that internal readout models (or codes) of TD observers are tuned to typical actions and internal readout models of ASD observers are tuned to autistic actions (9).

What are the exact tuning properties of typical and autistic models? Are internal readout models “feature based,” such that TD (ASD) readers assign more weight to those individual features that encode intention information in TD (ASD) movement kinematics? Or is visual kinematics more likely to be processed as a perceived whole, such that similar to face processing (21), changes in configural information (i.e., relationship between individual features) influence the identification of individual features?

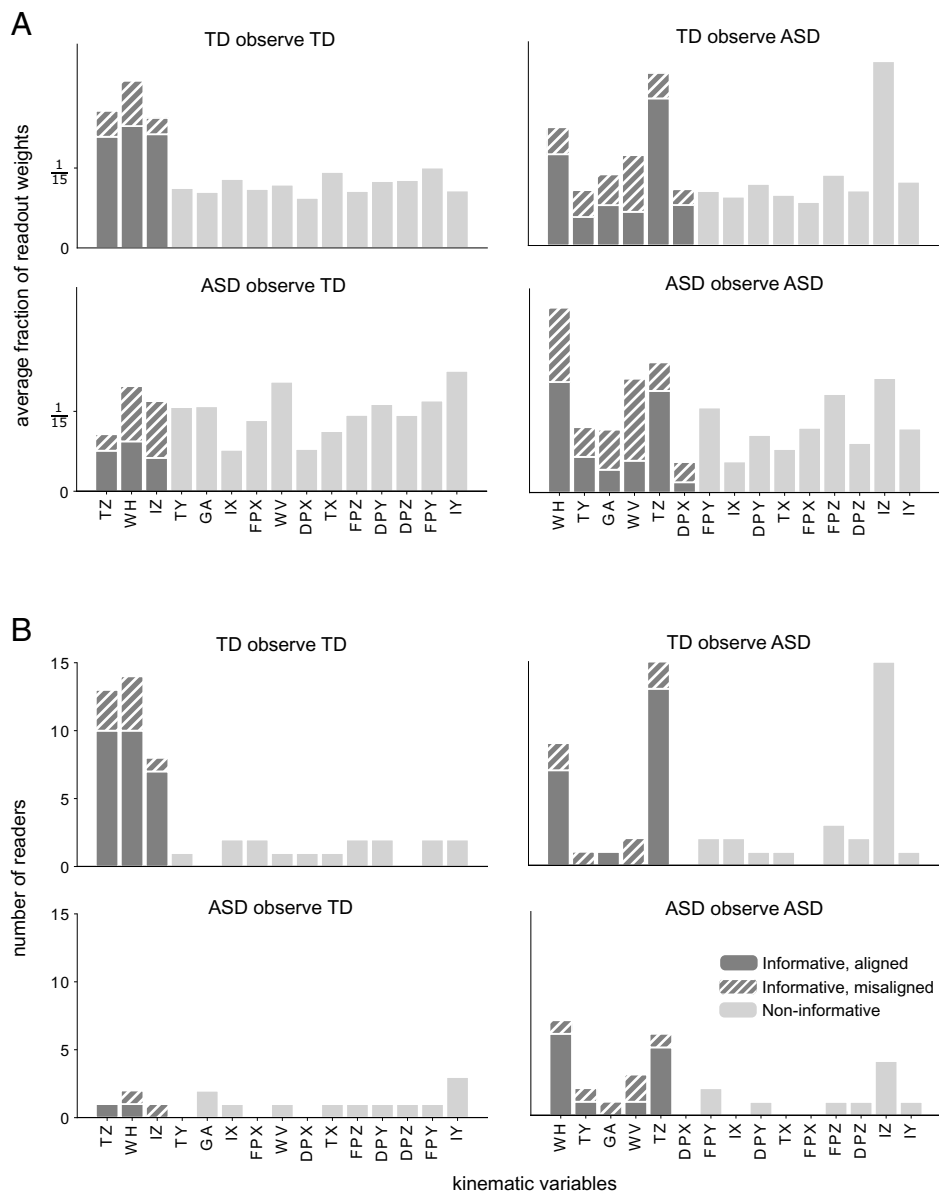


Fig. 5. Distribution and sign of readout weights across kinematic variables. (A) Bar graphs of the average fraction of kinematic readout weights across readers for each observer group and observed action. Kinematic variables are ordered from left to right from most informative to least informative. Variables carrying intention information are shown as dark bars. Filled bars indicate correct alignment of kinematic readout weights relative to encoding weights. Striped bars indicate incorrect alignment. (B) Bar plots of the number of readers who read each kinematic variable. A given variable is read by a reader if the corresponding readout weight is significantly different from zero. The order of variables and color coding are the same as in A. DPY, y dorsum plane; DPZ, z dorsum plane; FPX, x finger plane; FPY, y finger plane; FPZ, z finger plane; IX, x index; IY, y index; TX, x thumb.

