Where does the flow go? Humans automatically predict liquid pathing with coarse-grained simulation

Yuting Zhang, Mario Belledonne, Tristan Yates, Ilker Yildirim
{yuting.zhang, mario.belledonne, tristan.yates, ilker.yildirim}@yale.edu
Department of Psychology, Yale University
New Haven, CT 06511 USA

Abstract

Bodies of water manifest rich physical interactions via non-linear dynamics. Yet, humans can successfully perceive and negotiate such systems in everyday life. Here, we hypothesize that liquid bodies play such an integral role in human life that the mind automatically computes their approximate flow-paths, with attention dynamically deployed to efficiently predict flow trajectories using coarse mental simulation. When viewing animations of liquids flowing through maze-like scenes, we asked participants to detect temporary slowdowns embedded in these animations. This task, without any overt prompt of path or prediction, reveals that detection rates vary with the moment-to-moment changes in coarse flow-path predictions. Critically, coarse predictions better explain trial-level detection rates than a finer-grained alternative, independently of bottom-up salience of slowdowns. This work suggests liquid flow-path prediction as an implicit task in the mind, and introduces rich attentional dynamics as a new window into intuitive physics computations.

Keywords: intuitive physics; prediction; attention; liquid perception; mental simulation

Introduction

As a body of water splashes and flows, complex and dynamic interactions occur between the liquid and the rigid surfaces it collides with. Despite this, humans have no trouble perceiving liquids (Van Assen, Barla, & Fleming, 2018) or interacting with them. How can the mind keep up with a system as dynamic as liquids during perception or planning? A common computational principle, which is used to both explain the nature of perception (Clark, 2013) and to build artifacts for controlling complex systems (Holkar & Waghmare, 2010), is prediction. Often these predictions need not to be precise (Clark, 2015) – approximate predictions can guide adaptive processing (Gerstenberg, Peterson, Goodman, Lagnado, & Tenenbaum, 2017) and inform an initial plan in a controller (Tassa, Erez, & Todorov, 2012).

Here, we ask: Does the mind, during visual processing of liquid flow, make spontaneous coarse-grained predictions to estimate the overall trajectory the liquid will take? Given the importance of liquid bodies in human life (and in biology more generally) and the peril it can cause to miscalculate the path such a body of liquid will take, it is plausible that the mind makes such predictions automatically. We hypothesize that the mind makes spontaneous flow-path predictions, and how these predictions change at small time increments drive attentional dynamics; e.g., the higher the momentary changes in the estimated flow-path, the more attention is to be deployed.

We consider perception of liquid flow through maze-like configurations of rigid surfaces and introduce a new performance-based psychophysics task to measure behavioral signals related to intuitive physics computations. Instead of a task that overtly queries intuitive physics, such as prediction or reasoning, we implement a new task, inspired by recent work (Yasuda, Yates, & Yildirim, 2021), in which prediction or path of the liquid is not necessarily relevant nor explicitly mentioned. This task measures the fine temporal dynamics of attentional deployment during spontaneous visual processing of video stimuli, which we posit provides a new window onto intuitive physics related computations in the mind.

In particular, we simulated nine brief naturalistic animations of liquids interacting with 25 planks in a 5-by-5 maze formation (Fig. 1A). The liquid flowed from the top of the planks. We transiently (120 ms) slowed down each animation at one of 18 different time points and recruited participants (N=30 per scene and time point combination) to press a key when they noticed a pause, with the idea that detection accuracy would reflect the degree of attentional deployment.

Crucially, this task does not overtly involve intuitive physics nor mentions anything specific about the perception of liquid flow – it merely asks participants the tangential task of noticing a temporary slow-down in the video. We compared the temporal patterns of human's detection rates to four different covariates derived from either predictive or purely sensory computational mechanisms. We found that coarse-predictions about the liquid flow-path explains many subtleties of human detection rates, and does so when accounting for a hybrid prediction alternative as well as bottom-up factors (pixels and the later layers of a pretrained deep convolutional neural network).

