The Energy Metabolic Footprint of Predictive Processing in the Human Brain

André Hechler

Graduate School of
Systemic Neurosciences
LMU Munich

Dissertation der
Graduate School of Systemic Neurosciences der
Ludwig-Maximilians-Universität München

December 2023

Supervisor

Prof. Dr. med. Valentin Riedl, PhD

Department of Neuroradiology

Klinkum rechts der Isar der Technischen Universität München

First Reviewer: Prof. Dr. med. Valentin Riedl, PhD

Second Reviewer: Prof. Dr. Markus Ploner

External Reviewer: Prof. Dr. med. Philipp Sterzer

Date of Submission: December 18th, 2023

Date of Defense: May 14th, 2024

ACKNOWLEDGEMENTS

This work would have never seen the light of day without the support of the people around me. I count myself incredibly lucky to have found this wonderful group of colleagues who soon became my friends. The way your dedication to research and excellence never stood in the way of your empathy and understanding in stressful times was invaluable. I don't think there was a single day where I did not enjoy going to the lab, no matter which tasks awaited. Most importantly, no deadline could keep us from being excited about having lunch together.

I want to thank my supervisor, Valentin Riedl, for his unending support and sense of vision. Your leadership provided a great balance of guidance, empowerment and inspiration that kept me motivated throughout these years. You created a working environment that was as fun as it was exciting, as relaxed as it was stimulating. Our lab retreats to the Gruttenhütte will be the stuff of legend.

Thank you, Antonia, for the countless laughs we shared and your ability to switch from the most absurd topics to in-depth technical discussions with ease. Also for the chocolate. Thank you, Samira, for sharing your sense of commitment and compassion for topics that go way beyond our everyday life. Thank you, Gabriel, for dealing with my endless barrage of technical questions and the support on my way to eventually answering some of them on my own. My thanks further go to Mahnaz for great discussions about statistics and beyond, to Laura for sharing a sense of critical thinking about all things research and of course to everyone else that made my time in the lab unforgettable. This includes, but is not limited to, Dani, Alyssa, Roman, and Katarzyna.

There are many friends outside of the lab that deserve my thanks, and I here especially want to mention Randolf and Yvonne. I could not imagine a better friendship.

I want to thank my parents, without whom I would not be where and who I am today. Your support on my winding road to these pages made all this possible.

Finally, thank you, Sarah. Words cannot express how much it means to have you by my side. I can't wait for the next part of our journey.

ABSTRACT

To maximize our chances of survival and procreation, we need to process our environment in a highly sophisticated and accurate manner. In their limit, these two demands are mutually exclusive: While better sound localization, quicker reflexes or more accurate vision could improve survivability, the necessary energy consumption might not be sustainable. Luckily, our sensory systems strike an impressive balance between performance and energetic cost. In a both active and passive process, we learn about the rules that determine our experience and use them to form expectations. Efficient brain activity is then achieved by limiting the forward transmission of signals to deviations from what we predicted.

In the visual domain, this means that our perception is dominated by our expectations when we are in a familiar environment. Research in cognitive neuroscience has shown that expected input elicits weaker brain activity than surprising input, without any behavioral disadvantages. However, knowledge about associated energetic efficiency is limited by three gaps in the current literature. First, conventional imaging techniques do not provide direct measurements of energy metabolism. Second, previous research has focused on localizing areas of maximal effect, potentially missing weaker, but more widespread patterns. Third, our knowledge about the world is imperfect, leading to uncertain expectations. This has rarely been accounted for.

Neuronal activity is fueled by ATP, most of which is produced with chemical reactions that need oxygen. In the present work, I assessed energy metabolism with a novel imaging method that measures the rate of oxygen consumption across all parts of the brain. I used an experimental design during which participants saw visual object sequences that were either predictable, random, or surprising. Behavioral tests indicated that predictable sequences were learned without any feedback which resulted in anticipation of upcoming objects. I further found that participants varied in the confidence of their expectations. This had a major impact on oxygen consumption when viewing predictable sequences: The lowest energy usage was found for high levels of confidence. This effect was not limited to sensory regions but extended across large parts of the brain. Interestingly, my results suggest that confidence led to energy savings even when the visual input was objectively random. In conclusion, this work provides the first evidence that our expectations are a major promoter of efficient processing, which is crucial for any organism with limited energy availability.

TABLE OF CONTENTS

1 General introduction

The brain is an energy intense organ that consumes approximately 20 percent of the energy available to the body (Rolfe & Brown, 1997). Most of that energy is used to fuel the electrochemical activity of neurons. The action potential is the basic signaling unit in our nervous system and its propagation across the axon consumes 16% of the neuronal energy budget. Even more costly is the integration of multiple incoming signals at the dendrites: 44% of the energy budget is used towards postsynaptic potentials. Lastly, the resting potential itself needs to be upheld to enable action potentials and accounts for 15% (Howarth et al., 2012). The main energy currency is adenosine-triphosphate (ATP), the majority of which is created in the mitochondria during cellular respiration. Here, one mol of glucose is metabolized with six moles of oxygen, yielding approximately 36 mol of ATP. Healthy brain function is heavily dependent on the constant provision of these compounds: Limited availability of glucose quickly impacts cognitive function (Warren & Frier, 2005) and sustained oxygen deprivation is a driver of inflammation and metabolic diseases (Eltzschig & Carmeliet, 2011; Shobatake et al., 2022). Consequently, the management of energetic resources in the brain is of vital importance (Quintela-López et al., 2022).

Given the link between metabolic cost and signaling, efficient use of energetic resources rests on efficient neural activity. Neurons are thought to process (*encode*) sensory information in a way that minimizes the transmission of redundant information. This is known as the *efficient coding hypothesis* (Barlow, 1961; Simoncelli, 2003). Under this assumption, the brain preferentially processes deviations from previous input, or unique aspects of what is currently perceived. Both in space and in time, repeating patterns govern our experience, endowing us with the ability to process a wealth of information without encoding all details (Olshausen & Field, 2004). This logic has been extended from the level of neurons and neuronal populations to the interaction of large regions on a whole-brain level (Bullmore & Sporns, 2012; Zhou et al., 2022). Surprisingly, empirical evidence of corresponding energy efficiency is rare. This might be due to the difficulty of measuring metabolic processes across the human brain.

1.1 Measuring energy consumption in the human brain

While ATP metabolism cannot easily be measured with routine human brain imaging techniques, glucose and oxygen can be assessed using positron-emission-tomography (PET) and magnetic resonance imaging (MRI) respectively. PET provides direct access to glucose

consumption (or rather, its uptake rate into cells) but necessitates the use of a radioactive tracer which limits its applicability in healthy subjects. Standard MRI techniques are much more prevalent due to their non-invasive nature and sensitivity to changes in the oxygen content of blood. The dominant method of measuring neural activity in the brain is functional MRI (fMRI) based on the blood oxygenation level dependent (BOLD) signal (Ogawa et al., 1990). BOLD imaging exploits the differential effect of oxygen-carrying hemoglobin (Hb_{O2}) and deoxygenized hemoglobin (dHb) on the magnetic field. To fuel cellular respiration, arterial blood saturated with Hb_{O2} perfuses the local vasculature, and oxygen is taken up into cells by passive diffusion. Consequently, the local dHb concentration increases, leading to characteristic distortions of the magnetic field. In isolation, this triggers a signal decrease (lower brightness in the resulting image). However, a disproportionate increase in cortical blood flow (CBF) provides Hb_{O2} in excess of the amount of oxygen being consumed (Buxton et al., 2004). This supercompensation means that neuronal activity results in higher brightness in BOLD images. fMRI therefore measures a compound signal of hemodynamics and cellular respiration, complicating its interpretation in terms of energy consumption or neuronal activity (Drew, 2019; Kim & Ogawa, 2012). While its correlation to local electrophysiological activity has been established in primary sensory cortices (Logothetis et al., 2001), the link between neural activity and hemodynamic response (*neurovascular coupling*) varies across anatomical regions (Ances et al., 2008; Buxton et al., 2014; Devonshire et al., 2012).

Measuring energy consumption with MRI therefore necessitates dedicated imaging techniques, both in acquisition and processing. A promising set of methods extended BOLD imaging with the goal of quantifying the cerebral metabolic rate of oxygen consumption (CMR_{O2}). These multiparametric MRI approaches measure or estimate both hemodynamic effects and deoxygenation to disentangle the constituent parameters of the BOLD signal (Bright et al., 2019). Specifically, multiparametric quantitative BOLD (mqBOLD; Christen et al., 2012; Hirsch et al., 2014; Kaczmarz et al., 2020) combines separate MR measurements of CBF, cerebral blood volume (CBV) and deoxygenation to calculate CMR_{O2} at high spatial resolution. The MR data is combined with measures of arterial oxygen concentration (Ca_{O2}) using pulse oximetry and Hb concentration using blood sampling. The underlying logic of calculating CMR_{O2} is summed up by *Fick's Principle* (Fick, 1870):

$$\mathbf{CMR_{O2} = Ca_{O2} * CBF * OEF}$$

The amount of oxygen carried in the arteries (Ca_{O2}) is combined with the volume of blood flowing to a given region (CBF) and finally multiplied by the relative amount of oxygen used (oxygen extraction fraction, OEF). This yields the absolute amount of consumed oxygen in a brain region, in units of µmol O_2/min/100g. I describe the complete process in the corresponding methods section of project II.

Previous studies using mqBOLD focused on comparisons between healthy and clinical samples in experiments without explicit tasks (*resting state*) (Christen et al., 2013; Göttler et al., 2019). The resulting CMR_{O2} measurements are in good agreement with the literature across acquisition methods (Christen et al., 2013; F. Xu et al., 2009a). Basic research is mainly available for the visual domain. In the visual cortex, oxygen consumption increases by 10-30% during visual stimulation with a much larger increase in CBF (Lin et al., 2010). This supports classical accounts of neurovascular coupling. However, blood flow and oxygen consumption can decouple (Fox & Raichle, 1986) and CBF, CMR_{O2} and BOLD show adaptation (i.e. a reduction in response across longer stimulation) to different degrees (Moradi & Buxton, 2013). The relationship between blood flow and oxygen consumption also depends on attention: While unattended visual stimuli elicit coupled increases in line with classical neurovascular coupling, attended stimuli elicit outsized increases in CMR_{O2} (Moradi et al., 2012). Lastly, there is an ongoing discussion about the reason for the observed decoupling. A line of research found an increase in non-oxidative metabolism during task activation (Fox et al., 1988; Paulson et al., 2010). In terms of energy production, this metabolic pathway is much less energy-efficient (yielding only 2 ATP per glucose as opposed to 36 with the full oxidative pathway). This means oxidative energy metabolism is still dominant, even during asymmetric increases of non-oxidative processes (Lin et al., 2010).

In summary, constant availability and efficient use of energy is vital to the function of our brain. Glucose and oxygen are the main resources of energy metabolism, and their presence and use can be measured with neuroimaging methods. However, classical imaging techniques either rest on the application of radioactive compounds (PET) or are ill-equipped to isolate metabolic aspects of their signal (fMRI). In contrast, mqBOLD provides direct access to oxygen consumption rates across the brain. Since previous studies focused on clinical populations or simple sensory stimulation, many questions about the role of energy efficiency in human cognition are still unanswered.

1.2 Perception as inference

For a long time, visual processing (and sensory processing in general) was considered to be a bottom-up process carried out by independent modules. Early visual cortices were thought to identify low-level attributes like lines and edges which are combined into holistic percepts like objects later in the visual stream (Fodor, 1983; Pylyshyn, 1999). Over the last two decades, the mainstream view shifted towards an integrative account, where every level of processing is heavily impacted by later stages. This includes abstract knowledge about our environment and its objects based on our experience (Lupyan, 2015). Prior knowledge is assumed to be a necessary part of perception since our sensory input is often ambiguous. Crucially, there is no clear one-to-one correspondence between the concepts we think in and the current sensory information. Many of our concepts are abstractions that refer to a possibly infinite set of sensory patterns (e.g. a house or a cat). Furthermore, the same external object can lead to drastically different input to our retina depending on viewing distance, angle, lighting, occlusion, and many other factors. The fact that our brain is highly proficient at dealing with this inherent uncertainty led to the theory of perception as inference (Barlow, 1990; Helmholtz, 1867). According to this framework, perception is the continuous process of inferring the most likely source of our sensory input.

Inferring the likelihood of a hypothesis given some evidence is well-defined in Bayesian calculus. In Bayesian terminology, sensory input (also called *evidence*) is integrated with our previous best guess about what we see (the *prior belief*) to arrive at the best interpretation (the *posterior belief*). For any new evidence, this inference cycle continues (see Knill & Pouget, 2004; Pouget et al., 2013 for reviews of Bayesian inference in the brain). This process is inherently probabilistic, meaning all the parameters are captured as probability distributions. Beliefs can be described by a point estimate of the most likely hypothesis and a variance. The variance (often referred to by its inverse, *precision*) captures the uncertainty of our beliefs, while the point estimate relates to the *confidence* that this specific hypothesis is accurate. In an evolutionary setting, one could identify movement in our peripheral vision as most likely being a peaceful animal, but our concrete action would depend on how confident we are in that assessment. It follows that our perception of the world is highly subjective since it rests on individual prior knowledge. Even if a given belief is objectively wrong, high precision means that a lot more evidence is needed to make us change our mind.

A highly influential theory of Bayesian inference in the brain has been formulated in the seminal paper on predictive coding (Rao & Ballard, 1999). The authors showed that

characteristic features of visual cortex activity emerge from a computational model that actively predicts its own input while only forwarding deviations from predictions (termed *prediction errors,* PE) to higher levels. Under this account, the brain learns a so-called *generative model* (Dayan et al., 1995) that encodes how external causes (e.g. the presence of a cat) generate observable sensory input (e.g. the line-like percepts of fur and whiskers). The inferred cause then engenders predictions about further observations (e.g. the shape of the ears or the rhythm of movement). Since veridical encoding of these stimuli does not add to our inference, predicted aspects can be subtracted from bottom-up input. If the sensory input does not conform to the predictions, PEs are used to improve (*update*) the generative model. In this context, the term "prediction" is not synonymous with forecasting (i.e. predicting a future state from current and past ones). Rather, predictive coding aims to infer the most likely source of the current percept. However, prediction in time can be accommodated by the same model (Jiang & Rao, 2022). Extending the logic of predictive coding, the *Free Energy Principle* (FEP; Friston, 2010) posits that reduction of PEs is the guiding principle of all perception and behavior. More precisely, intelligent lifeforms can be described as optimizing their generative model by minimizing a quantity related to PEs: Free energy. In this contest, the generative model is considered a *world model*, since it encompasses all beliefs about our experience (Friston et al., 2021). Intriguingly, the core mathematical formulation of free energy minimization can be expressed as balancing accuracy and complexity: While the model should represent the world as accurately as possible, its complexity should be minimal. This tradeoff suggests that perception and cognition aim for efficiency. Indeed, under the assumption that our central nervous system represents this model, it has been suggested that minimizing free energy achieves accurate inference with minimal metabolic cost (Sengupta et al., 2013).

In summary, vision and perception in general can be understood as inferring what happens in the external world from imperfect sensory information. The uncertainty inherent in linking retinal images to abstract cognitive concepts can be resolved by employing Bayesian inference. Predictive coding suggests that the brain constantly improves the accuracy of our inference by representing and improving a generative model of the world. Instead of only aiming for perfect inference, the FEP suggests that the model is optimized for efficiency.