Our kinematic readout model results provide an initial opportunity to answer these questions. If feature identification is integrated into the overall kinematic configuration, then TD intention-informative features should be weighted less when presented in the context of ASD visual kinematics than in TD visual kinematics. Conversely, ASD intention-informative features should be weighted less when presented in the context of TD visual kinematics compared with ASD visual kinematics. Consistent with this prediction, ASD readers weighted less ASD intention-informative features when observing TD actions compared with ASD actions. Configural effects in the TD readout were less clear. In contrast to ASD readers, TD readers appeared to weight TD intention-informative features equally in TD and ASD visual kinematics. In particular, IZ—a feature

that carries intention information in TD visual kinematics but not in ASD visual kinematics—was weighted similarly during observation of TD and ASD actions. Combined, these data may indicate a difference in the properties of TD and ASD internal readout models, with ASD internal models being more sensitive to the overall visual kinematics in which informative features are embedded.

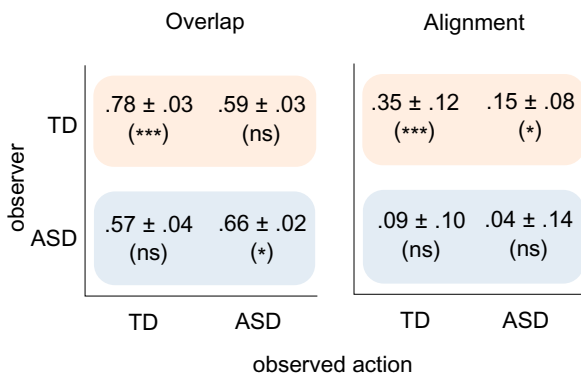
A third implication of our results is that, unlike TD readers, ASD readers lacked the ability to link kinematic variations to the correct intention. Interestingly, in both TD and ASD readers, (mis-)alignment of kinematic readout relative to kinematic encoding was comparable for TD and ASD visual kinematic, suggesting that, unlike overlap, alignment was little, if at all, affected by kinematic similarity. These data point to a selective

A Intersection measures

$$\text{overlap} = \frac{\text{abs}(\beta_{\text{read}}) \cdot \text{abs}(\beta_{\text{enc}})}{|\beta_{\text{read}}| |\beta_{\text{enc}}|}$$

$$\text{alignment} = \frac{\beta_{\text{read}} \cdot \beta_{\text{enc}}}{|\beta_{\text{read}}| |\beta_{\text{enc}}|}$$

B



C

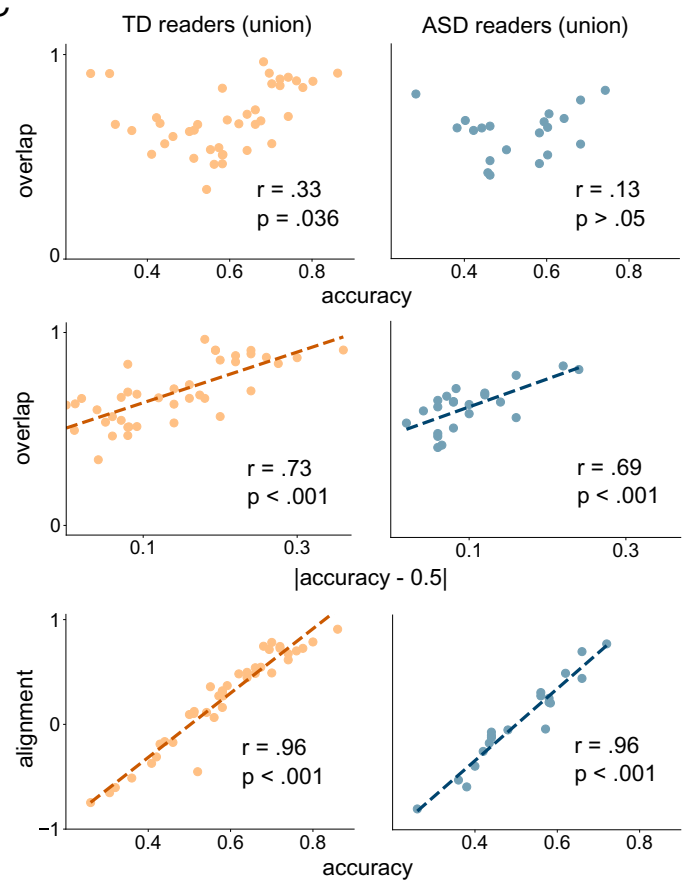


Fig. 6. Indices of the intersection between kinematic intention encoding and readout. (A) Overlap and alignment indices. (B) Average values of the overlap and alignment indices, reported as mean \pm SEM across readers, for each observer group and observed action. (C) Relationship between overlap and intention discrimination accuracy (Top), overlap and deviation of intention discrimination accuracy from the 0.5-chance level (Middle), and alignment and intention discrimination accuracy (Bottom). Pearson's correlation coefficients (r) and their significance values (P) are reported for significant linear relationships. ns indicates $P > 0.05$. * $P < 0.05$; *** $P < 0.001$.

impairment of ASD readers in interpreting informative variations in movement kinematics.