Our study is inspired by and builds on the work of Bates, Yildirim, Tenenbaum, and Battaglia (2019). They used prediction tasks (e.g., "which bin will have most liquid") and computational modeling to study human intuitions of how liquids flow in similarly complex scenes as ours. Building on the earlier classic work by Battaglia, Hamrick, and Tenenbaum (2013), they provided evidence for a dynamics-based, simulatable representation of liquids as the driver of these predictions. Here, unlike the explicit, reasoning-like behavioral task setting their study explored, we focus on automatic processes

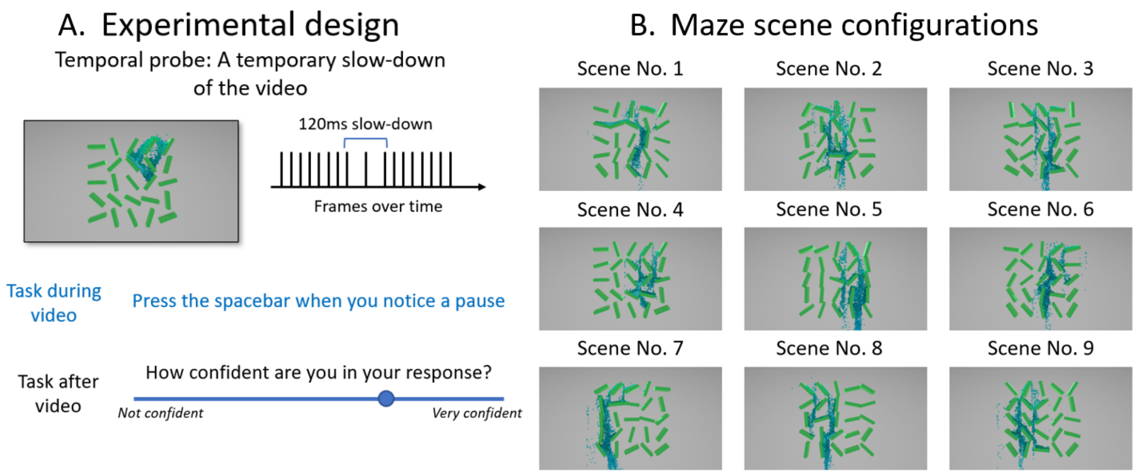

Figure 1: Behavioral experiment and stimuli. (A) The temporal probe detection experiment. (B) Screenshots of the approximately middle frames of each of the nine base maze animations we studied.

that unfold during the perception of complex liquid flow, afforded by the structure of our psychophysical task. Indeed, we suggest that estimating the coarse path the liquid will take might be an implicit goal in the mind, which drives attentional dynamics during spontaneous visual processing. The computational mechanisms identified in this previous work might also partially underlie the behavior we measure here; in such a case, by showing that these simulations can be automatically deployed during spontaneous processing, our work confers a heightened degree of realism to simulation-based representations for intuitive physics.

Experiment: Temporal probe detection

We aimed to measure moment-to-moment attentional distribution during spontaneous visual processing of liquid flow. Following recent work on event segmentation (Yasuda et al., 2021), we showed participants brief animations of liquids flowing through a maze-like obstacle course, and asked them to indicate when they notice a temporary slowing of the animation.

Participants

We recruited 180 participants using Prolific (`https://www.prolific.co`), a crowdsourcing platform. The experiment was conducted using Psiturk (Gureckis et al., 2016). Participants agreed electronically to the consent form before the experiment began. The experiment lasted less than 15 minutes on average and participants received a compensation of $2 upon completion. This study was approved by an Institutional Review Board (IRB).

Stimuli

The stimuli are based on brief animations of a water-like liquid flowing through a maze of 25 randomly rotated planks, configured 5 in a row with 5 rows in total (e.g., Fig. 1A). The animations are initialized with a blue-colored liquid ball at the top of the maze, positioned either just above the second, third, or fourth columns of planks. The liquid flow was simulated using SPlisHSPlasH (Bender, Kugelstadt, Weiler, & Koschier, 2020), a Smoothed Particle Hydrodynamics (SPH) solver. In SPH, behavior of fluid is approximated with a collection of particles, each representing a small package of fluid. We chose the simulation parameters to mimic water behavior as closely as possible (Gravity=[0,-0.98,0], viscosity = 0.01, surface tension = 0.05). After the simulation, the scenes were rendered in Blender (ver. 2.83) (`www.blender.org`). The length of each animation is 125 frames, formed into a video at 25 FPS, making a total duration of 5 seconds.

We made nine "base" animations, each with a different maze configuration. The initial liquid configuration was equally distributed across the three possible locations mentioned above. Maze configurations were hand-curated to ensure the liquid did not branch too much to become distracting.