1.3 Learning temporal patterns in visual input

An increasing body of research substantiates the claim that perception and behavior are strongly driven by prior expectations and contextual information in space and time

(reviewed in de Lange et al., 2018). To study the proposed efficiency of perception, experiments have been developed that present participants with a chance to learn patterns in sensory input. Comparing naïve to familiarized stages across time in the same individual, or across individuals, should reveal optimized models and efficient use of energy. In an established experimental paradigm, participants are presented with sequences of visual stimuli that are determined by conditional probabilities, i.e. which stimulus is shown depends on the one(s) that preceded it (Conway, 2020; Schapiro & Turk-Browne, 2015). To stress the fact that stimuli follow each other in time, the rules governing the sequences are referred to as *transitional* probabilities. The same Bayesian inference as described before applies here: Instead of mapping between sensory input and a percept, the relationship between current and future percepts is learned. After some exposure, given a current percept and a generative model of the transitional probabilities, the most likely future percept can be inferred (Fiser & Lengyel, 2022).

Our ability to learn transitions and their probabilistic structure has been studied using a wide range of related paradigms. *Statistical learning* focuses on the continuous process of perceptual inference that is thought to drive our experience (Fiser & Aslin, 2002). *Sequence* (or sequential) *learning* is closely related, but while statistical learning often includes the segmentation of continuous streams into sequences, the patterns may explicitly be presented during sequence learning (Gheysen & Fias, 2012). Lastly, some authors differentiate *rule learning* from statistical learning based on whether transitions are deterministic or probabilistic (Maheu et al., 2022).

Statistical learning in particular refers to the *implicit* acquisition of transitional patterns, without intention and awareness (Schapiro & Turk-Browne, 2015). After passively viewing the visual input for a few minutes up to an hour, behavioral effects emerge: Participants can differentiate sequences that follow the transition patterns from the learning phase and detect expected stimuli more quickly (Turk-Browne et al., 2005). However, while some studies provided evidence for the lack of awareness during statistical learning (Alamia et al., 2016; Destrebecqz & Cleeremans, 2001; Turk-Browne et al., 2005), the validity of these assumptions has been debated: First, a study using post-experimental questionnaires found that many participants did report awareness of underlying patterns (Dale et al., 2012). Second, the statistical validation of the absence of an effect is problematic with many standard tests and previous studies might have been underpowered to assert it (Vadillo et al., 2016). Opposing

an all-or-nothing position on this issue, it has been shown that both implicit and explicit processes contribute to statistical learning (Batterink et al., 2015). Nevertheless, dissociations have been found that differentiate implicit and explicit designs, especially with respect to neuroanatomical correlates in fMRI studies (Aizenstein, 2004; Batterink et al., 2019). While previous literature reflects considerable heterogeneity in paradigms, sometimes leading to disparate findings, recent years saw the development of unifying frameworks. Conway (2020) argued that most approaches reflect similar phenomena and vary gradually on three axes: Amount of exposure, amount of structure and amount of explicit instruction. The underlying process employed by the brain is thought to be the same across all variants: Bayesian inference (Fiser & Lengyel, 2022).

Studies using fMRI have provided evidence that statistical learning leads to prediction and anticipation on a neuronal level: In the visual cortex, expected events evoke corresponding activity even when they are omitted (Ekman et al., 2017; Kok et al., 2014) and characteristic signal patterns can be detected ahead of presentation (Summerfield et al., 2006). Furthermore, the activity patterns of neuronal populations increase in similarity for mutually predictive stimuli (Schapiro et al., 2012). The analysis of stimulus representations in fMRI data also showed that patterns representing expected stimuli are more stable, enabling the "readout" (*decoding*) of stimulus identities from activation patterns (Kok, Jehee & de Lange, 2012). Lastly, contrasting imaging findings of predictable input with unpredictable (or surprising) input yielded a highly consistent finding: The same stimulus elicits lower activity when it confirms a learned pattern compared to a violation or the absence of a pattern (Alink et al., 2010; Egner et al., 2010; Richter et al., 2018a).

This phenomenon has been referred to as *expectation suppression* and is usually interpreted in the context of predictive coding: Bottom-up signals (PEs) that are already accounted for by top-down predictions are suppressed because they don't improve our inference (Richter et al., 2018a). An alternative view is *surprise enhancement*, which posits that the forwarded PE signal is upregulated for surprising stimuli (without the predictable stimulus being affected). A single contrast of predictable and surprising stimuli cannot arbitrate between these explanations, because a measured difference can be based on either one of them. Consequently, some studies introduced a third type of stimulation, where the temporal contingencies are random (Davis & Hasson, 2018; for studies in animals see Kaposvari et al., 2018; Ramachandran et al., 2017). In these studies, random stimulation is assumed to provide

a baseline (or neutral) condition, which isn't subject to either expectation suppression or surprise enhancement. The former should lead to relatively lower activity for the predictable condition, while the latter is expected to lead to relatively higher activity for the surprising condition. A recent review found that evidence for expectation suppression is predominantly found in animal studies, but generally lacking in humans (Feuerriegel, Vogels, et al., 2021). This suggests that increases in PE signaling in reaction to surprising stimuli represent a more likely explanation.

Interpreting the link between stimulus predictability and neuronal response is further complicated when taking confounding factors into account: Activity reductions are absent when a competing stimulus is attended instead (Richter & de Lange, 2019) and activity can be increased for predictable stimuli when they are relevant to the task performed by participants (Kok, Rahnev, et al., 2012). Following these findings, it has been argued that attentional mechanisms provide an alternative explanation to predictive processing (Alink & Blank, 2021). However, attention can be accommodated within the predictive coding framework as reflecting the precision of evidence and beliefs (Hohwy, 2012; Yon & Frith, 2021). *Precision-weighting* means that the effect of observations on our beliefs depends on the reliability we attribute them with. Importantly, precision is not only based on the signal to noise ratio of a stimulus (an exogenous source of attention) but also our beliefs about their importance and reliability (an endogenous source of attention). This mechanism might explain the apparent discrepancy of the initially mentioned studies when assuming that task-relevancy increases precision (Alamia & Zénon, 2016). Consequently, a surprising stimulus that is not relevant elicits weak PE signals because of low precision, while an expected stimulus that is relevant elicits strong activation due to high precision.

In summary, the efficiency of predictive coding can be studied through our ability to extract regularities from sequences of visual input. Humans quickly learn the rules and probabilities and react more quickly and accurately to input that confirms these patterns. Statistical learning is the dominant experimental paradigm, but most approaches can be described through the lens of Bayesian inference. fMRI studies consistently find different responses to the same stimulus when presented in a predictable compared to a surprising context. Competing findings can be accounted for when factoring in the precision of evidence and beliefs.

1.4 Subjective confidence as a proxy for uncertainty

The importance of precision in determining the integration of beliefs and evidence led to the search for its behavioral and neural indicators. However, the assessment of precision in the brain is complicated since, mathematically, it reflects a probability distribution over all competing hypotheses (Pouget et al., 2016). A promising approximation is the confidence of a participant in a given decision. In the context of Bayesian models, confidence refers to the (posterior) probability that a given hypothesis is correct, given the evidence (Hangya et al., 2016). While this does not account for the full complexity of probabilistic beliefs, confidence is informative when the number of competing hypotheses is small or centered around a dominant one (Pouget et al., 2016). Confident decisions are thus assumed to reflect situations where our observations are strongly in favor of a given assumption, with the alternatives much less likely.

Independent of statistical formulations, human confidence reports are a central concept in research on metacognition. Metacognition describes our ability to reflect on our own thoughts and decisions, providing us with a source of self-regulation without needing explicit feedback (see Fleming (2024) for a recent review). In this context, confidence is a subjective feeling regarding the quality of our judgements, decisions and actions. Subjective ratings are usually assessed by asking participants to indicate their level of confidence in a given decision on a linear scale. These ratings exhibit good within-subject reliability across days and their distribution is highly specific to an individual, akin to a “fingerprint” (Ais et al., 2016). While this subjective assessment is generally correlated with objective performance, it often deviates from it (Shekhar & Rahnev, 2021). While it has been argued that this dissociation indicates suboptimal inference in humans (Acerbi et al., 2014; Beck et al., 2012), a recent study showed that it can be explained by optimal inference with incomplete knowledge about the environment (Khalvati et al., 2021). Furthermore, studies comparing statistical formulations of confidence to subjective ratings report general agreement (Meyniel, Schlunegger, et al., 2015; Sanders et al., 2016). While Bayesian accounts of confidence are prevalent, the exact nature of how the sense of confidence emerges is still unknown, and multiple alternative models exist (Adler & Ma, 2018; Boundy-Singer et al., 2022; Pleskac & Busemeyer, 2010).

A promising neuronal source of subjective confidence is probabilistic population coding (Knill & Pouget, 2004). This idea rests on the tuning property of neurons, where the frequency of the firing rate indicates how strongly the sensory input matches the preferred stimulus characteristics (Salinas & Abbott, 1994). A population of differently tuned neurons represents

a probability density estimate over a given variable (Ma et al., 2006). Confidence ratings can then be seen as a summary statistic over this distribution, providing an estimate of the precision (Meyniel, Sigman, et al., 2015). Experimental results suggest that the sense of confidence reflects this statistic and informs adaptive behavior (Geurts et al., 2022; Van Bergen et al., 2015). Neuroanatomical areas where BOLD activity covaries with confidence are predominantly found in the frontal cortex, especially the anterior cingulate (ACC), the inferior and superior parietal cortex as well as the anterior insula (Bang & Fleming, 2018; Bounmy et al., 2023; Hebart et al., 2016). Interestingly, activity in these regions generally decreased with increasing confidence. The opposite was found for the link between uncertainty and BOLD activity (Geurts et al., 2022), suggesting that the brain increases its activity when information determining a decision is scarce or unreliable.

In summary, research on confidence suggests a link between subjective judgements, Bayesian inference and neuronal activity. Neuronal population codes can be framed as probability distributions, a summary of which is consciously accessible. High confidence has been linked to lower fMRI activity, indicating a central role in decreased neural responses to predictable input.

1.5 Evidence of predictive processing in the human brain

Due to its explanatory power, Bayesian predictive processing is a leading framework in psychology and neuroscience (Clark, 2013; Friston, 2018; Hohwy, 2013; Yon et al., 2020). However, some authors warned against equating its influence with its truthfulness, citing missing evidence for neural representations of uncertainty and Bayesian calculus (Colombo et al., 2021). While the value of the *Bayesian brain hypothesis* in explaining empirical data is less contended, the question whether Bayesian inference is neurobiologically implemented led to concerns about its validity (Rahnev, 2019; Rescorla, 2021). Some of these concerns have been addressed by works on predictive coding, especially in the context of the FEP. Hierarchical predictive coding has been mapped onto cortical circuits (Bastos et al., 2012) and empirical data shows that predictions are processed in deep cortical layers while errors travel in superficial layers (Bastos et al., 2020; Kok et al., 2016). The problem of representing complex probability distributions in the cortex can be solved by appealing to approximations via sampling methods (Sanborn & Chater, 2016) or iterative optimizations (*variational* Bayes; Friston, 2010). A promising approach to validating predictive processing theories of the mind is computational modeling. Vincent et al. (2019) used the Bayesian inference calculus

postulated by the FEP to model the perception of visual sequences. The authors manipulated both the transitional probabilities and the precision of the evidence (i.e. how reliably the stimuli represented the true underlying sequence). Their model reproduced empirical pupil diameter data, which has been shown to be indicative of subjective uncertainty (Fan et al., 2023; Lavín et al., 2014). While this application is highly specific, the same logic has been applied to electroencephalography (EEG) data during speech perception (Friston et al., 2021), supporting the general applicability of the FEP. While it has been argued that predictive processing theories are difficult to falsify due to their flexibility (Kogo & Trengove, 2015; Miłkowski & Litwin, 2022), empirical evidence for predictive processing has accumulated across methodologies (Walsh et al., 2020).

In summary, Bayesian models of the mind are the dominant framework in neuroscience and psychology. While many questions about the neural implementation of Bayesian predictive coding remain, evidence has been provided across multiple paradigms and imaging methods.

1.6 Aims of the current work

In the present work, I tested the hypothesis that the inference of predictive structure in visual sequences leads to a decrease in energy metabolic activity. Furthermore, I aimed to establish individual confidence as a moderator of the link between objective predictability and energy consumption. I adapted an established statistical learning design to be compatible with mqBOLD, a novel metabolic imaging method. mqBOLD comes with a methodological challenge that necessitates changes to previous designs. While conventional BOLD has a temporal resolution of up to one second, mqBOLD results in averages over long blocks of stimulation. Therefore, I developed a task with stable task demands across this timespan and monitored the continuous engagement of the participants. I also ensured compatibility with previous findings by validating my design with conventional fMRI. To this end, I assessed BOLD data during long blocks of stimulation. The present work was divided into two projects which served five underlying aims. To my knowledge, this is the first investigation of the proposed energetic efficiency of predictive processing.

- **Project I**. Develop and validate a new statistical learning design that assesses subjective confidence and is compatible with long blocks of stimulation.
    - **Aim I**. Establish a link between confidence ratings and markers of statistical learning after a learning phase.

- **Aim II**. Replicate previous BOLD findings of predictive processing using long stimulation blocks.
- **Aim III**. Establish a link between confidence ratings and BOLD signals representative of prediction and prediction error signaling.

- **Project II**. Measure the energy consumption during visual stimulation of three levels of predictability and investigate the effect of subjective confidence.
  - **Aim IV**. Implement a statistical analysis that allows for the integration and summation of energy consumption across the brain.
  - **Aim V**. Quantify the energy use on multiple levels of analysis, accounting for both functionally specific networks and net effects on the cortex.

2 Manuscript I: Subjective confidence modulates individual BOLD patterns of predictive processing

André Hechler[1,2], Floris P. de Lange[3], Valentin Riedl[1,2]

[1] Department of Neuroradiology, Neuroimaging Center, Technical University of Munich, Munich, Germany
[2] Graduate School of Systemic Neurosciences, Ludwig-Maximilians-University, Munich, Germany
[3] Donders Institute for Brain, Cognition and Behaviour, Radboud University Nijmegen, Nijmegen, Netherlands

Abstract

Humans are adept at extracting patterns from sequentially presented information. This ability enables predictions about future states, resulting in anticipation both on a behavioral and neural level. Stimuli deviating from predictions usually evoke higher neural and hemodynamic activity than predicted stimuli, an effect which has been termed prediction error. However, interindividual differences in learning and uncertainty have rarely been taken into account. Under Bayesian formulations of cortical function, prediction errors should be stronger if a subject was highly confident that another stimulus should have appeared. While this is supported by studies employing Bayesian observers, prediction error signals in brain imaging have not been shown to scale with subjective confidence ratings.

In the present study, participants viewed visual object sequences of varying predictability over multiple days. Additionally, we intermittently prompted them to rate their confidence after completing partially presented sequences, resulting in a range of confidence levels after the training phase. In the following MR scanning phase, participants saw sequences that either confirmed predictions, deviated from them or were completely random. We replicated BOLD findings of prediction error activity in the ventral visual stream and found that their magnitude increased with confidence ratings. Furthermore, we observed the opposite for the contrast between predictable and random sequences: In the anterior cingulate, predictable sequences elicited higher activity for low level of confidence, but lower activity for high levels of confidence. These results show that group-level findings of predictive processing contain considerable variance that can be explained by differences in learning and inference.

Introduction

Our visual experience evolves as a continuous stream of sensory states. In the natural world, states close in time tend to be correlated, providing a probabilistic mapping of visual input from one moment to the next. Studies presenting temporal sequences of visual stimuli show that humans and animals track the underlying transitional probabilities (Sherman et al., 2020; Turk-Browne et al., 2009a). After a training phase, both behavior and neural patterns indicative of prediction (or anticipation) emerge: Participants detect and categorize expected stimuli faster (Turk-Browne et al., 2010) and stimulus templates of expected stimuli can be decoded from fMRI activity (Kok et al., 2014; Kok, Jehee, et al., 2012). Interestingly, the activity elicited by a given stimulus is increased when its occurrence violates previously presented patterns (Kaposvari et al., 2018; Manahova et al., 2018; Meyer & Olson, 2011; Richter et al., 2018b). This increase likely represents a prediction error (PE) – the mismatch between predictions and evidence based on an internal model of the process (Keller & Mrsic-Flogel, 2018; Rao & Ballard, 1999).