These results expand existing conceptions of mind reading in autism by pointing to distinct profiles of intention discrimination impairment in ASD observers. Some observers with ASD cannot read trial-to-trial variations in visual kinematics. Other observers with ASD, while reading trial-to-trial variations in movement kinematics, fail nevertheless to discriminate intention. Our single-trial results suggest that in this subtype of ASD readers, difficulties in mapping visual kinematics to intention may reflect both an interaction failure and an individual failure. The interaction failure manifests in poor identification of intention-informative features in TD visual kinematics by ASD readers and conversely, in poor identification of intention-informative features in ASD kinematics by TD readers, as measured by overlap. The individual failure manifests in poor interpretation of the extracted information specific to ASD readers. That is, while TD readers are generally able to link intention-informative variations in movement kinematics to the correct intention, ASD readers are unable to do so, regardless of whether the information is extracted from TD or from ASD visual kinematics.

In this study, we developed an experimental and analytic framework to decompose the process components of intention to action attribution and to investigate how intention encoding and readout intersect in TD and ASD observers who observe TD and ASD actions. This framework forms a powerful, general approach to test how information is encoded and read out in movement kinematics at the single-trial, single-subject level.

In the present study, we asked participants to simply judge the intention of the observed actions. An important direction for future research will be to investigate intention readout during active participation in social interaction: specifically, whether different patterns of readout emerge when individuals are asked not only to observe but also, to respond to the actions of others (15, 22). Moreover, by decomposing the component process of intention reading, our approach could be useful for identifying targets for intervention. There is evidence that TD observers can be explicitly guided to attend to potentially diagnostic features in visual kinematics (23). Based on the findings of the current study, a promising direction will be to investigate whether tutoring (either explicit or implicit) can promote alignment in observers with autism.

Materials and Methods

The research protocol was approved by the local ethics committee (ASL3 Genova) and complied with the principles of the revised Helsinki Declaration (24). Written informed consent was obtained from the parents of the children prior to participation in the experiment.

Participants. We report the results of 35 ASD children (29 males) without accompanying intellectual impairment and 35 TD children (29 males). Groups were matched for gender, age [TD mean \pm SD = 9.8 \pm 1.1 y; ASD mean \pm SD = 10.2 \pm 1.4 y; $t_{(68)} = -1.345$, $P = 0.183$], and full-scale IQ as measured by the Wechsler Scale of Intelligence (25) [TD mean \pm SD = 103.9 \pm 10.2; ASD mean \pm SD = 99.5 \pm 11.1; $t_{(68)} = 1.749$, $P = 0.085$]. Children with ASD were diagnosed according to the criteria of the Diagnostic and Statistical Manual of Mental Disorders (DSM-5) (26). The Autism Diagnostic Observation Scale (27) and the

Autism Diagnostic Interview-Revised (28) were administered by two experienced professionals. Autistic traits were assessed in 19 ASD children and 16 TD children using the SRS (29) and were more prevalent in ASD compared with TD children (TD mean \pm SD = 50.0 \pm 10.1; ASD mean \pm SD = 83.0 \pm 15.6; $P < 0.001$), with some overlap between the two groups in the moderate-range score (*SI Appendix, Fig. S3A*). All children had normal or corrected-to-normal vision and were screened for exclusion criteria (pharmacological treatment, epilepsy, and any other neurological and psychiatric conditions). All but three of the children (two in the ASD group and one in the TD group) were right handed according to the Edinburgh Handedness Inventory (30).

Experimental Design and Procedures.

Stimuli. Stimuli were selected from a dataset of 940 grasping actions obtained by recording 20 TD and 20 ASD children performing grasp-to-pour and grasp-to-place actions. For grasp-to-pour trials, a glass (height = 10 cm; diameter = 6.5 cm) was placed 19 cm from the bottle. Participants were instructed to reach for the bottle, lift it, and pour some water into the glass. A coexperimenter refilled the bottle on each trial. For grasp-to-place trials, a box (height = 6 cm; diameter = 10 cm) was placed 19 cm from the bottle. Participants were instructed to reach for the bottle, lift it, and place it in the box. Detailed procedures and the apparatus are described in ref. 13. Briefly, reach-to-grasp movements were tracked using a near-infrared camera motion capture system with six optical cameras (frame rate, 100 Hz; Vicon System) and simultaneously filmed from a lateral viewpoint using a video camera fully synchronized with the optical cameras (Vicon Vue; 100 frames/s, resolution 1,280 \times 720). As previously described (13), the child's right hand was outfitted with retroreflective hemispheric markers (6.5 mm in diameter) placed on the metacarpal joint and the tip of the index and the little finger, the trapezium bone and the tip of the thumb, the radial aspect of the wrist, and the hand dorsum. This marker set allowed us to finely track both the distal (hand shape) and proximal (transport) components of the motions. The kinematic data were run through a 6-Hz low-pass Butterworth filter. Based on previous studies investigating prospective action control in TD and ASD children (12, 13), we extracted 15 kinematic variables (*SI Appendix, Fig. S2*):