For each video, we introduced "temporal probes" – temporary slowdowns of the video at chosen time-points. At most one slowdown could occur in a video presented to participants. To generate a slowdown, we paused the video through one of the 18 evenly spaced 3-frame windows, by copying the middle of the these 3 frames across the window. This resulted in a 120-millisecond slowdown (Fig. 1A). The 18 windows consisted of every other 3-frame sequence and excluded the beginning and final 10 frames. This resulted in $9 \times 18 = 162$ stimuli videos, which we divided into 6 conditions. In each condition, there were 27 videos with temporal probes, in addition to 27 unmodified original animations without a probe (obtained by repeating each of the 9 original animations three times). In each condition and for each participant, videos were shuffled and presented in a random order.

Procedure

Before the experiment began, participants were asked to scale a bounding rectangle to the size of a ID card, to ensure that the size of the following video stimuli were approximately the same across participants. They then were shown the experiment instructions, during which participants were informed

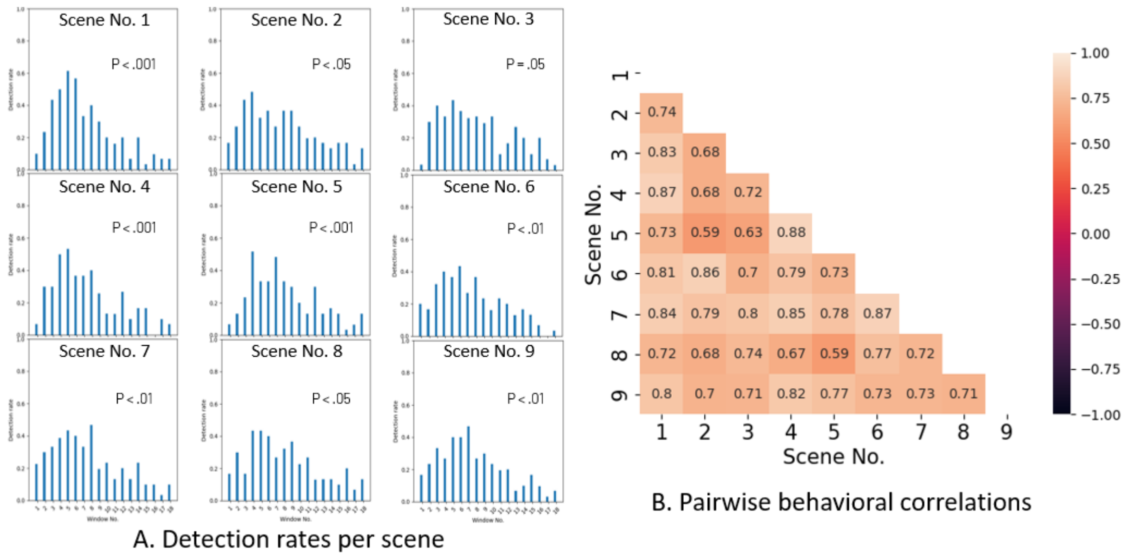

Figure 2: Behavioral results. (A) Detection rates for each scene as a function of time window. (B) The pairwise similarity structure of the detection rates across the nine scenes we studied.

that not all videos have a pause and were shown an example maze. Participants then completed two practice trials, and their understanding of the instructions were checked with a comprehension quiz. If either of the two questions on the quiz was answered incorrectly, the instructions and practice trials were repeated, and participants only proceeded to the main experimental trials once they obtained 100% on the comprehension test. During the experiment, participants were asked to view the stimuli video and press space bar whenever they noticed a pause in the video. Videos were shown centrally on the screen with a three seconds countdown before the start. The border around the video flashed red when it registered a spacebar response; otherwise, the border was black. After each video finished, participants rated their confidence on a sliding scale from 'Not very confident' to 'Very confident'. Progress through the experiment was shown with a trial counter on the bottom of the screen. After completing the experimental trials, participants were asked a series of debriefing questions about task difficulty, engagement, and any feedback they had about the task.