Most studies reporting PE activity focused on manipulating the parameters of the generative process, i.e. the true transitional probabilities of the sensory stream. However, the human inference process is subject to uncertainty (Hasson, 2017; O'Reilly, 2013), possibly implemented by a Bayesian integration of prior and evidence (Knill & Pouget, 2004; Pouget et al., 2013). Consequently, Bayesian predictive coding suggests that the magnitude of PEs depends both on the reliability of the evidence (the consistency of patterns and the signal to noise ratio) and the precision of our predictions (the probability of a prediction among all hypotheses) (Jiang & Rao, 2022 but see Aitchison & Lengyel, 2017 for non-Bayesian interpretations). Two approaches to assess the latter are prevalent in the literature: First, corresponding parameters can be extracted from a computational model performing the task at hand. Second, participants can be prompted to rate their subjective confidence in a given decision.

Confidence ratings are usually tied to an overt or covert decision by the participant. For this reason, they are usually not assessed when studying the learning of temporal patterns, which often aims at implicit processes (Schapiro & Turk-Browne, 2015; Turk-Browne et al., 2010). Nevertheless, a promising line of studies combined both aspects by interjecting explicit predictions and confidence ratings between long blocks of visual stimulation. The authors showed that these separately acquired ratings reflect the confidence of an ideal Bayesian observer during the stimulation phase (Meyniel, Schlunegger, et al., 2015; Meyniel & Dehaene,

2017). Crucially for the present study, BOLD activity covaries with both surprise and confidence in a range on neuroanatomical regions (Bounmy et al., 2023). However, direct evidence that imaging-derived PE activity scales with confidence is not available yet.

In the present study, we aimed to show that PE activity scales with subjective confidence, building on a previous study using visual streams of everyday objects (Richter et al., 2018b). Furthermore, we addressed whether PE activity emerges as a function of increasing activity for surprising input (surprise enhancement) or decreasing activity for predictable input (expectation suppression). Evidence for the latter is inconclusive in humans with only few fMRI studies looking into the alternative explanations (Feuerriegel, Vogels, et al., 2021). To address this gap, we included three experimental conditions regarding the visual stream: Fully predictable, fully unpredictable (random) and surprising. Following previous work, activity decreases for predictable compared to unpredictable stimulation served as evidence for expectation suppression (Manahova et al., 2018; Ramachandran et al., 2017). Interestingly, the same contrast has been used to study the brain's differential response to order and disorder in sensory input (Davis & Hasson, 2018; Nastase et al., 2014), a fundamental aspect of predictability. Areas that are more sensitive to structure might serve as prediction encoding hubs (Ficco et al., 2021). We expected BOLD activity in these areas to decrease with confidence, since new evidence will not lead to prediction updating.

Results

Confidence increases for predictable input and explains improvements in behavior

For the visual stimulation, we used sequences of visual images comprised of five everyday objects in full color. Prior to the scanning session, participants completed a three-day training phase with 20 minutes of stimulation per day (Figure 1). This phase only included predictable and unpredictable sequences. After each day, we sampled the learning process with a sequence completion task: Participants saw incomplete sequences and had to choose the correct trailing object from a set of five options. Each trial was followed by a confidence rating prompt. We found that average confidence regarding predictable sequences increased over the training phase while it remained constant for unpredictable ones (Figure 2a). Objective accuracy in completing predictable sequences also increased over days (Figure 2b) and was highly correlated with confidence ($r=0.8$, $p<0.001$). If not indicated otherwise, the following analyses used average confidence ratings in predictable sequences after the last training day. To ensure visual fixation and attentiveness during the stimulation blocks, we

instructed participants to react to upside-down images with a button-press (the *cover task*). Neither reaction time nor accuracy were significantly different between conditions (reaction time: t(*41*)=-1.2, p=.24; accuracy: t(*41*)=1.55, p=.13).

We also tested whether learning the transitional probabilities led to behavioral advantages. To this end, we included a short task before the MR acquisition. We presented the participants with a target object that would occur in the following object sequence. Sequence presentation followed the same design as in the training phase, but instead of performing the cover task, participants were instructed to react to the occurrence of the target object. Reaction times were nearly 10% faster for predictable objects (median RT change=-9.93%, t(41)=-5.66, p<0.001), with a maximum gain of 48%. We also found that subjects with higher confidence had significantly larger gains in reaction time (Figure 2c).

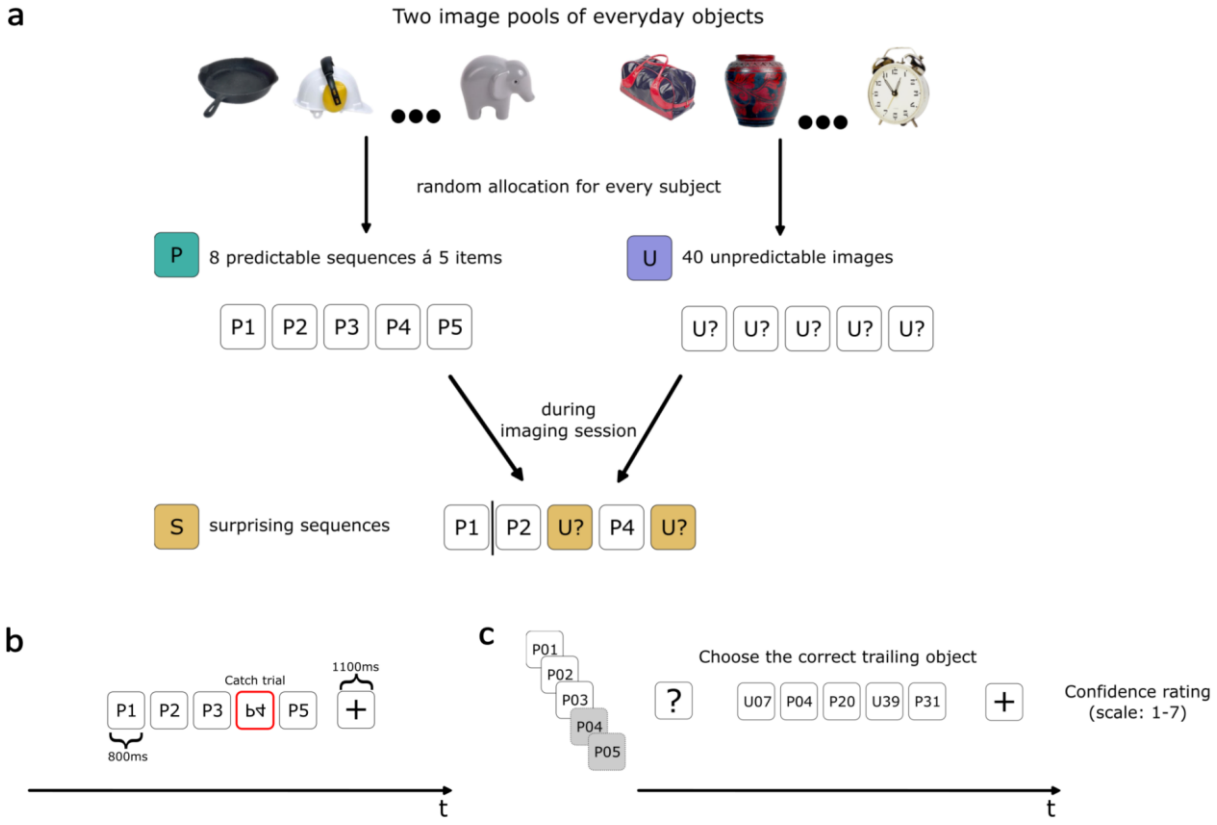

Figure 1. Experimental design. **a**. Creation of three types of object sequences. Predictable (P, green) and unpredictable (U, purple) sequences were shown during training. Surprising (S, yellow) sequences were only shown during the scanning session. See methods section for details. **b**. Objects were shown without ISI and whole sequences were separated by a fixation cross. Upside-down objects occurred with a probability of 10% irrespective of condition. **c**. Sequence-completion task. One to four of the leading objects were presented and the correct trailing object had to be selected from five options. The selection was followed by a confidence prompt. This task was not timed. Both predictable and unpredictable sequences were shown to keep object exposure comparable. We presented eight sequences per condition.

During fMRI scanning, we presented participants with blocks of four sequences per condition, with block order randomized. In addition to the predictable and unpredictable conditions, we included surprising blocks. These were based on predictable sequences from the learning phase but had one to three objects replaced (Figure 1a). The scanning phase included no sequence-completion task or confidence ratings. Analysis of the cover task revealed that reaction times were faster in predictable than unpredictable blocks (Figure 2d). The same trend was present, but not significant, compared to the surprising blocks ($t(41)=1.94$, $p_{FWE}=0.18$). The accuracy of reactions did not vary between conditions. Since the cover task was independent of the predictability of the sequences, we explored whether reaction times were impacted by confidence levels. We calculated condition-wise correlations and found that high confidence in the predictable sequences was associated with quicker reaction times (Figure 2e, left). This link was not significant in the other conditions. Seeing this possibly confounding effect, we included reaction times in the following BOLD analyses.

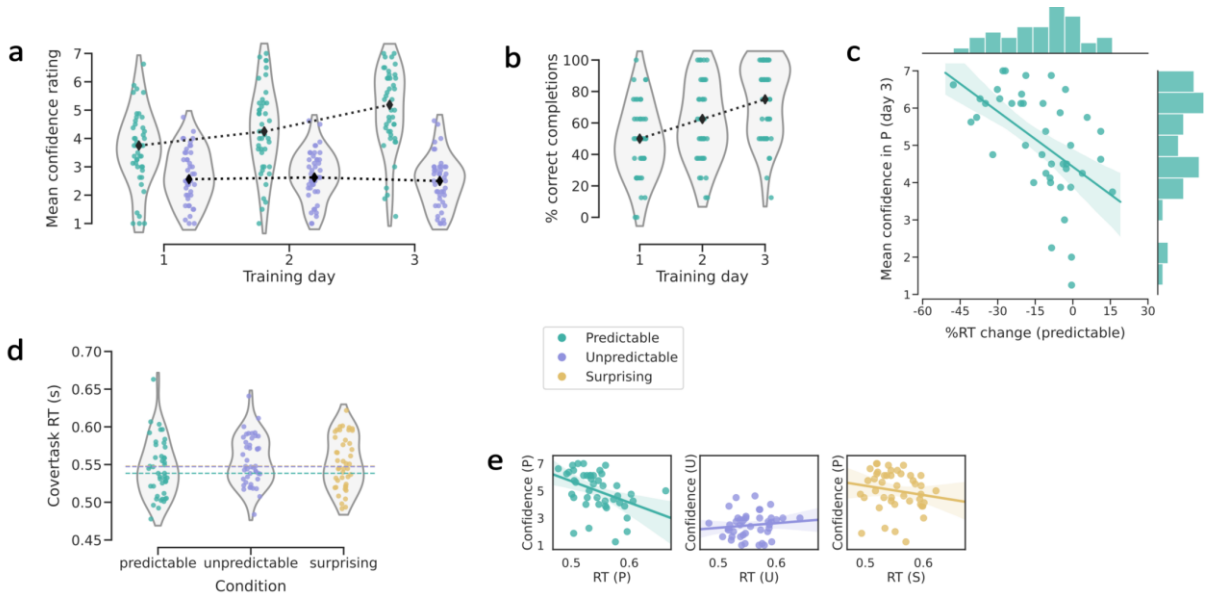

Figure 2. Training and cover task results. The central color legend applies to all subfigures. **a**. Average subject-wise confidence ratings within conditions after respective training days. Black markers show daily medians. **b**. Average percentage of correctly completed sequences. Unpredictable sequences could not be predicted over chance level and are not shown here. **c**. Significant linear regression between the relative difference in reacting to predictable versus unpredictable target objects and average confidence ratings for predictable sequences prior to scanning ($r=-0.536$, $p<0.001$). Negative values on the x-axis correspond to quicker reactions for predictable objects. **d**. Subject-wise averages of reaction time in the cover task during fMRI acquisition. Dashed horizontal lines show respective medians, with a significant difference between the predictable and unpredictable condition (post-hoc t-test: $t(41)=2.59$, $p_{FWE}=0.04$). Note that the lines for the unpredictable and surprising conditions overlap since the median is nearly the same. **e**. Condition-wise correlations between confidence ratings and reaction time in the cover task during scanning. Since surprising sequences were based on predictable ones during scanning, we used the corresponding confidence ratings. The relationship was only significant for the predictable condition ($r=-0.5$, $p_{FWE}=0.002$).

Prediction errors in sensory areas and prediction activity in parietal areas

We analyzed two contrasts with respect to the BOLD data: A PE contrast (surprising > predictable) and a prediction contrast (predictable > unpredictable). We adopt the terminology of "prediction contrast" in reference to non-error related aspects of predictive processing. Confidence was included as a regressor of interest and reaction time differences between the respective conditions as a confound regressor (Methods). First, we inspected the main contrasts between conditions. PE activity was present bilaterally throughout the ventral stream, in close correspondence to previous work using an event-based design after implicit statistical learning (Figure 3a, left; Richter et al., 2018). The prediction contrast did not provide evidence of expectation suppression, as we observed no significant negative clusters (i.e. predictable < unpredictable). Note that these results are based on a conservative whole-brain analysis. However, we found two positive clusters: A medial cluster in the posterior cingulate cortex and superior parietal cortex and a left lateral cluster covering parts of the superior parietal cortex, inferior parietal cortex and a small area of the dorsal stream (Figure 3b, left).

Confidence explains interindividual differences in error and prediction activity

We investigated the link between confidence and predictive processing patterns in both a confirmatory and an exploratory way. Since we assumed that PE activity increases with confidence, we regressed the subject-wise average contrast parameter estimates (COPEs) in the significant PE clusters on confidence ratings, controlling for reaction time differences. In accordance with our assumptions, we found that PE activity significantly increases with the participants' confidence level (Figure 3a, right). We then performed a corresponding analysis for the positive clusters of the prediction contrast. If these increases indicated PE activity, possibly due to uncertainty in the predictable sequences by low-confidence subjects, we would expect a similar relationship to confidence. However, we found no evidence for such a link (Figure 3b, right). A corresponding Bayesian correlation analysis revealed moderate evidence for the null hypothesis (BF=.195).

For the exploratory analysis, we analyzed the confidence regressor estimates of the second-level GLMs across the whole brain. No significant clusters emerged for the link between PE activity and confidence. The uncorrected maps indicate peaks in the ventral visual stream, anterior insula and inferior frontal cortex bilaterally (Figure 3c). Regarding the link to the prediction COPEs, we found a negative cluster in the anterior cingulate cortex (ACC; Figure 3d, left), although it did not survive FWE correction for two-sided cluster testing. Within this cluster, we found a strong negative association between COPEs and confidence. Interestingly,

the corresponding regression line crosses the zero value with respect to the COPEs (Figure 3d, right). Descriptively, this indicates that the relative difference flips sign – for low-confidence participan7ts, predictable blocks elicited higher BOLD signal compared to unpredictable blocks while the opposite happened for high-confidence subjects (Methods).

Finally, the directionality of the relationship with confidence generally depended on the contrast: Inspecting the uncorrected parameter maps, we see that PE activity increases with confidence while the inverse can be seen for the prediction contrast activity (with a single exception in the inferior frontal cortex, Figure 3d, left).

