Brain and Cognition 112 (2017) 78–83

Contents lists available at ScienceDirect

Brain and Cognition

journal homepage: www.elsevier.com/locate/b&c

# Disentangling signal and noise in autism spectrum disorder

Sander Van de Cruys [a,b,*], Ruth Van der Hallen [a,b,c], Johan Wagemans [a,b]

[a] Laboratory of Experimental Psychology, Brain & Cognition, KU Leuven, 3000 Leuven, Belgium
[b] Leuven Autism Research (LAuRes), KU Leuven, 3000 Leuven, Belgium
[c] Child and Adolescent Psychiatry, UPC KU Leuven, 3000 Leuven, Belgium

ARTICLE INFO

*Article history:*
Received 8 September 2015
Revised 11 August 2016
Accepted 11 August 2016
Available online 17 September 2016

*Keywords:*
Autism spectrum disorder
Motion perception
Coherent motion
Perceptual integration
Bayesian perception
Predictive coding

ABSTRACT

Predictive coding has recently been welcomed as a fruitful framework to understand autism spectrum disorder. Starting from an account centered on deficient differential weighting of prediction errors (based in so-called precision estimation), we illustrate that individuals with autism have particular difficulties with separating signal from noise, across different tasks. Specifically, we discuss how deficient precision-setting is detrimental for learning in unstable environments, for context-dependent assignment of salience to inputs, and for robustness in perception, as illustrated in coherent motion paradigms.


## 1. Introduction

As part of the wider trend in computational psychiatry (Friston, Stephan, Montague, & Dolan, 2014; Montague, Dolan, Friston, & Dayan, 2012; Stephan & Mathys, 2014), recent explanatory accounts of autism spectrum disorder (ASD) take inspiration from well-articulated information processing models of typical cognition (Hohwy, 2013; Lawson, Rees, & Friston, 2014; Pellicano & Burr, 2012; Qian & Lipkin, 2011; Quattrocki & Friston, 2014; Rosenberg, Patterson, & Angelaki, 2015; Sinha et al., 2014; Van de Cruys et al., 2014). Particularly influential in most of these new proposals for ASD is the predictive coding framework (Clark, 2013; Friston, 2010). Predictive coding assumes that the brain builds a so-called generative model about the environmental causes of the sensory inputs it receives. It infers these causes by making a best guess, or prediction, about incoming inputs at each point in time and evaluating whether the predicted sensory activity corresponds with that actually received through the senses. If not, the system will attempt to reduce this mismatch, or prediction error, by adjusting its prediction about the state of the environment and adapting its generative model for the current context accordingly. Within this scheme, these models are hierarchically structured (Rohe & Noppeney, 2015; Wacongne et al., 2011), where higher levels are capable of capturing patterns in sensory inputs that have larger spatial or temporal spans.

However, not all prediction errors are created equal. In order to appropriately weigh a prediction error, not only the mean (best estimate) is predicted at each level, also the precision (inverse variance) of the prediction error is estimated. The comparison with a statistical *t*-test makes clear why this is important: in a *t*-test a difference in means ("prediction error") is weighted by the variance or expected (standard) error (Friston, 2010). Otherwise, there is no way to interpret the importance (informative value) of the differences one finds. Technically, precisions are hyper-parameters which are estimated and learned with the same predictive coding machinery. Multiple types of uncertainties in the inputs we receive, make the task of predicting the world particularly challenging. There may be uncertainty because of our lack of knowledge about a particular regularity in the environment, either because we have not fully learned the regularity, or because the regularity has recently changed, which happens frequently in a volatile environment. Uncertainty can also be due to the probabilistic nature of a given regularity: By chance, an expected input pattern may not occur. All these types of uncertainties will result in prediction errors in the system. Unfortunately, we do not know a priori whether a given prediction error is actually relevant, i.e., corresponds to an actual learnable (change in) regularity in the environment, or not relevant, i.e., is due to probabilistic noise variability. In the first case the prediction error should be used to change inferences and learn new structure, but in the second case

* Corresponding author at: Brain & Cognition, KU Leuven, Tiensestraat 102, Box 3711, 3000 Leuven, Belgium.
*E-mail address:* sander.vandecruys@kuleuven.be (S. Van de Cruys).

http://dx.doi.org/10.1016/j.bandc.2016.08.004
0278-2626/

we should largely ignore it, using it only to improve future precision estimates. Optimally, precision or gain should be high (Yu & Dayan, 2005) when prediction errors correspond to reducible, learnable uncertainty or when confidence in the prediction is low.

In many of the recent proposals for ASD, deficient precision estimation is assumed to be key (Lawson et al., 2014; Palmer, Paton, Kirkovski, Enticott, & Hohwy, 2015; Pellicano & Burr, 2012; Van de Cruys, de-Wit, Evers, Boets, & Wagemans, 2013; Van de Cruys et al., 2014). Our own account, termed "HIPPEA" (for High, Inflexible Precision of Prediction Errors in Autism), assumes that bottom-up prediction errors are assigned a precision that is too high and not adapted (inflexible) to the uncertainty in the context (Van de Cruys et al., 2014). A crucial consequence of this is that prediction errors are taken at face value, hence there will be too little discounting of (prediction errors stemming from) noise or irrelevant variability in inputs. Non-repeating, accidental variations in the input will receive disproportionately high weight, resulting in overfitting to these irrelevant differences: models will be shaped by putative regularities that will not generalize. By setting precision invariably high, "training examples" will be more literally encoded (cf. veridical mapping, Mottron et al., 2013). According to HIPPEA, inefficient predictive updating will result in different top-down priors or predictions, namely overfitted, low-level ones that capture too much redundant inputs, and possibly an impoverished set of very high level ones that are not sufficiently informative. Such an incomplete hierarchical generative model would result from the fact that unduly high precision of inputs induces predictive matching that takes place on (and might not get beyond) very local (lower) levels. While this happens at the expense of detecting more abstract regularities, note that the basic capacity of forming predictions remains unaffected. Rather, encoding of noise hampers discovery of regularities when these are embedded in more complex, noisy inputs.