Behavioral results

Time course of detection accuracy is non-uniform over time and systematic across scenes For our main behavioral analysis, we asked whether detection accuracy across 18 evenly spaced temporal probes are significantly non-uniform within each scene. Detection accuracy is defined as the ratio of the number of participants who detected the probe to the number of responses on that trial. To count the number of detection for each temporal probe, we calculated the proportion of participants who indicated the slowdown within the 1000ms interval after the corresponding slowdown occurred. To determine whether detection rates are uniformly distributed across time windows, we performed chi-squared tests for goodness of fit, separately for each scene. We found that all nine scenes showed significantly variable detection rates (Fig. 2A), reflecting a non-uniform distribution of attentional deployment.

We also found that the detection performance has a stereotypical shape across scenes, with the detection rates rising early on in the video and decaying with a long-tail later on in the video. The average correlation between the detection rates of pairs of scenes is $r = .76$ (range = 0.59 - 0.87; Fig. 2B), indicating that despite the stereotypical shape of detections, there is also appreciable variability from scene to scene.

In our models, we will aim to explain the fine-grained patterns of detection rates in Fig. 2A as well as the similarity structure of slowdown detectability across the nine animations in Fig. 2B.

Confidence ratings Finally, we also analyzed participants' confidence ratings to test whether confidence varies as a function of scene identity or time window. For the latter, we coded it as a binary variable with each time window belonging either to the first or second-half of the video. A two-way ANOVA returned no main of effect of scene identity, a marginally significant effect of time window, and no interaction. This result suggests that participants' confidence ratings are minimally impacted by their performance (decreasing slightly from the first to second half).

Models

Here, we consider four models to understand the observed dynamics of detection rates. Two of these predictors – coarse simulation and hybrid simulation – allow us to quantify what lies ahead of the observer: the *sensitivity* of predicted flow-path to increments in time, based on intuitive physics. The other two predictors – pixels and VGG16 – allow us to account for what is right in front of the observer: the bottom-up salience of the slow-down of the video.

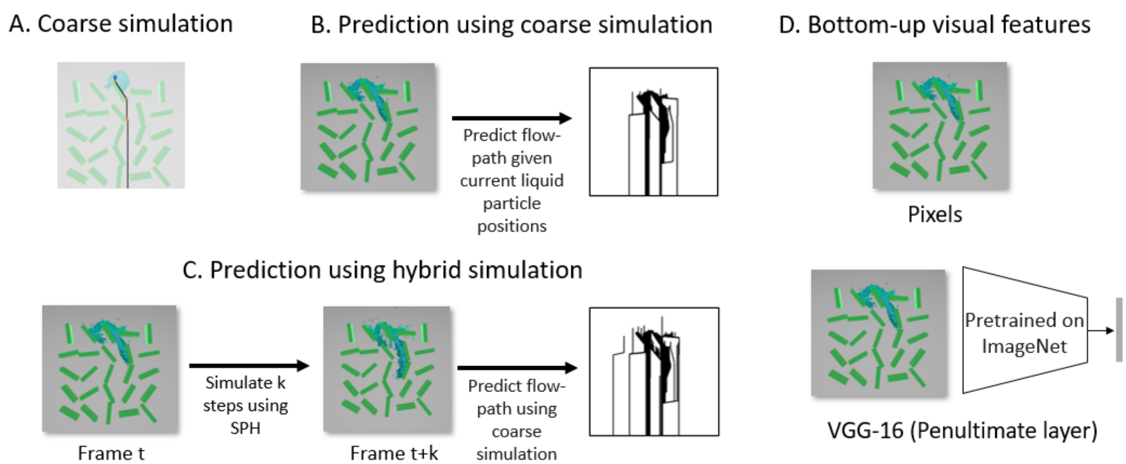

Figure 3: Models. (A) An example of a single liquid particle's predicted path according to the kinematics-only coarse simulation. (B) Predicted path given the ground-truth liquid particle positions from the SPH simulator. (C) Hybrid model. (D) Bottom-up visual features including pixels and the penultimate layer of a pretrained VGG16 network.

Coarse simulation

We hypothesize that the mind spontaneously predicts the path the liquid will take, and the sensitivity of this prediction to increments in time drives attentional deployment. Such adaptive deployment of attention should inform the actual perceptual processing, leading to more precise representations of the liquid if the flow-path predictions are too variable over time or resulting in noisier representations if the flow-path predictions seem to not change very much. We note that this paper does not implement this perceptual processing, but instead considers flow-path prediction as an "implicit task" that might inform the adaptive deployment of attention.