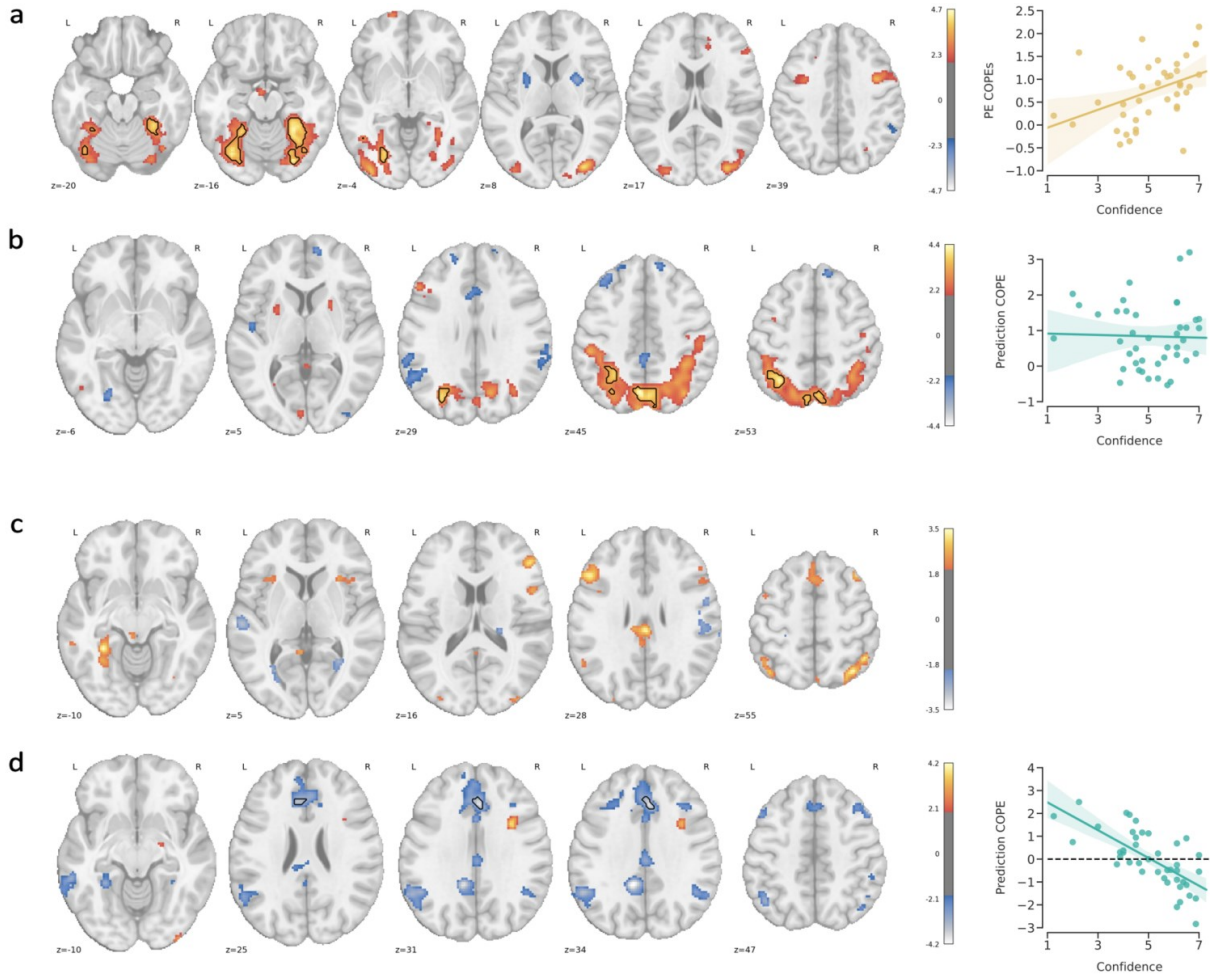

Figure 3. GLM group level results for main effects (a and b) and the confidence regressor (c and d). Shown are uncorrected t-values thresholded at z=2.33 (p<0.01). Clusters surviving cluster correction (cluster-forming threshold: z=3.09, cluster threshold: p_{FWE}=0.025) are shown in black contours. **a**. Left: Surprising > predictable contrast maps. Right: Confidence significantly explains differences in average COPEs ($R^2_{partial}$=.15, t(*41*)=2.54, p=.015). Data shown are subject-wise average confidence ratings for the predictable condition (x-axis) and subject-wise COPE averages within the significant clusters shown left. **b**. Corresponding results for the predictable > unpredictable contrast. The linear model yielded no significant results ($R^2_{partial}$=.001, t(*41*)=-0.24, p=.82). **c**. Statistical map for the confidence regressor of the prediction error GLM . Parameter estimates indicate how strongly PE COPEs vary with confidence levels. **d**. Left: Corresponding statistical map for the prediction contrast. Cluster contours show a negative cluster (p=0.043) which did not survive FWE correction for two-sided cluster

correction. Right: Confidence explained considerable variance of the COPEs in this cluster ($R^2_{partial}$=.49). No test for significance was carried out because the ROI was based on the GLM results.

Discussion

In the present paper, we established a link between brain activation patterns of predictive processing during the presentation of temporal sequences and subjective confidence in knowing the underlying patterns. Importantly, confidence ratings were derived separately from the visual stimulation. While this approach deviates from similar work using statistical learning, we replicated BOLD findings from a recent study where participants were kept unaware of the patterns (Richter et al., 2018b). In line with our assumptions, PE activity in the ventral visual stream increased with confidence in knowing the (previously reliably predictable) object sequences. In the ACC, we found tentative evidence that the activity during predictable input relative to unpredictable input depended on confidence levels: For high confidence levels, activity for predictable input was lower. Conversely, for low confidence levels, unpredictable input elicited lower activity levels. These results stress the importance of accounting for interindividual differences in uncertainty during partly implicit predictive processes.

Paradigms for learning of temporal regularities

Our ability to extract regularities in the visual domain has been studied from various perspectives, leading to subtly different paradigms (Conway, 2020). These include statistical learning (Fiser & Aslin, 2002), sequence or sequential learning (Gheysen & Fias, 2012), probabilistic learning (Meyniel, Schlunegger, et al., 2015) or temporal community learning (Schapiro et al., 2013). Instead of being fundamentally different, these approaches can be described as gradually varying on three axes (Conway, 2020): Amount of exposure, amount of structure and amount of instruction or feedback. Importantly, unifying frameworks have been developed, arguing that Bayesian inference models can account for most variants of acquiring regularities from sensory stimulation (Fiser & Lengyel, 2019; Konovalov & Krajbich, 2018). While the present study did not aim to resolve competing perspectives, it focused on the latter account. Here, we combined an extended learning phase of continuous sensory stimulation, using initially deterministic sequences and light instructions with a separate explicit performance task.

Prediction error activity

We first discuss the results of our main contrasts with respect to previous studies. Firstly, we found evidence for PE activity throughout the ventral visual stream. Similar to

previous studies in humans, there was no evidence of expectation suppression (Feuerriegel, Vogels, et al., 2021). In our data, the process behind PE activity is thus better explained as upregulated processing for surprising stimuli as opposed to downregulated processing for predictable ones. However, we here used a whole-brain approach that is more conservative than targeted analyses of the visual stream.

Our findings bear striking similarity to a previous study (Richter et al., 2018b) despite differences in pattern length, instruction and participant awareness. This suggests that the underlying neural processes are not exclusive to incidental or implicit learning. Differences between neuroanatomical substrates of implicit and explicit learning point to more frontal activations for the latter, while implicit conditions predominantly affect sensory areas (Gheysen & Fias, 2012). However, these systems may work in parallel, with the implicit processes as a prerequisite for the explicit ones (Batterink et al., 2015). Consequently, one might expect additional activations in higher cognitive areas for our design. While not surviving cluster correction, we did find corresponding clusters in the precentral gyrus and inferior frontal cortex bilaterally (Figure 3a). It is possible that our study is underpowered to detect these effects. However, our GLM contrast was not designed to detect general patterns of statistical learning, but specifically isolated PE processing after a multi-day training phase. Related to the implicit versus explicit divide, there is evidence that surprising stimuli only evoke higher activity when unattended (Kok et al., 2012). In contrast, our data suggests that PE activity persists even when participants were explicitly asked to reproduce the patterns. This was corroborated by reaction time data in our cover task: Reaction times were shortest for predictable sequences, suggesting that, even though participants were highly engaged here, surprising sequences still evoked a stronger signal. Of note, our study differs from previous work by using a block design. That means that our surprising condition represents a visual stream including both confirmations and violations of expectation. The subtleties of transient, event-based responses might thus deviate from our findings.

Prediction activity

Previous research on the contrast between predictable and unpredictable stimulation produced varied results. Stimuli predictive of other stimuli can elicit either increased activity (Egner et al., 2008) or decreased activity (Davis & Hasson, 2018), potentially dependent on task relevance (Richter & de Lange, 2019). Neuroanatomical findings partly overlap with the ACC being implicated irrespective of effect directionality. In our data, we found extended increases for predictable input focusing on the superior parietal cortex, intraparietal sulcus

and posterior cingulate. The intraparietal sulcus has been found to increase activity in response to predictive stimuli (Egner et al., 2008) and is generally sensitive to the entropy of visual and auditory sequences (Nastase et al., 2014). Interestingly, the superior parietal sulcus is central to a wide range of predictive processing tasks, as shown by a recent modality- and task-general meta-study (Ficco et al., 2021). As an interim summary of our main effects, we corroborated various findings in the literature that establish modality-specific areas of PE activity and modality-general areas of predictive processing. In the following paragraphs, we introduce subjective confidence as a moderator of the direction and magnitude of these effects.

Interpreting the effect of confidence

Confidence ratings are often studied in the context of decisions, where they reflect the posterior probability of being correct, given evidence and an internal model (Pouget et al., 2016, but see Adler & Ma, 2018 for a non-Bayesian view). However, a line of studies (Bounmy et al., 2023; Meyniel, Schlunegger, et al., 2015; Meyniel & Dehaene, 2017) showed that confidence ratings can also be used to sample the uncertainty during continuous perceptual inference. This suggests that confidence is an ever-present aspect of perception, which can be cast as ongoing inference about the correspondence of our internal model (and the associated predictions) with sensory input (Friston, 2010). However, while many studies explored imaging patterns of expectation and surprise, surprising trials were often defined by the objective deviation of a presented stimulus from preceding patterns (Kaposvari et al., 2018; Manahova et al., 2018; Meyer & Olson, 2011; Richter et al., 2018b). We suggest that much of the variance in individual response can be explained by accounting for individual (un)certainty in the presented patterns. As a simple example in the context of simpler paradigm, presenting a sequence of nine house images followed by a face image might lead to varying levels of surprise if observers did not infer the same conditional probability of observing a tenth house image. Given more complex patterns, as used in our study, interindividual variability is expected to increase. While the underlying reasons are beyond the scope of this study, the frequent divergence of human inference and behavior from (Bayes-)optimal models has been discussed previously (Acerbi et al., 2014; Beck et al., 2012).

Confirming our main hypothesis, PE activity in the significant clusters across the visual stream scaled with confidence. Previous studies investigating this link used model parameters derived from an ideal Bayesian observer, finding the strongest links between surprise and confidence outside of the visual cortex (Bounmy et al., 2023). To our knowledge, only two

other studies explored the link between imaging indicators of PE and confidence ratings: In an EEG analysis of a perceptual decision task, parietal event related potentials covaried with confidence ratings both during stimulus presentation and response (Boldt & Yeung, 2015). An fMRI study using a visual search paradigm found no evidence of PE scaling during stimulus presentation but did find an inverse relationship during the participants response in the inferior frontal gyrus (Sherman et al., 2016). To extend our findings, we performed an exploratory whole-brain analysis but did not find significant clusters.

We tested whether prediction activity scaled with confidence but found no relationship within the clusters of the contrast between predictable and unpredictable sequences. Mirroring the previous analysis, we then performed a whole-brain analysis on the voxel level. This yielded a cluster in the ACC where nearly 50 percent of the variance was explained by confidence levels. The direction of this effect was opposite to the PE contrast: Here, contrast estimates decreased as a function of confidence. Since this cluster did not survive correction for multiple comparisons, the results should be interpreted with care. However, evidence for the role of the ACC in predictive processing and conflict processing in general is widespread throughout the literature (Alexander & Brown, 2019; Garrison et al., 2013; Vanveen & Carter, 2002). Consequently, we discuss some implications of this finding in the following lines. Recently, the ACC has been described as part of a prefrontal network that anticipates PEs across cortical hierarchies (Alexander & Brown, 2018). This suggests a decrease in activity with increasing confidence, which has been confirmed for probabilistic learning tasks (Bounmy et al., 2023; Davis & Hasson, 2018). However, if the ACC is central to PE processing, we would expect corresponding results regarding our PE contrast. Possible explanations for this dissociation might lie in the hierarchical nature of predictive processing (Alexander & Brown, 2018): PEs in the visual stream indicate low-level mismatches without necessary awareness, while errors in higher cortices are subject to conscious reflection and, possibly, anticipation. Future studies are necessary to differentiate the modulating effect of confidence across the hierarchy of predictive processing.

Lastly, the confidence parameter estimates in the ACC flipped sign at medium-to-high levels of confidence. This means that predictable sequences elicited higher activity than unpredictable ones for low confidence levels, while the opposite was true for high levels. Regarding low levels of confidence, this can be explained by precision weighting of expected errors (Yon & Frith, 2021): If participants inferred that there is a ground truth to the predictable

sequences while not being able to accurately predict the order, PEs of high precision can be expected. Due to the stochastic nature of the unpredictable condition, stimuli can never be accurately predicted, but the precision of errors is low. Nevertheless, if the unpredictable condition did elicit residual error activity, it might explain the relatively higher activity for high confidence levels. Here, the correspondence between predictions and sensory information is (close to) perfect in the predictable condition, triggering little to no error anticipation or processing. Finally, it is possible that even after three days of exposure, subjects still assumed that there is some order in randomness (Huettel et al., 2002), leading to incorrect predictions of comparatively high precision. Tentative evidence for this assumption can be found in the confidence ratings for the unpredictable condition. While these did not change during training, they were still far from minimal, indicating that some of the subjects did not correctly infer that these sequences were inherently unpredictable.

Methods

Participants

We recruited 44 participants including students, doctoral researchers, clinic staff and the general population of Munich. The sample size was based on a-priori analysis, aiming to detect at least a medium effect size (d>=0.5, alpha=0.05, beta=0.9). All participants took part in a familiarization MRI session (20 minutes), an online training phase over three days and the main MRI session (70 minutes) on the day after training completion. Two participants were excluded from the analysis due to technical problems during data acquisition. The remaining 42 participants (19 female, age [mean(std)] = 27.1(3.9)) were included for analysis if not indicated otherwise. The study was approved by the ethics board of the Technical University of Munich (TUM), and we acquired written informed consent from all participants.

Visual stimuli

We selected 224 full-color images of everyday objects from a larger image set (Brady et al., 2008). The stimuli were chosen to be maximally homogenous regarding salience, e.g., by excluding food items and bright colors. If not stated otherwise, all allocations of stimuli to subjects and conditions were random. A set of 80 images was drawn for every subject, with half assigned to the predictable condition and the other half to the unpredictable condition.

Experimental design

Implementation. We used Psychopy (Peirce et al., 2019) to implement the design. The online training sessions were realized using pavlovia.org, where Javascript-translated Psychopy

experiments can be run online with millisecond precision (Bridges et al., 2020; Sauter et al., 2020).

Main task. Participants saw sequences of five objects each, each stimulus presented for 800ms with no inter-stimulus interval and a 1100ms fixation cross between sequences. Three types of sequences were presented: Predictable sequences had the same order for every repetition. Unpredictable sequences were randomly drawn from the subject-specific set of unpredictable images. After all unpredictable images were shown, new sequences were formed. Surprising sequences were based on predictable sequences but had one to three stimuli replaced with random unpredictable images upon every repetition. The visual stream was reshuffled prior to repeated presentation based on two rules: Objects could only repeat after all objects in the respective condition were presented, and objects (or sequences) could not appear twice in a row. This was done to minimize the confounding effects of repetition suppression. To ensure fixation and concentration, participants were instructed to quickly react to occurrences of upside-down objects. These appeared with a probability of 10%, independent of condition.