It follows that we should be careful in stating that individuals with ASD cannot form informative priors (only "weak" or low precision priors), as several authors seem to suggest (Manning, Tibber, Charman, Dakin, & Pellicano, 2015; Pellicano & Burr, 2012; Sinha et al., 2014; Zaidel, Goin-Kochel, & Angelaki, 2015). Surely, in lots of cases individuals with ASD can detect and learn to use regularities (in the form of informative priors). A recent study by Spanò, Peterson, Nadel, Rhoads, and Edgin (2015) convincingly showed that children with ASD use both low-level priors (on convexity and surface integration) and higher level priors (based on form/object memories) in a basic visual figure-ground segregation task, to the same extent as typically developing children. As pointed out before, even probabilistic and implicit regularities can be learned in ASD (e.g., Nemeth et al., 2010; Roser, Aslin, McKenzie, Zahra, & Fiser, 2015; Solomon, Smith, Frank, Ly, & Carter, 2011). If the task makes clear which stimuli or dimensions are relevant, people with ASD may even be more sensitive to changes in environmental patterns (Westerfield, Zinni, Vo, & Townsend, 2015) as seen in the P3 ERP component, consistent with our proposal of increased precision of prediction errors.

Still, there is important commonality between the HIPPEA account and the weak priors account (Pellicano & Burr, 2012), because both are talking about *relative* precisions of bottom-up and top-down information *on each level of the hierarchy* (see also Lawson et al., 2014). The weight of new evidence (and thus the change in prediction) is defined as the precision of input divided by the precision of the prior (see Mathys et al., 2014 for the full computational details). Hence, weaker (i.e., lower precision or less confidence in) top-down predictions would also lead to increased reliance on bottom-up information as in the "weak priors" account (Pellicano & Burr, 2012). The importance of relative precisions, however, also implies that studies that find reduced adaptation in behavior or reduced repetition suppression in fMRI responses in ASD (or high autism traits) (e.g., Ewbank et al., 2014; Molesworth, Chevallier, Happé, & Hampton, 2015; Turi et al., 2015) cannot simply be considered evidence for the weaker priors thesis, even though both adaptation and repetition suppression are thought to be the result of (top-down) predictive activity (Chopin & Mamassian, 2012; Summerfield, Trittschuh, Monti, Mesulam, & Egner, 2008).

Both ways of shifting the balance in inference (higher bottom-up precision vs. lower top-down precision) should be dissociable, at least in principle, when considering the result of inference. Specifically, one would expect higher precision of the posterior estimate for ASD in case of higher precision prediction errors. At first sight, this seems to be a testable prediction, for example by directly or indirectly probing for decision confidence (Meyniel, Schlunegger, & Dehaene, 2015; Yeung & Summerfield, 2012). Additionally, one should expect less trial-by-trial fluctuations in confidence according to the current proposal. However, simple confidence measures may not be able to satisfactorily answer these questions, given that (1) they provide one measure for something (posterior) that takes place on multiple hierarchical levels, (2) they might be affected by (executive) post-perceptual processes, and (3) they require the capacity to explicitly reflect on one's own thought processes (explicit metacognition), which may be particularly deficient in ASD (Grainger, Williams, & Lind, 2014).

To make progress on these Bayesian accounts of ASD, it will be important to study the updating (learning) and application of priors on a trial-by-trial basis, precisely quantifying varying uncertainties and to what extent they are taken into account for future inference. In the next section, we will discuss studies in ASD that are beginning this enterprise, in the context of learning in unstable or volatile environments. We will see that forming higher order expectations on volatility are necessary to restrain the effect of noise.

In a second section, we will look at studies on salience in perception in ASD. Even if priors are learned, they may fail to adequately smooth out the variability that is inherent to all natural stimuli, because of the weight that deviations receive. This increased sensitivity to variability, irrespective of its origin or relevance, means that the informative value or salience of different pieces of input is not properly determined. Again, this can be seen as an inability to disentangle relevant (signal) and irrelevant (noise) inputs, dependent on a given context.

In a final section, we will discuss the sensitivity to variability in the input and how that leads to the lack of robustness in inference. We will discuss studies that suggest that coherent motion perception and motor behavior is more vulnerable to noise in ASD. Here too, second-order estimations of to be expected variability, learned across different experiences, would typically help rein in noise, but because of the precision-based mechanism described above this seems to be hampered in ASD.

## 2. Learning in unstable environments

A mix of uncertainties (actual and accidental changes) is particularly present in a probabilistic learning task, where the governing rule (e.g., stimulus A is rewarded) has to be learned based on imperfect (probabilistic) feedback on your choices, and the governing rule can change unexpectedly (e.g., not A but B is rewarded from now on). If one accurately estimates the intrinsic level of uncertainty across multiple trials (i.e., the expected amount of prediction errors one will encounter), it is easy to weight feedback that exceeds this expected uncertainty, such that subsequent predictions about which rule holds, will be updated. Probabilistic reversal learning studies in ASD participants show that while they seem perfectly able to learn a probabilistic rule initially,