We implement these flow-path predictions using a coarse, kinematics-only simulation model (Fig. 3A, B). At each time point, this model takes as input particle positions from the SPH simulator underlying our stimuli (SPlisHSPlasH), and completes the path each of these liquid particles will take. This prediction is not based on a dynamics simulation. Rather, it follows the heuristic that liquid particle will move in the direction of gravity, until it hits a rigid surface, in which case it will move along the rigid surface. This is an adaptation of the "gravity heuristic" model from Bates et al. (2019), which was inspired by Gardin and Meltzer (1989) and Forbus (1984).[1] Specifically, for each particle position at the current time point, we solve a set of linear equations to get the paths each particle will take through the scene. We aggregate these paths across particles by overlaying them, resulting in the predicted flow-path from the current time point (Fig. 3B).

To relate this model to behavior, we predict the detection rate for a given 3-frame window as the following. In the given 3-frame window, we calculate the average change in the predicted flow-path between each consecutive frame (i.e., from the first frame to the second frame, and from the second frame to the third frame) using L^2-distance. We calculate this average change for each of the 18 such windows throughout the video. The resulting vector of average changes in the predicted flow-path is then compared to behavioral detection rates.

Hybrid simulation

We also explore a more finer-grained hybrid simulation model. In this model, to obtain a prediction of flow-path at a time point, we first play forward the liquid simultion k time steps using SPH, and then execute the coarse simulation model based on the positions of the particles at that time point (Fig. 3C). In contrast to coarse simulation that only considers gravity and maze geometry, this model incorporates dynamics for the initial k time steps of prediction (including momentum and particle-particle interactions). (In our simulations, we set $k = 5$.) Due to the stochasticity in the SPlisHSPlasH engine, we simulate each scene for 5 times and calculate the average position per particle in each frame for each scene, before proceeding to the coarse portion of the simulation.

To relate this model to behavior, we follow the identical procedure as the coarse model. That is, we use the average L^2-distance between the flow-path predictions of the consecutive frames of each 3-frame window, which is done for each of the 18 windows.

Pixels

In addition to top-down, goal-driven processes, visual attention is also attracted to salient stimuli (Pashler, 2000). It is expected that considerable changes in lower-level visual features (a sudden color or shape change) will influence attention deployment. Thus, we wish to establish that above and beyond the influence of lower-level visual features, forward simulation can explain additional variance in human attention deployment. We therefore considered pixel changes as an additional covariate (Fig. 3D). To do this, for a given 3-frame window, we calculated average of L^2-distance between the pixel images of the consecutive frames (like the coarse

[1]We note that in this paper, we have no commitment in terms of the exact nature of the coarse simulation (e.g., whether it is dynamics-based or kinematics-only); instead, we are more generally interested in exploring approximate forms of simulation.