Training phase. Over the course of three days prior to the main scanning session, participants followed an online implementation of the main design. This phase only included predictable und unpredictable sequences. We instructed every participant to perform the training in a quiet environment without distractions. Stimulation blocks lasted 20 minutes per day, with a break of one minute after 10 minutes. Instructions were shown on screen before the stimulation began and the first day included a five-minute familiarization block with on-screen feedback regarding cover-task button presses. To measure objective and subjective progress in statistical learning, a testing phase followed each training day. Here, participants saw eight incomplete sequences for both conditions. After every sequence, participants chose what they assumed would be the correct trailing object from five options and gave a confidence rating on a scale of one to seven. No feedback was given regarding performance and no information on the underlying conditions was disclosed.

Imaging phase. During scanning, stimuli were presented against a grey background and subtended 4° of visual angle. We chose a block design with four successive sequences of one condition forming one block. The order of blocks was randomized, as were sequences within blocks. Participants saw 12 blocks per condition with a duration of 20.4s each, adding up to a

run time of 12 minutes and 23.4 seconds. We disclosed no information regarding the sequence patterns and the imaging session included no sequence completion test.

MRI acquisition

High-resolution structural scans were acquired with a T1-weighted 3D MPRAGE sequence (170 slices, voxel size = 1.0x1.0x1.0mm³, FOV = 250x256x170mm³, TR/TE/flip angle=9ms/4ms/8°). fMRI data was acquired using single-shot EPI (40 slices, voxel size = 3.0x3.0.3.0mm^3, FOV 192x192x127.8mm^3, TR/TE/flip angle=1200ms/30ms/70°) with 612 dynamic scans plus 2 dummy scans (total duration: 768 seconds).

MRI data preprocessing

We preprocessed both structural and functional MRI data with fMRIPrep (Esteban, 2019) which produces an automated processing description that we provide in the supplement.

Statistical analysis

We performed whole-brain BOLD data analysis using FSL FEAT via a Python interface (adapted from Esteban et al. 2020). Regarding the first level analysis, we modelled the experimental manipulation as blocks, following their presentation timing with a length of 20.4 seconds (four object sequences) and included the respective temporal derivatives. The resulting boxcar functions were convolved with a double gamma hemodynamic response function. We added eight motion regressors: Six translations and rotations as well as two measures of bulk head motion (DVARS and framewise displacement). Finally, we included the average global CSF and white matter signal. All regressors were taken from the fMRIPrep results. The data were grand mean scaled, smoothed with a 6mm³ kernel and high pass filtered at 120 seconds. Since we did not have a resting condition, we created two GLMs: One leaving the surprising blocks unmodelled (used to create the prediction contrast between the modelled blocks) and one leaving the unpredictable blocks unmodelled (used to create the PE contrast). This was done to prevent overparametrization of the design matrix which can lead to unstable estimates.

Subsequent group-level analyses were run using FSL's FLAMEO and included two covariates: Confidence ratings from predictable sequences (individual average after the last training day) and the percent change in reaction time to catch trials between the conditions being contrasted (individual average over all blocks). Since we acquired separate statistical maps from the first level analyses, we ran two corresponding analyses on the group level. This influences the interpretation of the covariates: The effect of confidence is now specific to the

respective contrast (i.e. explained variance regarding the prediction contrast or PE contrast). Statistical maps of significant activation were obtained using FSLs cluster correction with a voxel threshold of z=3.09 and a cluster threshold of p=0.05 (Woo et al., 2014). Since cluster correction is applied for negative and positive clusters separately, we used an FWE-corrected threshold of 0.025. Our interpretation of contrast parameter estimates (COPEs) rests on the underlying subtraction of two parameter estimates. The resulting COPE thus informs both about the magnitude of the difference as well as the direction (which parameter weight is larger). Comparisons regarding the anatomical localization were based on relative voxel coverage over regions in the Harvard-Oxford Cortical Probabilistic Atlas.

All analyses on behavioral data were performed using the python package pingouin (Vallat, 2018). In linear regressions, the partial R^2 values for confidence predictors were based on a partitioning of explained variance following Grömping (2006).

References

Acerbi, L., Vijayakumar, S., & Wolpert, D. M. (2014). On the Origins of Suboptimality in Human Probabilistic Inference. *PLoS Computational Biology*, *10*(6), e1003661. https://doi.org/10.1371/journal.pcbi.1003661

Adler, W. T., & Ma, W. J. (2018). Comparing Bayesian and non-Bayesian accounts of human confidence reports. *PLOS Computational Biology*, *14*(11), e1006572. https://doi.org/10.1371/journal.pcbi.1006572

Aitchison, L., & Lengyel, M. (2017). With or without you: Predictive coding and Bayesian inference in the brain. *Current Opinion in Neurobiology*, *46*, 219–227. https://doi.org/10.1016/j.conb.2017.08.010

Alexander, W. H., & Brown, J. W. (2018). Frontal cortex function as derived from hierarchical predictive coding. *Scientific Reports*, *8*(1). https://doi.org/10.1038/s41598-018-21407-9

Alexander, W. H., & Brown, J. W. (2019). The Role of the Anterior Cingulate Cortex in Prediction Error and Signaling Surprise. *Topics in Cognitive Science*.

Batterink, L. J., Reber, P. J., Neville, H. J., & Paller, K. A. (2015). Implicit and explicit contributions to statistical learning. *Journal of Memory and Language*, *83*, 62–78. https://doi.org/10.1016/j.jml.2015.04.004

Beck, J. M., Ma, W. J., Pitkow, X., Latham, P. E., & Pouget, A. (2012). Not Noisy, Just Wrong: The Role of Suboptimal Inference in Behavioral Variability. *Neuron*, *74*(1), 30–39. https://doi.org/10.1016/j.neuron.2012.03.016

Boldt, A., & Yeung, N. (2015). Shared Neural Markers of Decision Confidence and Error Detection. *The Journal of Neuroscience*, *35*(8), 3478–3484. https://doi.org/10.1523/JNEUROSCI.0797-14.2015

Bounmy, T., Eger, E., & Meyniel, F. (2023). A characterization of the neural representation of confidence during probabilistic learning. *NeuroImage*, *268*, 119849. https://doi.org/10.1016/j.neuroimage.2022.119849

Brady, T., Konkle, T., Alvarez, G. A., & Oliva, A. (2008). Visual long-term memory has a massive storage capacity for object details. *Proceedings of the National Academy of Sciences*, *105*(38), 14325–14329. https://doi.org/10.1073/pnas.0803390105

Bridges, D., Pitiot, A., MacAskill, M. R., & Peirce, J. W. (2020). The timing mega-study: Comparing a range of experiment generators, both lab-based and online. *PeerJ*, *8*, e9414. https://doi.org/10.7717/peerj.9414

Conway, C. M. (2020). How does the brain learn environmental structure? Ten core principles for understanding the neurocognitive mechanisms of statistical learning. *Neuroscience & Biobehavioral Reviews*, *112*, 279–299. https://doi.org/10.1016/j.neubiorev.2020.01.032

Davis, B., & Hasson, U. (2018). Predictability of what or where reduces brain activity, but a bottleneck occurs when both are predictable. *NeuroImage*, *167*, 224–236. https://doi.org/10.1016/j.neuroimage.2016.06.001

Egner, T., Monti, J. M. P., Trittschuh, E. H., Wieneke, C. A., Hirsch, J., & Mesulam, M.-M. (2008). Neural Integration of Top-Down Spatial and Feature-Based Information in Visual Search. *Journal of Neuroscience*, *28*(24), 6141–6151. https://doi.org/10.1523/JNEUROSCI.1262-08.2008

Esteban, O. (2019). fMRIPrep: A robust preprocessing pipeline for functional MRI. *Nature Methods*, *16*, 14.

Esteban, O., Ciric, R., Finc, K., Blair, R. W., Markiewicz, C. J., Moodie, C. A., Kent, J. D., Goncalves, M., DuPre, E., Gomez, D. E. P., Ye, Z., Salo, T., Valabregue, R., Amlien, I. K., Liem, F., Jacoby, N., Stojić, H., Cieslak, M., Urchs, S., … Gorgolewski, K. J. (2020).

Analysis of task-based functional MRI data preprocessed with fMRIPrep. *Nature Protocols*, *15*(7), 2186–2202. https://doi.org/10.1038/s41596-020-0327-3

Feuerriegel, D., Vogels, R., & Kovács, G. (2021). Evaluating the evidence for expectation suppression in the visual system. *Neuroscience & Biobehavioral Reviews*, *126*, 368–381. https://doi.org/10.1016/j.neubiorev.2021.04.002

Ficco, L., Mancuso, L., Manuello, J., Teneggi, A., Liloia, D., Duca, S., Costa, T., Kovacs, G. Z., & Cauda, F. (2021). Disentangling predictive processing in the brain: A meta-analytic study in favour of a predictive network. *Scientific Reports*, *11*(1), 16258. https://doi.org/10.1038/s41598-021-95603-5

Fiser, J., & Aslin, R. N. (2002). Statistical learning of higher-order temporal structure from visual shape sequences. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *28*(3), 458–467. https://doi.org/10.1037//0278-7393.28.3.458

Fiser, J., & Lengyel, G. (2019). A common probabilistic framework for perceptual and statistical learning. *Current Opinion in Neurobiology*, *58*, 218–228. https://doi.org/10.1016/j.conb.2019.09.007

Friston, K. (2010). The free-energy principle: A unified brain theory? *Nature Reviews Neuroscience*, *11*(2), 127–138. https://doi.org/10.1038/nrn2787

Garrison, J., Erdeniz, B., & Done, J. (2013). Prediction error in reinforcement learning: A meta-analysis of neuroimaging studies. *Neuroscience & Biobehavioral Reviews*, *37*(7), 1297–1310. https://doi.org/10.1016/j.neubiorev.2013.03.023

Gheysen, F., & Fias, W. (2012). Dissociable neural systems of sequence learning. *Advances in Cognitive Psychology*, *8*(2), 73–82. https://doi.org/10.5709/acp-0105-1

Grömping, U. (2006). Relative Importance for Linear Regression in *R*: The Package **relaimpo**. *Journal of Statistical Software*, *17*(1). https://doi.org/10.18637/jss.v017.i01

Hasson, U. (2017). The neurobiology of uncertainty: Implications for statistical learning. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *372*(1711), 20160048. https://doi.org/10.1098/rstb.2016.0048

Huettel, S. A., Mack, P. B., & McCarthy, G. (2002). Perceiving patterns in random series: Dynamic processing of sequence in prefrontal cortex. *Nature Neuroscience*, *5*(5), 485–490. https://doi.org/10.1038/nn841

Jiang, L. P., & Rao, R. P. N. (2022). Predictive Coding Theories of Cortical Function. In L. P. Jiang & R. P. N. Rao, *Oxford Research Encyclopedia of Neuroscience*. Oxford University Press. https://doi.org/10.1093/acrefore/9780190264086.013.328

Kaposvari, P., Kumar, S., & Vogels, R. (2018). Statistical Learning Signals in Macaque Inferior Temporal Cortex. *Cerebral Cortex*, *28*(1), 250–266. https://doi.org/10.1093/cercor/bhw374

Keller, G. B., & Mrsic-Flogel, T. D. (2018). Predictive Processing: A Canonical Cortical Computation. *Neuron*, *100*(2), 424–435. https://doi.org/10.1016/j.neuron.2018.10.003

Knill, D. C., & Pouget, A. (2004). The Bayesian brain: The role of uncertainty in neural coding and computation. *Trends in Neurosciences*, *27*(12), 712–719. https://doi.org/10.1016/j.tins.2004.10.007

Kok, P., Failing, M. F., & de Lange, F. P. (2014). Prior Expectations Evoke Stimulus Templates in the Primary Visual Cortex. *Journal of Cognitive Neuroscience*, *26*(7), 1546–1554. https://doi.org/10.1162/jocn_a_00562

Kok, P., Jehee, J. F. M., & de Lange, F. P. (2012). Less Is More: Expectation Sharpens Representations in the Primary Visual Cortex. *Neuron*, *75*(2), 265–270. https://doi.org/10.1016/j.neuron.2012.04.034

Kok, P., Rahnev, D., Jehee, J. F. M., Lau, H. C., & de Lange, F. P. (2012). Attention Reverses the Effect of Prediction in Silencing Sensory Signals. *Cerebral Cortex*, *22*(9), 2197–2206. https://doi.org/10.1093/cercor/bhr310

Konovalov, A., & Krajbich, I. (2018). Neurocomputational Dynamics of Sequence Learning. *Neuron*, *98*(6), 1282-1293.e4. https://doi.org/10.1016/j.neuron.2018.05.013

Manahova, M. E., Mostert, P., Kok, P., Schoffelen, J.-M., & de Lange, F. P. (2018). Stimulus Familiarity and Expectation Jointly Modulate Neural Activity in the Visual Ventral Stream. *Journal of Cognitive Neuroscience*, *30*(9), 1366–1377. https://doi.org/10.1162/jocn_a_01281

Meyer, T., & Olson, C. R. (2011). Statistical learning of visual transitions in monkey inferotemporal cortex. *Proceedings of the National Academy of Sciences*, *108*(48), 19401–19406. https://doi.org/10.1073/pnas.1112895108

Meyniel, F., & Dehaene, S. (2017). Brain networks for confidence weighting and hierarchical inference during probabilistic learning. *Proceedings of the National Academy of Sciences*, *114*(19). https://doi.org/10.1073/pnas.1615773114

Meyniel, F., Schlunegger, D., & Dehaene, S. (2015). The Sense of Confidence during Probabilistic Learning: A Normative Account. *PLOS Computational Biology*, *11*(6), e1004305. https://doi.org/10.1371/journal.pcbi.1004305

Nastase, S., Iacovella, V., & Hasson, U. (2014). Uncertainty in visual and auditory series is coded by modality-general and modality-specific neural systems: Uncertainty in Visual and Auditory Series. *Human Brain Mapping*, *35*(4), 1111–1128. https://doi.org/10.1002/hbm.22238

O'Reilly, J. X. (2013). Making predictions in a changing world—Inference, uncertainty, and learning. *Frontiers in Neuroscience*, *7*. https://doi.org/10.3389/fnins.2013.00105

Peirce, J., Gray, J. R., Simpson, S., MacAskill, M., Höchenberger, R., Sogo, H., Kastman, E., & Lindeløv, J. K. (2019). PsychoPy2: Experiments in behavior made easy. *Behavior Research Methods*, *51*(1), 195–203. https://doi.org/10.3758/s13428-018-01193-y

Pouget, A., Beck, J. M., Ma, W. J., & Latham, P. E. (2013). Probabilistic brains: Knowns and unknowns. *Nature Neuroscience*, *16*(9), 1170–1178. https://doi.org/10.1038/nn.3495

Pouget, A., Drugowitsch, J., & Kepecs, A. (2016). Confidence and certainty: Distinct probabilistic quantities for different goals. *Nature Neuroscience*, *19*(3), 366–374. https://doi.org/10.1038/nn.4240

Ramachandran, S., Meyer, T., & Olson, C. R. (2017). Prediction suppression and surprise enhancement in monkey inferotemporal cortex. *Journal of Neurophysiology*, *118*(1), 374–382. https://doi.org/10.1152/jn.00136.2017

Rao, R. P. N., & Ballard, D. H. (1999). Predictive coding in the visual cortex: A functional interpretation of some extra-classical receptive-field effects. *Nature Neuroscience*, *2*(1), 79–87. https://doi.org/10.1038/4580

Richter, D., & de Lange, F. P. (2019). Statistical learning attenuates visual activity only for attended stimuli. *eLife*, *8*, e47869. https://doi.org/10.7554/eLife.47869