- WV defined as the module of the velocity of the wrist marker (millimeters per second);
- WH defined as the z component of the wrist marker (millimeters);
- GA defined as the distance between the marker placed on the thumb tip and the one placed on the tip of the index finger (millimeters);
- x thumb, TY, and TZ defined as x, y, and z coordinates of the tip of the thumb (millimeters);
- x index, y index, and IZ defined as x, y, and z coordinates of the tip of the index (millimeters);
- x, y, and z finger plane defined as x, y, and z components of the thumb-index plane (i.e., the three-dimensional components of the vector that is orthogonal to the plane; this plane provides information about the abduction/adduction movement of the thumb and index finger independent of the effects of wrist rotation and of finger flexion/extension); and
- DPX, y, and z dorsum plane defined as x, y, and z components of the radius-phalanx plane (this plane provides information about the abduction, adduction, and rotation of the hand dorsum independent of the effects of wrist rotation).

Custom software (MATLAB; MathWorks Inc.) was used to extract the selected variables. Each variable was calculated at intervals of 10% of the movement duration from reach onset to reach offset.

Selection of action stimuli. From the dataset of grasping actions, we selected, for each group, 50 grasping actions (grasp to place, $n = 25$; grasp to pour, $n = 25$) according to the following criteria: 1) minimized within-intention distance (using the metric reported in ref. 19) and 2) mean duration of movements not significantly different between intentions. Video clips that corresponded to the selected reach-to-grasp actions were used as stimuli for the intention discrimination task. Each video clip began with reach onset and ended with the contact between the hand and the bottle. To allow participants sufficient time to focus on the reach onset, static frames ranging in duration from 160 to 800 ms (in 160-ms increments) were randomly added to the beginning of each video.

Intention discrimination task. Participants were seated in front of a 24-inch computer monitor (resolution 1,280 \times 720; 100 Hz) at a viewing distance of 50 cm. The task structure conformed to a one-interval forced choice task with binary choice (to place vs. to pour). Each trial began with the presentation of a white central fixation cross for 1,000 ms. Then, a video clip showing the reach-to-grasp action was presented. After the video (followed by a waiting window of 80 ms), a screen prompted participants to indicate the action (to place or to pour) that would follow the observed grasp (5,000 ms). For half of

the participants, the Italian word “mettere” (to place) on the left prompted a button press with the index finger of the left hand, and the word “versare” (to pour) on the right prompted a button press with the index finger of the right hand. The position of the two words was counterbalanced across participants. Participants completed two sessions in which they observed reach-to-grasp actions performed by TD children and ASD children in counterbalanced order. Each session consisted of 50 experimental trials performed in three blocks (10, 20, and 20 trials), with a 2-min break between each block. Participants received no feedback either during the experimental blocks or during the practice block.

The task was revised and administered by a clinically experienced experimenter. During the experiment, the experimenter positioned herself behind the child. The experimenter was the same for all participants. Participants were introduced to the stimuli and were given both written and verbal instructions. Practice trials ($n = 20$) were included before the experimental session to familiarize the child with the task and ensure that they had understood the task. Participants were instructed to respond as accurately and quickly as possible during the presentation of the prompt screen. If they did not comply with the instructions (e.g., they responded during video presentation), the experimenter would ask them to repeat practice trials. We verified that the number of participants who repeated the practice did not differ between groups (six in the ASD group and six in the TD group). The proportion of responses during video presentation or within the first 100 ms of the response window was generally low (TD observers, mean \pm SEM = 0.021 \pm 0.007; ASD observers, mean \pm SEM = 0.015 \pm 0.004) and did not differ between groups [$t_{(68)} = 0.752$, $P = 0.455$]. Stimuli presentation, timing, and randomization were controlled using E-prime V2.0 software (Psychology Software Tools).