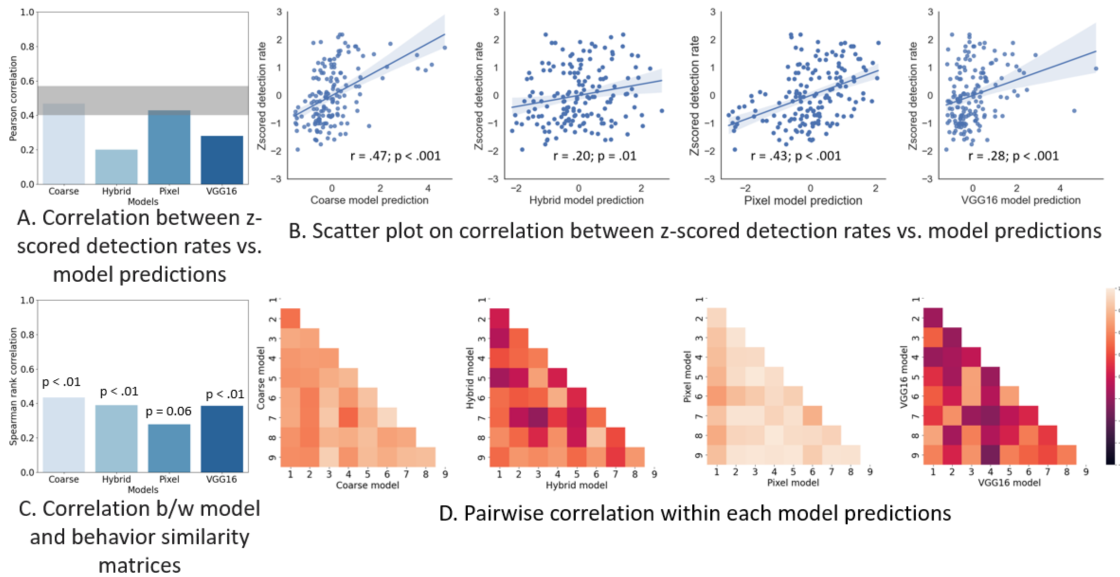

Figure 4: Comparison of models to behavior. (A, B) A bar chart and scatter plots showing the fine-grained relationship between models and behavioral detection rates. (C) Spearman's rank correlation between the similarity structure of behavioral vs. model-predicted detection rates. (D) Similarity matrices of each model (compare to Fig. 2).

simulation model), and we did so for each of the 18 windows. (Notice that this is the identical procedure we use to relate prediction models to behavior.) The resulting vector of average pixel-level differences quantifies our low-level visual features predictor of attention.

VGG16

Following previous work (R. Zhang, Isola, Efros, Shechtman, & Wang, 2018), in addition to low-level pixels, we also tested more abstract or complex visual features using the penultimate layer of the VGG16 model (Simonyan & Zisserman, 2014),[2] a deep convolutional neural network pre-trained on the ImageNet dataset (Deng et al., 2009) (Fig. 3D). To relate this model to behavior, we again followed the same procedure as other models: For a given 3-frame window, we calculated average of L^2-distance between the VGG16 activations of the consecutive frames of this window, and we did so for each of the 18 windows.

Model vs. behavior comparisons

Explaining trial-level detection rates We first compared behavioral detection rates with the four models introduced – coarse simulation, hybrid simulation, pixels and VGG16 – at the level of individual temporal probe trials (Fig. 4A, B). Using regression, we test whether each model's detection rate predictions correlate with participants' detection rates ($n = 162$ trials). Strikingly, we found that some models we explored (coarse simulation and pixels) reached human split-half correlation confidence interval (CI [0.40, 0.57], shaded region in Fig. 4A, left panel). The coarse simulation model (but not pixels) explained more variance than each of the remaining two model alternatives ($p < .05$ two-tailed pairwise comparisons using Fisher r-to-z transformation). These results suggest that beyond bottom-up salience, spontaneous predictions about a system as complex as liquid flow drive dynamic allocation of attention.

[2] We also considered the earlier fully connected layer, which led to a slightly worse fit.

Explaining maze-level similarity of detection rates As we illustrated in our behavioral results in Fig. 2B, there is some measure of similarity between pairs of mazes in the way detection rates change over time. Here, we ask whether this similarity structure can be accounted for by the models, using representational similarity analysis (Kriegeskorte, Mur, & Bandettini, 2008). Paralleling our pairwise behavioral similarity matrix from Fig. 2B, a similarity matrix for each model is obtained by calculating the correlation between pairs of scenes using model-predicted detection rates. These model-based similarity matrices are presented in Fig. 4D. We found the coarse simulation model significantly correlated with the behavioral similarity matrix ($\rho = 0.43, p < .01$; Fig. 4C, using Spearman's rank correlation). Other models also correlated significantly, but to lesser degree numerically (hybrid: $\rho = 0.39, p < .01$; VGG16: $\rho = 0.39, p < .01$), except the pixels-based similarity matrix, which was only marginally correlated with data ($\rho = 0.28, p = .06$; Fig. 4C). These results suggest that beyond low-level visual features, flow-path predictions explain the similarity structure of how liquid flow is experienced across different scene configuration.

Coarse model explains unique variance in detection rates and similarity structure Lastly we ask whether the coarse simulation model explains additional variance when controlling for other models. First, focusing on Fig. 4A, we calculate 1) the partial correlation between coarse model and behavioral detection rates, after removing the effect of other