Richter, D., Ekman, M., & de Lange, F. P. (2018). Suppressed Sensory Response to Predictable Object Stimuli throughout the Ventral Visual Stream. *The Journal of Neuroscience*, *38*(34), 7452–7461. https://doi.org/10.1523/JNEUROSCI.3421-17.2018

Sauter, M., Draschkow, D., & Mack, W. (2020). Building, Hosting and Recruiting: A Brief Introduction to Running Behavioral Experiments Online. *Brain Sciences*, *10*(4), 251. https://doi.org/10.3390/brainsci10040251

Schapiro, A. C., Rogers, T. T., Cordova, N. I., Turk-Browne, N. B., & Botvinick, M. M. (2013). Neural representations of events arise from temporal community structure. *Nature Neuroscience*, *16*(4), 486–492. https://doi.org/10.1038/nn.3331

Schapiro, A., & Turk-Browne, N. (2015). Statistical Learning. In *Brain Mapping* (S. 501–506). Elsevier. https://doi.org/10.1016/B978-0-12-397025-1.00276-1

Sherman, B. E., Graves, K. N., & Turk-Browne, N. B. (2020). The prevalence and importance of statistical learning in human cognition and behavior. *Current Opinion in Behavioral Sciences*, *32*, 15–20. https://doi.org/10.1016/j.cobeha.2020.01.015

Sherman, M. T., Seth, A. K., & Kanai, R. (2016). Predictions Shape Confidence in Right Inferior Frontal Gyrus. *The Journal of Neuroscience*, *36*(40), 10323–10336. https://doi.org/10.1523/JNEUROSCI.1092-16.2016

Turk-Browne, N. B., Scholl, B. J., Chun, M. M., & Johnson, M. K. (2009). Neural Evidence of Statistical Learning: Efficient Detection of Visual Regularities Without Awareness. *Journal of Cognitive Neuroscience*, *21*(10), 1934–1945. https://doi.org/10.1162/jocn.2009.21131

Turk-Browne, N. B., Scholl, B. J., Johnson, M. K., & Chun, M. M. (2010). Implicit Perceptual Anticipation Triggered by Statistical Learning. *Journal of Neuroscience*, *30*(33), 11177–11187. https://doi.org/10.1523/JNEUROSCI.0858-10.2010

Vallat, R. (2018). Pingouin: Statistics in Python. *Journal of Open Source Software*, *3*(31), 1026. https://doi.org/10.21105/joss.01026

Vanveen, V., & Carter, C. (2002). The anterior cingulate as a conflict monitor: fMRI and ERP studies. *Physiology & Behavior*, *77*(4–5), 477–482. https://doi.org/10.1016/S0031-9384(02)00930-7

Woo, C.-W., Krishnan, A., & Wager, T. D. (2014). Cluster-extent based thresholding in fMRI analyses: Pitfalls and recommendations. *NeuroImage*, *91*, 412–419. https://doi.org/10.1016/j.neuroimage.2013.12.058

Yon, D., & Frith, C. D. (2021). Precision and the Bayesian brain. Current Biology, 31(17), R1026–R1032. https://doi.org/10.1016/j.cub.2021.07.044

Supplementary material

MRI data processing

The following paragraphs include the automatically created methods section created by fMRIPrep. No changes have been made by the authors.

Results included in this manuscript come from preprocessing performed using *fMRIPrep* 20.2.4 (Esteban, Markiewicz, et al. (2018); Esteban, Blair, et al. (2018); RRID:SCR_016216), which is based on *Nipype* 1.6.1 (Gorgolewski et al. (2011); Gorgolewski et al. (2018); RRID:SCR_002502).

Anatomical data preprocessing

A total of 1 T1-weighted (T1w) images were found within the input BIDS dataset. The T1-weighted (T1w) image was corrected for intensity non-uniformity (INU) with N4BiasFieldCorrection (Tustison et al. 2010), distributed with ANTs 2.3.3 (Avants et al. 2008, RRID:SCR_004757), and used as T1w-reference throughout the workflow. The T1w-reference was then skull-stripped with a *Nipype* implementation of the antsBrainExtraction.sh workflow (from ANTs), using OASIS30ANTs as target template. Brain tissue segmentation of cerebrospinal fluid (CSF), white-matter (WM) and gray-matter (GM) was performed on the brain-extracted T1w using fast (FSL 5.0.9, RRID:SCR_002823, Zhang, Brady, and Smith 2001). Volume-based spatial normalization to one standard space (MNI152NLin2009cAsym) was performed through nonlinear registration with antsRegistration (ANTs 2.3.3), using brain-extracted versions of both T1w reference and the T1w template. The following template was selected for spatial normalization: *ICBM 152 Nonlinear Asymmetrical template version 2009c* [Fonov et al. (2009), RRID:SCR_008796; TemplateFlow ID: MNI152NLin2009cAsym],

Functional data preprocessing

For each of the 1 BOLD runs found per subject (across all tasks and sessions), the following preprocessing was performed. First, a reference volume and its skull-stripped version were generated using a custom methodology of *fMRIPrep*. Susceptibility distortion correction (SDC) was omitted. The BOLD reference was then co-registered to the T1w reference using flirt (FSL 5.0.9, Jenkinson and Smith 2001) with the boundary-based registration (Greve and Fischl 2009) cost-function. Co-registration was configured with nine degrees of freedom to account for distortions remaining in the BOLD reference. Head-motion parameters with respect to the BOLD reference (transformation matrices, and six corresponding rotation and translation parameters) are estimated before any spatiotemporal filtering using mcflirt (FSL 5.0.9, Jenkinson et al. 2002). The BOLD time-series (including slice-timing correction when applied) were resampled onto their original, native space by applying the transforms to correct for head-motion. These resampled BOLD time-series will be referred to as *preprocessed BOLD in original space*, or just *preprocessed BOLD*. The BOLD time-series were resampled into standard space, generating a *preprocessed BOLD run in MNI152NLin2009cAsym space*. First, a reference volume and its skull-stripped version were generated using a custom methodology of *fMRIPrep*. Several confounding time-series were calculated based on the *preprocessed BOLD*:

framewise displacement (FD), DVARS and three region-wise global signals. FD was computed using two formulations following Power (absolute sum of relative motions, Power et al. (2014)) and Jenkinson (relative root mean square displacement between affines, Jenkinson et al. (2002)). FD and DVARS are calculated for each functional run, both using their implementations in *Nipype* (following the definitions by Power et al. 2014). The three global signals are extracted within the CSF, the WM, and the whole-brain masks. Additionally, a set of physiological regressors were extracted to allow for component-based noise correction (*CompCor*, Behzadi et al. 2007). Principal components are estimated after high-pass filtering the *preprocessed BOLD* time-series (using a discrete cosine filter with 128s cut-off) for the two *CompCor* variants: temporal (tCompCor) and anatomical (aCompCor). tCompCor components are then calculated from the top 2% variable voxels within the brain mask. For aCompCor, three probabilistic masks (CSF, WM and combined CSF+WM) are generated in anatomical space. The implementation differs from that of Behzadi et al. in that instead of eroding the masks by 2 pixels on BOLD space, the aCompCor masks are subtracted a mask of pixels that likely contain a volume fraction of GM. This mask is obtained by thresholding the corresponding partial volume map at 0.05, and it ensures components are not extracted from voxels containing a minimal fraction of GM. Finally, these masks are resampled into BOLD space and binarized by thresholding at 0.99 (as in the original implementation). Components are also calculated separately within the WM and CSF masks. For each CompCor decomposition, the *k* components with the largest singular values are retained, such that the retained components' time series are sufficient to explain 50 percent of variance across the nuisance mask (CSF, WM, combined, or temporal). The remaining components are dropped from consideration. The head-motion estimates calculated in the correction step were also placed within the corresponding confounds file. The confound time series derived from head motion estimates and global signals were expanded with the inclusion of temporal derivatives and quadratic terms for each (Satterthwaite et al. 2013). Frames that exceeded a threshold of 0.5 mm FD or 1.5 standardised DVARS were annotated as motion outliers. All resamplings can be performed with *a single interpolation step* by composing all the pertinent transformations (i.e. head-motion transform matrices, susceptibility distortion correction when available, and co-registrations to anatomical and output spaces). Gridded (volumetric) resamplings were performed using antsApplyTransforms (ANTs), configured with Lanczos interpolation to minimize the smoothing effects of other kernels (Lanczos 1964). Non-gridded (surface) resamplings were performed using mri_vol2surf (FreeSurfer).

References

Abraham, Alexandre, Fabian Pedregosa, Michael Eickenberg, Philippe Gervais, Andreas Mueller, Jean Kossaifi, Alexandre Gramfort, Bertrand Thirion, and Gael Varoquaux. 2014. "Machine Learning for Neuroimaging with Scikit-Learn." *Frontiers in Neuroinformatics* 8. https://doi.org/10.3389/fninf.2014.00014.

Avants, B.B., C.L. Epstein, M. Grossman, and J.C. Gee. 2008. "Symmetric Diffeomorphic Image Registration with Cross-Correlation: Evaluating Automated Labeling of Elderly and Neurodegenerative Brain." *Medical Image Analysis* 12 (1): 26–41. https://doi.org/10.1016/j.media.2007.06.004.

Behzadi, Yashar, Khaled Restom, Joy Liau, and Thomas T. Liu. 2007. “A Component Based Noise Correction Method (CompCor) for BOLD and Perfusion Based fMRI.” *NeuroImage* 37 (1): 90–101. https://doi.org/10.1016/j.neuroimage.2007.04.042.

Esteban, Oscar, Ross Blair, Christopher J. Markiewicz, Shoshana L. Berleant, Craig Moodie, Feilong Ma, Ayse Ilkay Isik, et al. 2018. “FMRIPrep.” *Software*. Zenodo. https://doi.org/10.5281/zenodo.852659.

Esteban, Oscar, Christopher Markiewicz, Ross W Blair, Craig Moodie, Ayse Ilkay Isik, Asier Erramuzpe Aliaga, James Kent, et al. 2018. “fMRIPrep: A Robust Preprocessing Pipeline for Functional MRI.” *Nature Methods*. https://doi.org/10.1038/s41592-018-0235-4.

Fonov, VS, AC Evans, RC McKinstry, CR Almli, and DL Collins. 2009. “Unbiased Nonlinear Average Age-Appropriate Brain Templates from Birth to Adulthood.” *NeuroImage* 47, Supplement 1: S102. https://doi.org/10.1016/S1053-8119(09)70884-5.

Gorgolewski, K., C. D. Burns, C. Madison, D. Clark, Y. O. Halchenko, M. L. Waskom, and S. Ghosh. 2011. “Nipype: A Flexible, Lightweight and Extensible Neuroimaging Data Processing Framework in Python.” *Frontiers in Neuroinformatics* 5: 13. https://doi.org/10.3389/fninf.2011.00013.

Gorgolewski, Krzysztof J., Oscar Esteban, Christopher J. Markiewicz, Erik Ziegler, David Gage Ellis, Michael Philipp Notter, Dorota Jarecka, et al. 2018. “Nipype.” *Software*. Zenodo. https://doi.org/10.5281/zenodo.596855.

Greve, Douglas N, and Bruce Fischl. 2009. “Accurate and Robust Brain Image Alignment Using Boundary-Based Registration.” *NeuroImage* 48 (1): 63–72. https://doi.org/10.1016/j.neuroimage.2009.06.060.

Jenkinson, Mark, Peter Bannister, Michael Brady, and Stephen Smith. 2002. “Improved Optimization for the Robust and Accurate Linear Registration and Motion Correction of Brain Images.” *NeuroImage* 17 (2): 825–41. https://doi.org/10.1006/nimg.2002.1132.

Jenkinson, Mark, and Stephen Smith. 2001. “A Global Optimisation Method for Robust Affine Registration of Brain Images.” *Medical Image Analysis* 5 (2): 143–56. https://doi.org/10.1016/S1361-8415(01)00036-6.

Lanczos, C. 1964. “Evaluation of Noisy Data.” *Journal of the Society for Industrial and Applied Mathematics Series B Numerical Analysis* 1 (1): 76–85. https://doi.org/10.1137/0701007.

Power, Jonathan D., Anish Mitra, Timothy O. Laumann, Abraham Z. Snyder, Bradley L. Schlaggar, and Steven E. Petersen. 2014. “Methods to Detect, Characterize, and Remove Motion Artifact in Resting State fMRI.” *NeuroImage* 84 (Supplement C): 320–41. https://doi.org/10.1016/j.neuroimage.2013.08.048.

Satterthwaite, Theodore D., Mark A. Elliott, Raphael T. Gerraty, Kosha Ruparel, James Loughead, Monica E. Calkins, Simon B. Eickhoff, et al. 2013. “An improved framework for confound regression and filtering for control of motion artifact in the preprocessing of resting-state functional connectivity data.” *NeuroImage* 64 (1): 240–56. https://doi.org/10.1016/j.neuroimage.2012.08.052.

Tustison, N. J., B. B. Avants, P. A. Cook, Y. Zheng, A. Egan, P. A. Yushkevich, and J. C. Gee. 2010. “N4ITK: Improved N3 Bias Correction.” *IEEE Transactions on Medical Imaging* 29 (6): 1310–20. https://doi.org/10.1109/TMI.2010.2046908.

Zhang, Y., M. Brady, and S. Smith. 2001. “Segmentation of Brain MR Images Through a Hidden Markov Random Field Model and the Expectation-Maximization Algorithm.” *IEEE Transactions on Medical Imaging* 20 (1): 45–57. https://doi.org/10.1109/42.906424.

3 Manuscript II: The energy metabolic footprint of predictive processing in the human brain

André Hechler[1,2], Floris P. de Lange[3], Valentin Riedl[1,2]

[1] Department of Neuroradiology, Neuroimaging Center, Technical University of Munich, Munich, Germany
[2] Graduate School of Systemic Neurosciences, Ludwig-Maximilians-University, Munich, Germany
[3] Donders Institute for Brain, Cognition and Behaviour, Radboud University Nijmegen, Nijmegen, Netherlands

Abstract

Neural activity is a highly energy-intensive process. In the human brain, signaling consumes up to 75% of the available energy resources with postsynaptic potentials as the largest factor. Visual processing is especially costly, with increases in energy consumption of up to 20% in the visual cortex. In recent years, vision has been cast as a constructive process, harnessing prior knowledge in a constant feedback loop of top-down prediction and bottom-up sensory input. Interestingly, sensory input that is in line with our predictions might be processed at lower energy metabolic cost. However, there is no evidence for this claim yet, possibly due to the scarcity of measures that quantify energy consumption in the human brain.

Here, we used a novel MR method measuring the cerebral metabolic rate of oxygen during sensory stimulation of visual sequences that varied in their predictability. Since predictive processing is driven by estimates of uncertainty, we assessed how confident subjects were in their knowledge of the underlying patterns. We found that processing predictable sequences steeply decreased in energetic cost with increasing confidence. Strikingly, these energetic effects were not limited to visual areas, summing up to a cortical difference of 13% between high and low levels of confidence. Furthermore, sequences deviating from expectations were energetically cheaper than predictable ones for low confidence levels, but costlier for high levels. These results speak for a major role of predictive processing in balancing the brain's energy budget and emphasize the impact of interindividual differences when learning predictive patterns.