Eye movement data. Gaze direction was measured with an infrared eye tracker (SMI RED500; SensoMotoric Instruments). The eye tracker suffered a fatal technical failure before testing was completed; moreover, calibration of the eye tracker was unsuccessful in some participants. Therefore, eye-tracking data are available for 35 TD participants and 23 ASD participants. For each observer and each trial, we extracted the sequence of spatial position coordinates (scan path). We computed the fraction of time in which the scan path was within the screen in each trial and then, averaged this value across trials for each observer. This fraction was overall high (TD observers, mean \pm SEM = 0.91 \pm 0.01; ASD observers, mean \pm SEM = 0.88 \pm 0.03) and did not differ between groups [$t_{(56)} = 0.860$, $P = 0.393$].

Quantification and Statistical Analysis.

Data preprocessing. Trials for which participants provided a response during the waiting window or within the first 100 ms of the response window were discarded from analyses (<2% of trials). We verified that the pattern of results and their significance remained similar even when all trials were included.

Mixed effects models to assess statistical differences in intention discrimination performance, response bias, and model performance. We used mixed effects models to assess the significance of differences in intention discrimination performance (Fig. 1D), response bias (*SI Appendix, Fig. S1B*), and encoding and readout model performance (Figs. 2C and 3B, respectively) compared with chance and across observer groups and observed actions. We used logistic regression with single-trial accuracy and response as the dependent variable to assess differences in intention discrimination performance and response bias and linear regression, with the fraction of correct predictions of each video across cross-validation repetitions as the dependent variable, to assess differences in encoding and readout model performance. The chance-level null hypothesis distribution for encoding and readout model performance was created by fitting the model after randomly permuting across trials the observer's choice labels.

To determine the fixed and random effects to include in the model, we applied a model selection procedure that started from the model with the most complex structure to arrive at a model that included only the significant predictors. We first selected the random effects structure of the model by keeping the full fixed effects structure and using the Bayesian Information Criterion (BIC) (31). The BIC rewards model fit and penalizes model complexity. We then retained the optimal random effects structure and selected the best fixed effects structure by conducting likelihood ratio tests between models differing only by the presence or absence of one predictor (32). Model selection results are reported in *SI Appendix, Table S1*. CIs for model coefficients and statistical comparisons for the effects are reported in *SI Appendix, Table S2*. We performed model fitting using the R package lme4 (33). We performed comparisons across levels of the selected models using the glht command from the R package multcomp (34). The multcomp package estimates the value and SE of each effect, from which a z value (to calculate two-sided P values) is computed. The results are reported in *SI Appendix, Table S2* along

with CIs for the estimates of the regression coefficients and for the SD of random effects of the selected models, computed using the bootstrap option in the R function `confint`.

Quantifying and assessing the significance of individual task performance. Because there was no significant response bias, we quantified individual performance on the intention discrimination task as the fraction of correct intention choices. For each participant, we assessed the significance of discrimination performance against chance using a binomial test separately for TD observed actions and ASD observed actions (SI Appendix, Table S5).

Single-trial kinematic vector. To quantify single-trial kinematics, the 15 kinematic variables of interest were averaged, for each grasping action, over 10 epochs (t), each spanning 10% of the normalized movement time (0 to 10, 10 to 20%, etc., of the movement duration from reach onset to reach offset). Next, for each trial and epoch, we created a 15-dimensional, single-trial time-dependent kinematic vector, $\vec{K}(t)$, whose entries, for each trial, were the 15 kinematic variables averaged over that time epoch. We used this kinematic vector for all logistic regressions (see below). We verified that increasing the number of kinematic features by considering the x , y , and z component of all the markers used to compute the kinematic variables of interest (3×6 retroreflective markers) did not improve the performance of the kinematic encoding model. This observation held true for both TD and ASD actions and even when using a finer time windowing (25 or 50 movement epochs rather 10 movement epochs, as in the analyses reported in the main text; $P < 0.02$ for all movement epochs number). These control analyses suggest that our kinematic encoding model provided adequate spatial and temporal resolution to capture intention-related variations in TD and ASD kinematics.

Logistic regression models of kinematic encoding and readout. To determine the dependence of intention (kinematic encoding model) and intention choice (kinematic readout model) on kinematics over time, we used a logistic regression to estimate the single-trial cumulative probability $y(t)$ (i.e., the cumulative evidence) in favor of the intention to place as function of the time-dependent kinematic vector in that trial up to time t . Specifically, we modeled $y(t)$ as a sigmoid transformation of the sum of two terms: a linear transformation of the kinematic vector $\vec{K}(t)$, describing the evidence provided by the single-trial kinematic vector at the current time epoch (t), and a drift term, describing the contribution of the cumulated evidence $y(t-1)$ provided by the kinematic vectors up to the previous time epoch ($t-1$). More precisely, the equation of the logistic model was as follows:

$$P([y(0)=\text{'to pour'}]) = P([y(0)=\text{'to place'}]) = \frac{1}{2}$$

$$P([y(t)=\text{'to pour'}] | \vec{K}(1), \dots, \vec{K}(t)) = \sigma\left(\vec{K}(t) \cdot \vec{\beta} + w \cdot \left(y(t-1) - \frac{1}{2}\right) + \beta_0\right)$$