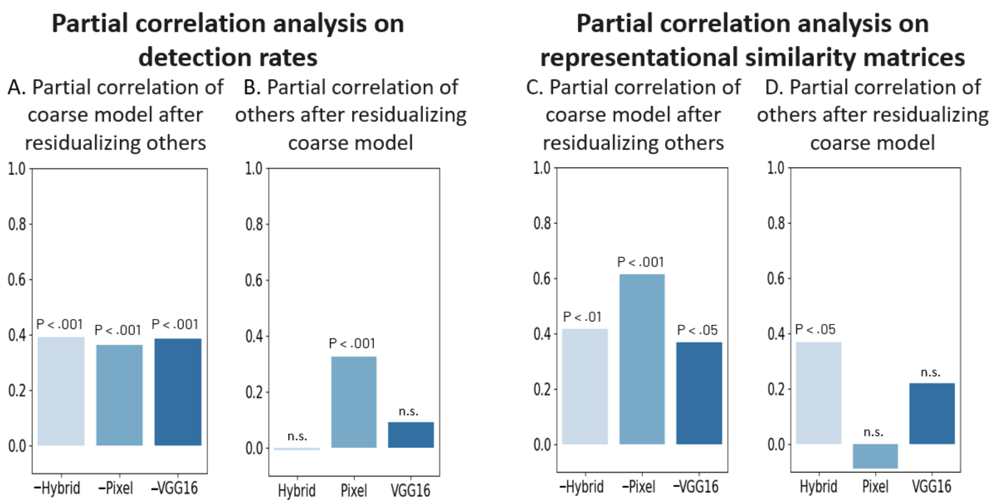

Figure 5: Partial correlation analysis. (A, B) Partial correlation analyses using detection rates. (C, D) Partial correlation analyses using the similarity structure of detection rates across scene pairs.

models, and 2) the partial correlation between other models and behavioral detection rates, after removing the coarse model's effect. We found that the coarse simulation model explains similar amounts of variance after residualizing each of the other models (hybrid: $r = 0.39, p < 0.001$; pixels: $r = 0.36, p < 0.001$; VGG16: $r = 0.39, p < 0.001$; Figure 5A). In the other direction, i.e., after residualizing the effect of the coarse simulation model, only pixels continued to significantly explain variance in detection rates (hybrid: $r = -0.009, p = 0.91$; pixels: $r = 0.33, p < 0.001$; VGG16: $r = 0.09, p = 0.24$; Fig. 5B). These results establish that coarse model flow-path predictions subsume all of the variance explained by the hybrid model in detection rates, and explain additional variance in attentional patterns above and beyond bottom-up salience.

Second, focusing on Fig. 4C, we calculate 1) the partial correlation between the similarity of scenes observed in behavior vs. predicted by the coarse model, after removing the effect of other models, and 2) the partial correlation between the similarity of scenes observed in behavior vs. predicted by other models, after removing the effect of the coarse model. For this partial correlation analysis, we use linear correlation (Pearson's r). As Fig. 5C shows, the coarse model continues to explain unique additional variance relative to all of the alternatives. However, in the other direction, only the hybrid model survives this partial regression analysis (Fig. 5D). Together, these results establish the coarse flow-path predictions as the only model that consistently accounts for behavior across detection rates, scene similarity, and partialing anlaysis.

Discussion

In this study, using computational probes and a performance-based task, which does not overtly relate to prediction or liquid flow, we presented evidence in support of the hypothesis that the mind automatically computes flow-paths of liquid bodies. During liquid-flow perception, we speculate that attention is dynamically deployed to efficiently resolve liquid trajectories using coarse mental simulation. In comparisons to plausible alternatives, we showed that flow-path predictions arising from coarse simulations, which are based on simple geometric heuristics, explains both detailed trial-level detection rates, as well as the similarity structure of these detection rates across the nine maze animations we considered. Coarse model predicted variance beyond both bottom-up salience (operationalized using pixels and VGG16), as well as a finer hybrid model. These results suggest flow-path prediction as an implicit task in the mind, and introduces attentional dynamics as a revealing window onto intuitive physics.

There are multiple limitations to our work. As we noted earlier, our work does not provide a computational mechanism to actually produce these attentional dynamics or to use them for guiding perceptual processing. Existing work on perception of liquids focus on the perception of liquid viscosity (Van Assen et al., 2018; van Assen, Nishida, & Fleming, 2020; Y. Zhang, Bi, & Yildirim, 2022), instead of tracking the liquid as it flows. We leave the development of an integrated attention/perception model of liquid flow perception as future work, however we point to a recent study that considered similar goal-driven perceptual processing problems in the domain of place perception and navigation-related goals (Belledonne & Yildirim, 2021). Our work explored four computational probes and their relation with behavioral signals we measures. However, we can build more on the work of Bates et al. (2019) and Sanchez-Gonzalez et al. (2020) to explore a wider range of simulation substrates, in addition to the coarse simulation presented here. Even though our results suggest that we explain data essentially at the noise-ceiling level using this kinematics-only method, future work should explore a wider range of approximate simulations. Finally, future psychophysical work should consider extending this work to measuring the spatial distribution of attention in relation to the spatial pattern of coarse predictions.