Introduction

To produce the energy that fuels neural activity, the brain needs a steady supply of oxygen and glucose. Neural signaling is estimated to account for 75% of the brain's energy consumption in grey matter with 50% of that fraction going towards postsynaptic potentials (Howarth et al., 2012). The rate of energy consumption is strongly affected by sensory stimulation: In the visual cortex, up to 20% more energy is used during stimulation (Lin et al., 2010). Such effects cannot be reliably measured by conventional BOLD imaging, which is strongly driven by hemodynamic effects (Drew, 2019). While blood flow and energy consumption are tightly coupled at rest, the ratio differs across the cortex (Devonshire et al., 2012; Drew, 2022; Hyder, 2010) and is impacted by attentional state (Moradi et al., 2012) as well as stimulation duration (Moradi & Buxton, 2013). Addressing this problem, multiparametric quantitative BOLD imaging (mqBOLD; Christen et al., 2012; Hirsch et al., 2014; Kaczmarz et al., 2020) provides direct access to energy metabolic processes by measuring the cerebral metabolic rate of oxygen on a voxel level (CMR_{O2}). CMR_{O2} has three key advantages for research on energy consumption: First, it represents the main resource of ATP production in the brain (Dienel, 2014; Harris et al., 2012) and is biologically interpretable. Second, it combines separate measurements of hemodynamics and blood (de-)oxygenation. This accounts for differences in neurovascular coupling, allowing the comparison and integration of CMR_{O2} across the brain. Third, due to the quantitative approach, CMR_{O2} can be analyzed in absolute values, enabling inference on conditions instead of contrasts while staying comparable between subjects.

Previous studies on the energetic cost of visual perception purely focused on stimulus characteristics, presenting simple or complex stimuli and limiting the analysis to the visual cortex (Griffeth et al., 2015; Lin et al., 2010). However, perception can be understood as a constructive process that heavily relies on prior knowledge from higher cognitive areas to interpret sensory input (Clark, 2013; de Lange et al., 2018; Teufel & Fletcher, 2020). The predictive coding framework posits that the brain is constantly optimizing a model of the world with the goal of minimizing the surprise of sensory observations (Rao & Ballard, 1999). Deviations from the model lead to error signals and subsequent model updating which are implemented in the brain as hierarchical feedback loops (Bastos et al., 2012; Walsh et al., 2020). A cornerstone of predictive processing is the extraction of spatial and temporal patterns to infer causes from incomplete observations and predict future events from past and current

ones. The latter ability is exemplified by studies investigating statistical learning, where a sensory stream is generated from specific transitional probabilities. These provide a source of prediction after repeated exposure (Fiser & Lengyel, 2022; Turk-Browne et al., 2009b, 2010). Studies using such a design revealed lower activation for expected compared to unexpected stimuli across conventional imaging modalities (Manahova et al., 2018; Richter et al., 2018a; Stefanics et al., 2011). This is expected when understanding predictive coding as a theory of efficient coding (Chalk et al., 2018; Quintela-López et al., 2022): By learning the patterns behind sensory input, our predictions can be optimized to the point where little to no error signals or changes to the internal model are necessary (for a mathematical derivation see Sengupta et al., 2013). However, there is no empirical evidence available that this results in reduced energy consumption.

In the present study, we extended an established visual statistical learning design with a multi-day training phase, maximizing consolidation of the underlying patterns. Additionally, we accounted for individual differences in the learning process by assessing participants' confidence in knowing the underlying patterns (Geurts et al., 2022; Meyniel, Sigman, et al., 2015; Sanders et al., 2016). Predictive processing theories heavily draw from Bayesian inference models where uncertainty about predictions has a major effect on perception (Yon & Frith, 2021). While statistical learning is based on continuous, partly implicit processes (but see Dale et al., 2012; Vadillo et al., 2016), a line of studies has shown that intermittent confidence ratings separate from the learning process reflect the statistical confidence of a Bayesian observer during the learning blocks (Bounmy et al., 2023; Meyniel, Schlunegger & Dehaene, 2015; Meyniel & Dehaene, 2017).

To summarize, we here used a MR-derived measure of oxygen consumption to study how predictive characteristics of external input and internal beliefs impact the brain's energy balance. We hypothesized that predictable stimuli are cheaper to process than unexpected stimuli and that this difference scales with subjective confidence in the predictable patterns. To preview our findings, the energetic cost of predictable input indeed decreased with confidence in both sensory and higher cognitive regions. Intriguingly, a general energetic advantage only emerged at high confidence levels.

Results

Confidence ratings and processing speed increase with repeated exposure to visual patterns

Prior to the scanning session, participants completed a three-day online training phase. They were presented with image sequences of everyday objects that either followed a deterministic pattern (predictable condition, P) or a random pattern (unpredictable condition, U) (Figures 1a and b). To ensure vigilance and gaze fixation, participants had to react to upside-down objects, which occurred irrespective of sequence predictability. Average accuracy exceeded 95% and neither accuracy nor reaction time differed between conditions (reaction time: t(*40*)=-1.4, p=.17; accuracy: t(*40*)=1.51, p=.14). Separate from the statistical learning blocks, we assessed the trajectory of learning both on an objective and a subjective scale. To this end, we presented participants with incomplete sequences and prompted them to choose the correct follow-up object. For each choice, we assessed confidence ratings as an indicator of uncertainty regarding knowledge of the underlying patterns (Figure 1c). Over the training days, confidence ratings increased for predictable visual sequences while they remained constant for unpredictable input (Figure 1d, left). We also saw an objective improvement in pattern completion (Figure 1d, right) which was highly correlated with confidence (r=0.82, p<0.001). Since we did not provide any feedback, this suggests that participants were highly accurate in assessing their performance.

Directly prior to scanning, we included an additional task that aimed at detecting improvements in an implicit marker of learning. Similar to the training phase, we presented predictable and unpredictable sequences. Instead of performing the cover task, participants were instructed to react to the occurrence of a target object that was shown before each sequence. The target was randomly chosen from the upcoming objects, meaning it could be anticipated in the predictable condition. We found that participants detected predictable objects over 10% faster (median RT change=-10.75%, t=-5.7, p<0.001). Furthermore, this difference scaled with confidence ratings, meaning that confident subjects gained a stronger increase in processing speed (Figure 1e).

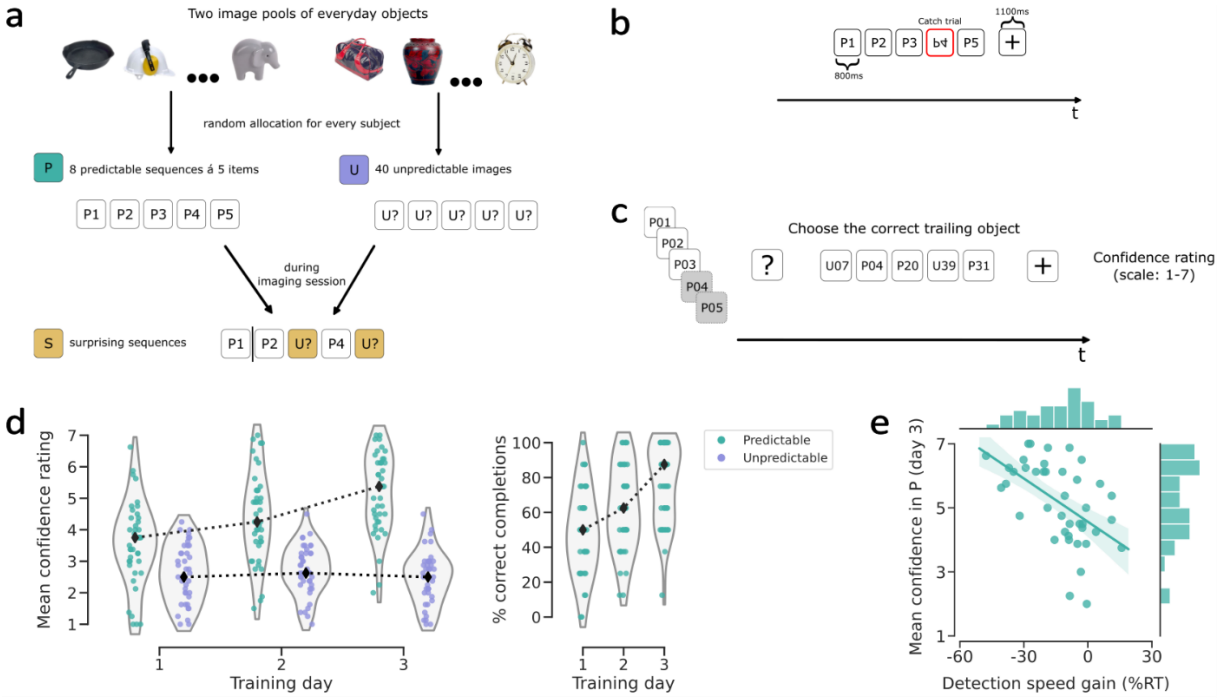

Figure 1. **a**. Experimental stimuli were sequences of five everyday objects. Predictable sequences (P) had consistent order and composition, unpredictable ones (U) were random. During MR acquisition, surprising sequences (S) composed of P and U were additionally shown. See Methods section for details. **b**. Stimulus presentation. Objects were shown sequentially and have no interstimulus interval. With a 10% chance, an image was horizontally flipped, and participants had to react with a button press. **c**. Sequence completion test after each training day. One to four images of a sequence were shown, and the correct trailing object had to be chosen from five options. This was followed by a confidence rating prompt. **d**. *Left*: Subject-wise average confidence ratings following respective training days. Black lines and markers indicate day-wise sample medians. *Right*: Percentage of correctly chosen trailing images in the completion test. Averaged completion percentage is in steps of 12.5% because eight sequences were tested per day. **e**. Results of separate object detection task. Detecting a predictable compared to an unpredictable target was significantly quicker for participants with high confidence (r=-0.54, p<0.001). Negative values indicate faster reactions for predictable objects.

The interaction of predictability and confidence drives energy consumption

For our main analysis, we first created voxel-wise brain maps of CMR_{O2} from our MR and blood sampling data (Figure 2a and Methods). We then averaged condition-wise CMR_{O2} values for every subject within 400 functional regions as defined by the Schaefer parcellation (Schaefer et al., 2018). This was done to control for voxel-level noise and to differentiate sensory from higher cognitive regions. Prior to analyzing the imaging data, we tested for potential confounding effects of the cover task. Unlike during the training phase, we found differences between conditions (repeated measures ANOVA: $F(2,80)=4.44$, $p=0.015$). Post-hoc tests showed that participants reacted most quickly to predictable items (Figure 2b). This is surprising, given the independence of the cover task from sequence predictability and the absence of reaction time differences in the training phase. Consequently, we controlled for these effects in the following analyses.

In contrast to conventional GLM analyses, our aim was not to localize regions of maximum effect, but to assess net energetic changes on a larger scale. As argued previously, evoked activity extends far beyond local peaks which can limit results to the "tip of the iceberg" (Noble et al., 2022, 2023). To address this, we analyzed metabolic effects on the level of functional networks. We used an established cortical atlas which defines regions based on functional coactivation (Schaefer et al., 2018). These regions form extended networks that differentiate sensory from higher cognitive areas. For our analysis, we used six networks covering most of the cortex, comprising 374 regions. The regional data was analyzed using a single linear mixed model with predictors of condition, confidence, functional network, and reaction time in the cover task (Methods). Essentially, this model allowed us to explain regional CMR_{O2} as a linear combination of predictability and confidence, with separate regressors for six functional networks while accounting for confounding effects of reaction time. The random term of the model addresses subject-specific baselines of CMR_{O2} on which the fixed terms are added. Confidence values correspond to the condition-wise average after the last training day, and we used ratings from the predictable sequences for the surprising condition. All resulting model estimates are relative to the reference categories of the respective predictors: We used the visual network and the predictable condition. Importantly, since CMR_{O2} values are quantitative, the model parameters for the reference levels can be interpreted without referring to a contrast estimate. Parameter estimates for reference levels can thus be interpreted for the visual network in the predictable condition. This extends to all predictor levels (i.e., conditions and networks) for which no significant interaction is present.

We first confirmed that our model significantly improves upon a null model that lacks the experimental variable (BF_{log}=131.51; $X^2(42)$=552.6, p<0.001). Table 1 includes an overview of all significant predictors and the full model results are provided in the supplement. The intercept of 136.33 represents the baseline CMR_{O2} across subjects, in units of µmol O_2/min/100g. Given our reference levels, this refers to the visual network during predictable input, for a hypothetical subject with average confidence and reaction time. Validating our approach, this is well within the range of 120-160 usually reported for healthy human subjects (summarized in Xu et al., 2009, see also Christen et al., 2012; Göttler et al., 2019). With respect to our main research question, we found no main effect of condition, but a significant interaction between condition and confidence. This means that, while energy consumption did not differ between levels of predictability per se, the conditions were differentially affected

by confidence levels. Overall, CMR_{O2} decreases with confidence (denoted by the negative estimate for the reference category) but this effect is offset in the surprising condition. The linear combination of these parameters is visualized in Figure 2e (for uncertainty of these estimates refer to Table 1). As the reference category corresponds to the predictable condition, our results show a steep decrease of energetic cost with confidence for predictable sequences. Interestingly, this effect was not significantly different for unpredictable sequences, suggesting that confidence impacted energetic cost even in the face of objectively random stimulation. However, the corresponding parameter estimate showed a quantitative trend towards a weaker effect of confidence.

High confidence reduces energy consumption across the whole cortex

Interestingly, while the functional networks differed in their CMR_{O2} baseline, the interaction between predictability and confidence was largely consistent across the cortex. The control network (also referred to as the fronto-parietal network) is a notable exception: Here, the decrease of CMR_{O2} with confidence was significantly weaker than in the other networks. Regarding unpredictable sequences, the control network showed another unique deviation. Irrespective of confidence levels, its energy consumption was significantly higher than in the other conditions.

Given that the interaction of condition and confidence was found across networks, we evaluated the net effect on the whole cortex. To achieve this, we used the significant model parameters of confidence and network to calculate the predicted CMR_{O2} for low (5^{th} percentile), average (50^{th} percentile) and high (95^{th} percentile) confidence. Since CMR_{O2} is a rate of consumption relative to a tissue mass of 100g, we scaled the predicted energy consumption to the gray-matter mass of the respective networks (Methods). Strikingly, our data suggest that high confidence reduces cortical energy consumption by 13.25% (CI: [0.34; 26.16]) relative to low levels of confidence (Figures 2d).