$$P([y(t)=\text{'to place'}] | \vec{K}(1), \dots, \vec{K}(t)) = 1 - P([y(t)=\text{'to pour'}] | \vec{K}(1), \dots, \vec{K}(t)),$$

where σ is the sigmoid function, $\vec{\beta}$ is the vector containing the values of the regression coefficients of each kinematic variable, w is a coefficient weighting the accumulation of information over time, and β_0 is a kinematic-independent bias term. The value of $y(t)$ computed at reach offset provides the final probability of intention (kinematic encoding model) or intention choice (kinematic readout model) associated with the kinematics of the whole trial. In this model, a single regression coefficient is assigned to each variable, meaning that the contribution of each kinematic variable is weighted equally across all time epochs. More complex models with different regression weights assigned to each variable at different epochs as in ref. 6 yielded no better performance ($P > 0.09$ for all observer groups and observed actions), confirming that, despite its simplicity, our model fit well both intention encoding and readout.

Training logistic regression models. Training and evaluation were performed in a similar manner for encoding and readout models. Each model was trained on a set of 50 trials. We z scored the single-trial kinematic vectors within each model to avoid penalizing predictors with larger value ranges. We trained the models by minimizing the negative binomial log likelihood with L^2 penalty via stochastic gradient descent with adaptive moment estimation (Adam) (35). The training was marginally improved by a data augmentation scheme based on small random deformations over the time dimension (SI Appendix, Data Augmentation Procedure for Training the Logistic Regressions has full details). The parameter λ , which controls the strength of the L^2 regularization term, was set to 0.05 for all models. A cross-validation approach for tuning this hyperparameter was also tested and yielded similar results. The kinematic encoding and readout models were implemented using Python/PyTorch (36).

Kinematic encoding model. The kinematic encoding model expressed the probability that a grasping action was performed with the intent to pour as a function of the kinematic vector of that action. We trained separate encoding models for TD and ASD actions. We used the encoding model to quantify the intention information encoded in movement kinematics (Fig. 2C) and to

identify the kinematic variables that carry intention information in TD and ASD movement kinematics (Fig. 2D).

Kinematic readout model. The kinematic readout model expressed the probability of intention choice in each trial as a function of the kinematic vector measured in that trial. We trained the readout model separately for each observer in each session.

Evaluation of model performance. We evaluated the performance of the encoding and readout models by repeated fivefold cross-validation (50 random splits) (37). We computed the most likely value of Y for each trial by taking the argmax over Y of $P(Y | \vec{K})$ in the equation of the logistic model. Model performance was computed as the fraction of correct trials averaged over folds and random splits.

Statistics on the proportion of readers. We used a binomial test to establish whether the number of readers and the fraction of good readers were statistically significant in each group. To assess the significance of differences between observer groups and observed actions in the proportion of readers and in the proportion of good readers, we used a nonparametric permutation test.

Estimate of CIs of model coefficients. For all kinematic encoding and readout models, we obtained estimates and 95% CIs for the regression coefficients from a bootstrap distribution obtained by fitting the models to data randomly sampled with replacement from the original training data.

Classification of individual kinematic variables as informative for encoding or readout. We assessed the informativeness of individual variables (Fig. 2D) by testing whether the corresponding encoding coefficients were significantly different from zero. We retained as informative those variables whose encoding coefficients (absolute value) were greater than the 95th percentile of the null hypothesis values obtained when training the kinematic encoding models with permuted trial labels. A similar procedure was used to determine the number of observers who read each variable during action observation (shown in Fig. 5B).

Computation of discrimination performance and confidence predicted by the kinematic readout model. In Fig. 3C, we used the kinematic readout model to estimate the intention discrimination performance of individual participants. Using the equation of the logistic model, the intention choice predicted as most likely by the readout model was computed for each trial and compared with the actual intention choice. The individual intention discrimination performance was obtained by averaging the probability of correct choice across all trials for a given participant. For the analysis in Fig. 3D, we computed the confidence of the single-trial model prediction as the deviation of the estimated probability of to pour from chance (0.5).

Computation of overlap and alignment. We computed two indices of intersection between encoding and readout: overlap and alignment. The index quantifying the overlap in kinematic space between encoding and readout was computed by taking the elementwise absolute value of β_{enc} and β_{read} and computing the normalized scalar product of the resulting vectors:

$$overlap(\beta_{enc}, \beta_{read}) = \frac{(\text{abs}(\beta_{enc}), \text{abs}(\beta_{read}))}{\|\beta_{enc}\| \|\beta_{read}\|} \in [0, 1].$$

The overlap index measures the amount of weight common to the two vectors, regardless of the sign of the coefficients.