References

Bates, C. J., Yildirim, I., Tenenbaum, J. B., & Battaglia, P. (2019). Modeling human intuitions about liquid flow with particle-based simulation. *PLoS computational biology*, *15*(7), e1007210.

Battaglia, P. W., Hamrick, J. B., & Tenenbaum, J. B. (2013). Simulation as an engine of physical scene understanding. *Proceedings of the National Academy of Sciences*, *110*(45), 18327–18332.

Belledonne, M., & Yildirim, I. (2021). Automatic computation of navigational affordances explains selective processing of geometry in scene perception: behavioral and computational evidence. In *Proceedings of the annual meeting of the cognitive science society* (Vol. 43).

Bender, J., Kugelstadt, T., Weiler, M., & Koschier, D. (2020). Implicit frictional boundary handling for sph. *IEEE Transactions on Visualization and Computer Graphics*, *26*(10), 2982–2993.

Clark, A. (2013). Whatever next? predictive brains, situated agents, and the future of cognitive science. *Behavioral and brain sciences*, *36*(3), 181–204.

Clark, A. (2015). Radical predictive processing. *The Southern Journal of Philosophy*, *53*, 3–27.

Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K., & Fei-Fei, L. (2009). Imagenet: A large-scale hierarchical image database. In *Computer vision and pattern recognition, ieee conference on* (pp. 248–255).

Forbus, K. D. (1984). Qualitative process theory. *Artificial intelligence*, *24*(1-3), 85–168.

Gardin, F., & Meltzer, B. (1989). Analogical representations of naive physics. *Artificial Intelligence*, *38*(2), 139–159.

Gerstenberg, T., Peterson, M. F., Goodman, N. D., Lagnado, D. A., & Tenenbaum, J. B. (2017). Eye-tracking causality. *Psychological science*, *28*(12), 1731–1744.

Gureckis, T. M., Martin, J., McDonnell, J., Rich, A. S., Markant, D., Coenen, A., … Chan, P. (2016). psiturk: An open-source framework for conducting replicable behavioral experiments online. *Behavior research methods*, *48*, 829–842.

Holkar, K., & Waghmare, L. M. (2010). An overview of model predictive control. *International Journal of control and automation*, *3*(4), 47–63.

Kriegeskorte, N., Mur, M., & Bandettini, P. A. (2008). Representational similarity analysis-connecting the branches of systems neuroscience. *Frontiers in systems neuroscience*, 4.

Pashler, H. (2000). The psychology of attention. *JOP*, *14*, 1.

Sanchez-Gonzalez, A., Godwin, J., Pfaff, T., Ying, R., Leskovec, J., & Battaglia, P. (2020). Learning to simulate complex physics with graph networks. In *International conference on machine learning* (pp. 8459–8468).

Simonyan, K., & Zisserman, A. (2014). Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*.

Tassa, Y., Erez, T., & Todorov, E. (2012). Synthesis and stabilization of complex behaviors through online trajectory optimization. In *2012 ieee/rsj international conference on intelligent robots and systems* (pp. 4906–4913).

Van Assen, J. J. R., Barla, P., & Fleming, R. W. (2018). Visual features in the perception of liquids. *Current biology*, *28*(3), 452–458.

van Assen, J. J. R., Nishida, S., & Fleming, R. W. (2020). Visual perception of liquids: Insights from deep neural networks. *PLoS computational biology*, *16*(8), e1008018.

Yasuda, S., Yates, T. S., & Yildirim, I. (2021). Physical event representations: Observers spontaneously impose discrete temporal structure in intuitive physical scene understanding. *Journal of Vision*, *21*(9), 2672–2672.

Zhang, R., Isola, P., Efros, A. A., Shechtman, E., & Wang, O. (2018, June). The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the ieee conference on computer vision and pattern recognition (cvpr).*

Zhang, Y., Bi, W., & Yildirim, I. (2022). Perception of liquids relies on generalizable, physics-based representations. In *Proceedings of the annual meeting of the cognitive science society* (Vol. 44).