The index quantifying the alignment of encoding and readout in kinematic space was computed as the normalized scalar product between the encoding and readout vectors:

$$alignment(\beta_{enc}, \beta_{read}) = \frac{(\beta_{enc}, \beta_{read})}{\|\beta_{enc}\| \|\beta_{read}\|} \in [-1, 1].$$

Note that the absolute value of the alignment index is bounded from above by the value of overlap. Alignment values close to zero can be found either with low overlap values (when the variables with nonzero weights differ between encoding and readout) or with high overlap values (when the two models select the same variables with nonzero weights but the signs of the weights are inconsistent).

In Fig. 6 and in SI Appendix, Table S4, the statistics of the overlap index and the alignment index were computed over the set of observers who were classified as readers when observing either TD or ASD actions.

Permutation test to assess the significance of overlap and alignment. To assess the significance of the overlap and alignment indices, we compared them with a null hypothesis distribution obtained by recomputing their values after random permutation ($n = 10^5$ random permutations) of the entries of the encoding vectors.

Conventions for P values. SI Appendix, Tables S1–S3 report details of logistic mixed effects models statistical tests, linear mixed effects models statistical tests, and nonparametric permutation tests. Reported P values are two sided and Holm–Bonferroni corrected. In the figures, * $P < 0.05$, ** $P < 0.01$, *** $P < 0.001$, and ns indicates $P > 0.05$. Following standard notation, asterisks above bars indicate significance of difference from chance of an individual quantity,

and asterisks above brackets indicate significance of difference between two quantities.

Statistical significance of correlations. The significance of correlation values was assessed using the `scipy.stats` Python module, with two-sided parametric Student statistics for Pearson correlation and two-sided permutation distribution for Spearman correlation (38). We used the `SciPy` package (39). Significance values are shown in Fig. 3D.

Data Availability. The code supporting the main results of this study is described in *Materials and Methods* and has been deposited in GitHub

(https://github.com/noemimontobbio/ASD_encoding_readout). Additional data are included in *Dataset S1*. This study did not generate new unique reagents or materials.

ACKNOWLEDGMENTS. We thank all participants and their families for their efforts to participate in the study. This research was supported by EnTimeMent EU H2020 FETPROACT Grant 824160 and by NIH BRAIN Initiative Grant R01NS109961. DiNOGMI contributed to this work within the framework of the DiNOGMI Department of Excellence MIUR 2018 to 2022 (legge 232/2016).

1. C. M. Heyes, C. D. Frith, The cultural evolution of mind reading. *Science* **344**, 1243091–1243091 (2014).
2. C. Becchio, A. Koul, C. Ansuini, C. Bertone, A. Cavallo, Seeing mental states: An experimental strategy for measuring the observability of other minds. *Phys. Life Rev.* **24**, 67–80 (2018).
3. R. Edey *et al.*, Interaction takes two: Typical adults exhibit mind-blindness towards those with autism spectrum disorder. *J. Abnorm. Psychol.* **125**, 879–885 (2016).
4. F. Castelli, C. Frith, F. Happé, U. Frith, Autism, Asperger syndrome and brain mechanisms for the attribution of mental states to animated shapes. *Brain* **125**, 1839–1849 (2002).
5. S. Boria *et al.*, Intention understanding in autism. *PLoS One* **4**, e5596 (2009).
6. J.-F. Patri *et al.*, Transient disruption of the inferior parietal lobule impairs the ability to attribute intention to action. *Curr. Biol.* **30**, 4594–4605.e7 (2020).
7. S. Van de Cruys *et al.*, Precise minds in uncertain worlds: Predictive coding in autism. *Psychol. Rev.* **121**, 649–675 (2014).
8. R. P. Lawson, C. Mathys, G. Rees, Adults with autism overestimate the volatility of the sensory environment. *Nat. Neurosci.* **20**, 1293–1299 (2017).
9. J. Cook, From movement kinematics to social cognition: The case of autism. *Philos. Trans. R. Soc. Lond. B Biol. Sci.* **371**, 20150372 (2016).
10. B. Hommel, J. Müsseler, G. Aschersleben, W. Prinz, The Theory of Event Coding (TEC): A framework for perception and action planning. *Behav. Brain Sci.* **24**, 849–878 (2001).
11. D. M. Wolpert, K. Doya, M. Kawato, A unifying computational framework for motor control and social interaction. *Philos. Trans. R. Soc. Lond. B Biol. Sci.* **358**, 593–602 (2003).
12. A. Cavallo *et al.*, Prospective motor control obeys to idiosyncratic strategies in autism. *Sci. Rep.* **8**, 13717 (2018).
13. A. Cavallo *et al.*, Identifying the signature of prospective motor control in children with autism. *Sci. Rep.* **11**, 3165 (2021).
14. J. L. Cook, S.-J. Blakemore, C. Press, Atypical basic movement kinematics in autism spectrum conditions. *Brain* **136**, 2816–2824 (2013).
15. L. Schilbach, Towards a second-person neuropsychiatry. *Philos. Trans. R. Soc. Lond. B Biol. Sci.* **371**, 20150081 (2016).
16. L. Casartelli *et al.*, Neurotypical individuals fail to understand action vitality form in children with autism spectrum disorder. *Proc. Natl. Acad. Sci. U.S.A.* **117**, 27712–27718 (2020).
17. S. Panzeri, C. D. Harvey, E. Piasini, P. E. Latham, T. Fellin, Cracking the neural code for sensory perception by combining statistics, intervention, and behavior. *Neuron* **93**, 491–507 (2017).
18. M. Valente *et al.*, Correlations enhance the behavioral readout of neural population activity in association cortex. *Nat. Neurosci.* **24**, 975–986 (2021).
19. A. Cavallo, A. Koul, C. Ansuini, F. Capozzi, C. Becchio, Decoding intentions from movement kinematics. *Sci. Rep.* **6**, 37036 (2016).
20. M. L. Latash, The bliss (not the problem) of motor abundance (not redundancy). *Exp. Brain Res.* **217**, 1–5 (2012).
21. G. Rhodes, A. Calder, M. Johnson, J. V. Haxby, Eds., *Oxford Handbook of Face Perception* (Oxford University Press, Oxford, United Kingdom, 2012).
22. L. Schilbach *et al.*, Toward a second-person neuroscience. *Behav. Brain Sci.* **36**, 393–414 (2013).
23. M. L. Slepian, S. G. Young, A. M. Rutchick, N. Ambady, Quality of professional players' poker hands is perceived accurately from arm motions. *Psychol. Sci.* **24**, 2335–2338 (2013).
24. World Medical Association, World Medical Association Declaration of Helsinki: Ethical principles for medical research involving human subjects. *JAMA* **310**, 2191–2194 (2013).
25. D. Wechsler, *Wechsler Intelligence Scale for Children* (PsychCorp, San Antonio, TX, ed. 4, 2003).
26. American Psychiatric Association, *Diagnostic and Statistical Manual of Mental Disorders* (American Psychiatric Association, Arlington, VA, ed. 5, 2013).
27. C. Lord, R. Luyster, K. Gotham, W. Guthrie, *Autism Diagnostic Observation Schedule, Second Edition (ADOS-2) Manual (Part I): Modules 1-4* (Western Psychological Services, Torrance, CA, 2012).
28. M. Rutter, A. Le Couteur, C. Lord, *The Autism Diagnostic Interview-Revised (ADI-R)* (Western Psychological Services, Los Angeles, CA, 2003).
29. J. Constantino, J. Gruber, *Social Responsiveness Scale (SRS) Manual* (Western Psychological Services, Los Angeles, CA, 2005).
30. R. C. Oldfield, The assessment and analysis of handedness: The Edinburgh inventory. *Neuropsychologia* **9**, 97–113 (1971).
31. G. Schwarz, Estimating the dimension of a model. *Ann. Stat.* **6**, 461–464 (1978).
32. A. Agresti, *An Introduction to the Categorical Data Analysis* (John Wiley & Sons, Inc., Hoboken, NJ, 2007).
33. D. Bates, M. Machler, B. M. Bolker, S. C. Walker, Fitting linear mixed-effects models using lme4. *J. Stat. Softw.* **67**, 1–48 (2015).
34. T. Hothorn, F. Bretz, P. Westfall, Simultaneous inference in general parametric models. *Biom. J.* **50**, 346–363 (2008).
35. D. P. Kingma, J. L. Ba, "Adam: A method for stochastic optimization" in *Proceedings of the 3rd International Conference on Learning Representations (ICLR)*, Y. Bengio, Y. LeCun, Eds., San Diego, CA, USA, 2015, Conference Track Proceedings (ICLR, 2015), pp. 1–15.
36. A. Paszke *et al.*, "PyTorch: An imperative style, high-performance deep learning library" in *Advances in Neural Information Processing Systems*, H. Wallach *et al.*, Eds. (Chapman & Hall, 2019), vol. 32, pp. 8024–8035.
37. J.-H. Kim, Estimating classification error rate: Repeated cross-validation, repeated hold-out and bootstrap. *Comput. Stat. Data Anal.* **53**, 3735–3745 (2009).
38. D. Best, D. Roberts, Algorithm AS 89: The upper tail probabilities of Spearman's rho. *J. R. Stat. Soc. Ser. C Appl. Stat.* **24**, 377–379 (1975).
39. P. Virtanen *et al.*, SciPy 1.0 Contributors, SciPy 1.0: Fundamental algorithms for scientific computing in Python. *Nat. Methods* **17**, 261–272 (2020).