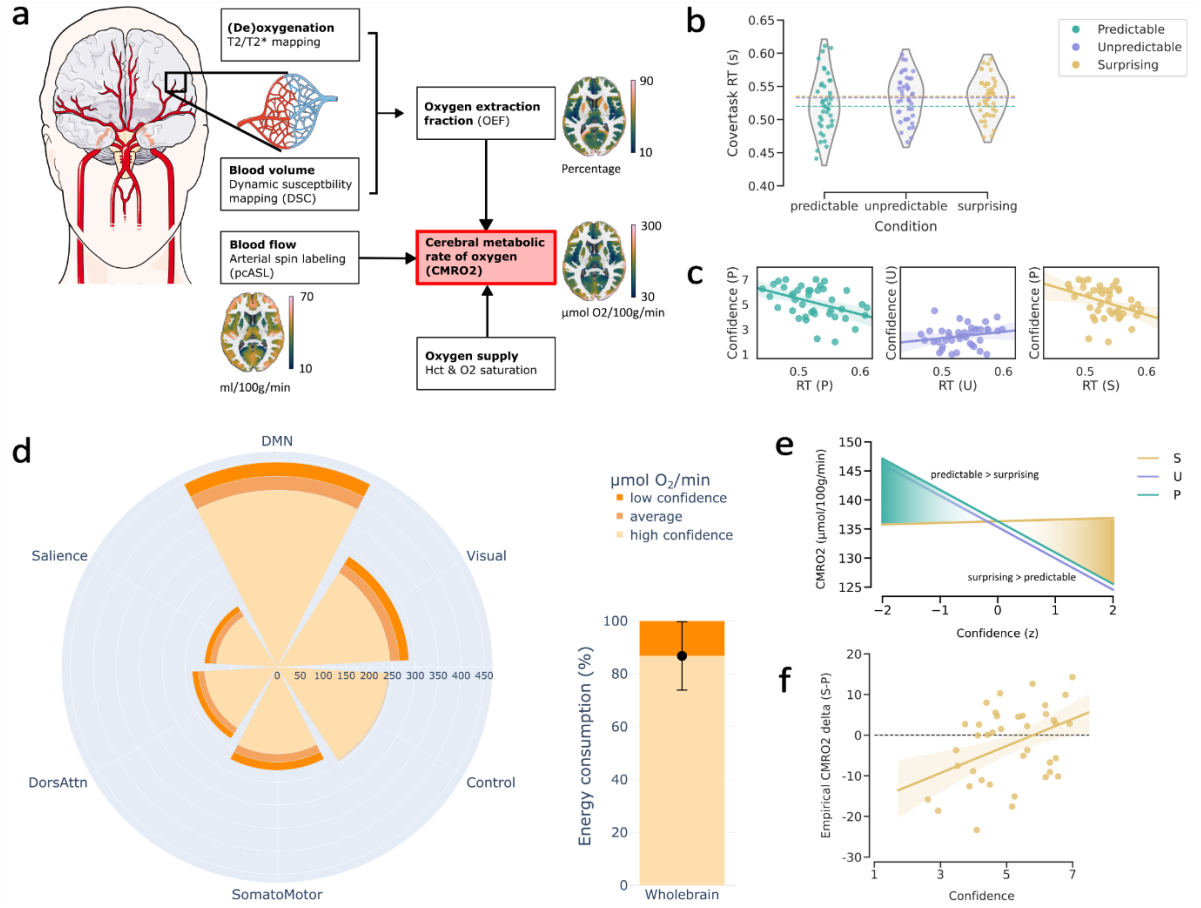

Figure 2. **a**. Flowchart depicting the workflow of CMR_{O2} calculation from experimental data*. Parameter maps show group averages in cerebral grey matter. Details are described in the method section. **b**. In the cover task, participants reacted significantly quicker to predictable than surprising objects (t(*40*)=-2.9, p_{FWE}=0.017) with a trend in the same direction compared to unpredictable objects (t(*40*)=-2.24, p_{FWE}=0.091). Points show subject-wise averages across the scanning session and dashed lines indicate condition-wise medians. **c**. Condition-wise correlations between confidence ratings and reaction time in the cover task during scanning. We found a significant negative association for the predictable condition (r=-0.42, p_{FWE}=0.018). The effect was similar for the surprising condition but did not survive FWE correction (r=-0.36, p_{FWE}=0.066). **d**. *Left*: Oxygen consumption per minute (scaled to network gray-matter mass) as predicted by the significant model terms of network and confidence. Low, average and high confidence correspond to the 5th, 50th and 95th percentile of z-scaled confidence ratings. *Right*: Predicted relative energy consumption for subjects with high confidence compared to low confidence, summed over all networks. The black error bar indicates the 95th confidence interval of estimated energy consumption in high confidence subjects based on the confidence interval of the model parameter (see Methods). **e**. Predicted CMR_{O2} across confidence levels for the reference network (Visual), based on our model slopes (Table 1). Shaded areas indicate whether predictable sequences (green) or surprising sequences (yellow) are more costly. Note that the slopes for predictable and unpredictable sequences entirely overlap according to our model. The regression lines are slightly offset for visualization purposes. **f**. Cortical CMR_{O2} difference between surprising and predictable sequences is explained by confidence (linear regression controlling for reaction time: $R^2_{partial}$=.18, t(*40*)=2.92, p=0.005). Data shown are subject-wise cortical CMR_{O2} differences between surprising and predictable sequences in the experimental data. **Visualizations of head and vessels were created using Servier Medical Art, provided by Servier, licensed under a Creative Commons Attribution 3.0 unported license*

Predictor level	Predictor (only significant[§])	Estimate (in µmol O2/min/100g)	95% CI	p-value
Reference effect*	Intercept	136.33	128.84 – 143.83	<0.001
	confidence	-5.39	-9.88 – -0.9	0.019
	reaction time	-4.84	-7.73 – -1.95	0.001
	network[control]	10.91	7.28 – 14.54	<0.001
	network[DMN]	8.26	5.08 – 11.43	<0.001
	network[SomMot]	-11.81	-15.08 – -8.54	<0.001
	network[salience]	-17.67	-21.37 – -13.96	<0.001
Interaction	condition[S]*confidence	5.67	-2.25 – 12.45	0.026
	condition[U]*confidence	5.10	-0.27 – 12.94	0.174[§]
	condition[U]*reaction time	3.12	0.16 – 6.08	0.039
	condition[U]*network[control]	6.51	0.8 – 12.21	0.025
	confidence*network[control]	5.03	1.7 – 8.35	0.003

Table 1. Significant predictors estimated by a robust linear mixed model with the following formula (following R conventions): condition*confidence*network + condition:reaction_time + (1|subject/condition). Estimates reflect absolute CMR_{O2} values and can be interpreted as a predicted change in energy consumption for a given change in (standardized) predictors. Non-significant predictors are omitted for brevity. We provide the full model output in the supplement. *Parameters are based on these reference categories: condition[predictable] and network[visual] [§]Additional inclusions due to importance for experimental interpretation

Relative cost of surprising input is a function of decreasing cost of predictable input

As non-quantitative imaging methods are limited to contrast-based analyses, prediction error activity is often formalized as a relative increase in activity for surprising over predictable input. Since our quantitative data resolves the cost of individual conditions, it informs us about the origins of relative effects. Looking at the condition-wise slopes (Figure 2e), our model suggested that relative prediction error activity emerges as a function of the decreasing cost of predictable input. We tested this assumption on our empirical data, regressing the difference between CMR_{O2} in the surprising and predictable condition against confidence. As expected, prediction error activity emerged at high levels of confidence (Figure

2f) across the cortex. Interestingly, the opposite was true for low levels of confidence and both conditions had similar energetic cost at slightly above average confidence.

Cover task reaction times play in major role for energy consumption

According to our model, reaction times in the cover task had a general effect on cortical energy consumption, with a magnitude similar to the effect of confidence. The slower subjects reacted to the upside-down stimuli, the lower their energy consumption was across conditions (Table 1). However, this effect was significantly weaker in the unpredictable condition. Since confidence was central to the effect of predictability on energy consumption, we wondered if it also explains differences in reaction time. We therefore correlated confidence ratings and reaction times for each condition separately.

Interestingly, reaction times decreased with confidence in the predictable and the surprising condition, although the latter association did not survive correction for multiple comparisons (Figure 2c). Although tentative, this pattern has an intriguing implication for the energy metabolic effects: While more confident subjects used less energy for predictable input, they also reacted quicker in the cover task, which in turn increased energy consumption. This led us to ask the question whether these effects even out according to our model. We recalculated the net cortical energy consumption with opposing effects of reaction time (methods). While the trend of a cortical CMR_{O2} decrease persisted (8.27%, CI: [-5.46; 22.03]), we found no significant evidence when factoring in opposing effects of reaction time.

Discussion

In the present study, we quantified the metabolic cost of processing visual sequences under different levels of objective and subjective predictability. According to the predictive coding framework, humans build a predictive model of the world that is updated when deviations are encountered (Rao & Ballard, 1999). Expected visual stimuli lead to weaker error signaling and model updating across processing hierarchies of the cortex (Bastos et al., 2012; Walsh et al., 2020). It has been argued that these predictive models balance accuracy and complexity (Friston, 2010). This might promote the minimization of energy usage in the brain (Sengupta et al., 2013). Due to the limited availability of in-vivo metabolic imaging methods, this assumption has never been tested.

We show that cortical energy consumption is driven by the interaction between the objective predictability of visual input and subjective confidence in knowing the patterns. When presented with highly predictable input, energy consumption decreased with increasing

confidence, up to a cortical difference of 13 percent. As a consequence, perceiving surprising input became more costly in comparison. These effects were remarkably consistent across the cortex, with only the control network showing slight deviations from the system-level effect. High confidence was also linked to quicker detection of predictable stimuli. In summary, we found that predictable patterns promote behavioral improvements and concurrent energy metabolic reductions.

Energy metabolism and efficiency

When formulating biologically realistic models of brain function, accounting for resource constraints is crucial (Roberts et al., 2014). Healthy brain function relies on constant oxygen and glucose delivery and shortages can have severe consequences ranging from cognitive deficits to cell death (Lee et al., 2020; Warren & Frier, 2005). Interestingly, the brain seems to be optimized for efficiency, from the firing patterns of neurons to the architecture of functional networks (Yu & Yu, 2017; Zhou et al., 2022). In related cognitive research, it has been argued that learning is an efficient process that maximizes performance while minimizing cost (Commins, 2018). These lines of research can be unified from an information-theoretic perspective: Akin to training algorithms in machine learning, the brain might aim to optimize behavioral accuracy while minimizing the complexity of internal representations (Zénon et al., 2019). This notion has been generalized under the Free Energy Principle (Friston, 2010): Humans build an internal model of the world that is continuously updated with new information, balancing model accuracy with model complexity. Assuming that the brain represents this model, the efficiency of the model might extend to the energetic efficiency of the underlying neural activity (Sengupta et al., 2013). Computational studies using neural networks are in support of this hypothesis: Networks trained to minimize their activity in a sequence prediction task developed predictive architectures (Ali et al., 2022) and predictive learning algorithms were expressed as energy-minimization algorithms (Luczak et al., 2022). Despite these theoretical and computational underpinnings, no direct evidence is available that predictive processing leads to metabolic efficiency.

The role of subjective confidence during statistical learning

The ability to extract transitional probabilities from sequential stimulation is often termed statistical learning and has a rich tradition in cognitive research (Schapiro & Turk-Browne, 2015). Previous work focused on the assumed automaticity of the learning process, which can happen in the absence of intention or awareness (Alamia et al., 2016; Turk-Browne et al., 2005). However, recent work showed that, while implicit processes are a prerequisite,

explicit knowledge can be acquired in parallel (Batterink et al., 2015; Dale et al., 2012). Blurring the lines further, Conway (2020) argued that most designs studying the learning process of transitional probabilities tap into the same underlying process. However, there is evidence for a difference in the neuroanatomical substrates of implicit and explicit statistical learning, so comparisons between studies should be drawn with care (Aizenstein, 2004). Recently, it has been suggested that a general Bayesian inference process underscores all probabilistic computation (Fiser & Lengyel, 2022). Consequently, instead of addressing a specific paradigm, the current work used streams of visual sequences as a tool to elicit probabilistic learning.

By including confidence ratings, we addressed a major parameter of Bayesian inference: Uncertainty (often referred to by its inverse, *precision*). Both perception and decisions are subject to uncertainty, stemming from both external sources (e.g. the visibility of a stimulus) and internal sources (e.g. noise in neural transmission) (Bach & Dolan, 2012). These sources have downstream effects on the identification of rules in our environment or the prediction of likely outcomes of a decision. Furthermore, while Bayesian inference is a powerful model of human cognition, individuals often deviate from the idealized performance (Acerbi et al., 2014; Beck et al., 2012). These lines of research suggest that considerable interindividual differences are to be expected when submitting a group of participants to the same task. To account for this variance, we used confidence ratings as a proxy for individual uncertainty.

The exact interpretation and implementation of confidence ratings is still debated. Previously, it has been defined as the posterior probability that a choice is correct, given the evidence (Pouget et al., 2016). For a statistically precise estimate, a neural representation of probability distributions is needed, an assumption that has been called into question (Koblinger et al., 2021). A promising explanation are probabilistic population codes (Knill & Pouget, 2004). Under this model, neurons are sensitive to different expressions of internal or external variables, adjusting their firing pattern or frequency according to the similarity of the stimulus to their preferred values. A population of neurons can then serve as a probability distribution over possible values of the variable (for a review of the evidence see Ma & Jazayeri, 2014). In this context, confidence ratings provide a summary statistic over this distribution, indicating the spread of the probability distribution around the central tendency (Meyniel, Sigman & Mainen, 2015). However, multiple alternative explanations have been developed, ranging from approximations to non-Bayesian accounts of confidence (Adler & Ma,

2018). Nevertheless, most formulations agree that some form of confidence-weighting is central to the human inference process.

The effect of subjective confidence on energetic cost

Our results confirm that confidence determines energy consumption during varying levels of objective predictability. The link between confidence and brain activation during probabilistic learning has been investigated in a previous line of studies using conventional BOLD (Bounmy et al., 2023; Meyniel, Schlunegger & Dehaene, 2015; Meyniel & Dehaene, 2017). Three caveats apply when comparing these studies to our results. First, we presented long blocks with stable transitional probabilities, while the underlying parameters were volatile in their design. Second, local BOLD effects do not necessarily correspond to energy consumption (Drew, 2019), especially when extending the analysis above local peaks as in our case. Lastly, our data represents an average of multiple minutes of visual input. Transient brain responses as revealed by BOLD analyses may not correspond to the net efficiency of a given neural process. Consequently, our data is reflective of the aggregate cost of continuous precision-weighted updating.

The aforementioned studies found that BOLD activity decreased with confidence while it increased with surprise in various areas. They also report results for regions most sensitive to precision-weighted updating: Here, surprising input elicited higher signal across confidence levels. In contrast, we found that predictable input leads to higher energy consumption than surprising input for low confidence, while the inverse was true for high confidence. Furthermore, the cost of surprising input did not change with confidence in our data. A possible explanation is the length of the presentation blocks in our study: While deviations from the learned sequences were initially surprising, participants might have learned the increased variance of the presentation. This would make the surprising sequences more akin to sequences of intermediate (probabilistic) predictability. In line with this argument, humans have been shown to track changes in environmental statistics, discounting previous observations (Beierholm et al., 2020; Maheu et al., 2022). In this case, confidence ratings from the previously predictable condition would have no lasting effect on the ”surprising” condition.

Nevertheless, we still reproduced the classical pattern of higher cost for surprising compared to predictable input. In our data, this effect only emerged at high confidence levels due to the decreasing cost of predictable input. This addresses the arbitration between two potential processes behind prediction errors: Surprise enhancement assumes that surprising stimuli evoke stronger signals while expectation suppression assumes decreased signals for

predictable input (Feuerriegel et al., 2021; Manahova et al., 2018). From the perspective of energy consumption, our data is suggestive of expectation suppression as the source of relative prediction errors. This contrasts with results from BOLD and EEG studies, where evidence for surprise enhancement is stronger (reviewed in Feuerriegel, Vogels & Kovács (2021)).

Interestingly, the effect of confidence on cortical cost was not significantly different between predictable and objectively random input. Since we gave no feedback during the sequence prediction task, it is possible that participants thought they detected a pattern in the random sequences. This is supported by the low, but far from minimal confidence ratings for unpredictable sequences during training. Humans tend to perceive structure even in random sequences (Huettel et al., 2002) which might drive perceptual inference in the absence of feedback. Future studies are needed to examine to which extent the effect of precise priors is independent of their objective accuracy. We also found that reaction times in the cover task increase with confidence during predictable and surprising sequences. This indicates that highly confident subjects improved their performance in a concurrent, but unrelated task. A tentative interpretation is that participants shifted their resources from model updating to lower-level perception of the objects' orientations. Under this assumption, the brain does not aim to decrease energy consumption per se, but rather reassigns resources dynamically under the constraints of energy availability (Christie & Schrater, 2015).

Energetic changes on a cortical level

In summary, our results show that an experimental manipulation as seemingly small as a change to transitional probabilities can lead to large differences in energy consumption. Surprisingly, these effects were highly consistent over the cortex, affecting both sensory and higher cognitive networks. Importantly, our study does not speak to the functional specificity of these energetic changes. Rather, we provide evidence for the net effect of perceptual inference on the cortex. The widespread changes in energy metabolism are in line with recent accounts of brain function as a highly integrated system (Pessoa, 2023). The focus on localization in neuroimaging research might have been supported by traditional mass-univariate analyses, which restrict the view of the brain to the "tip of the iceberg" (Noble et al., 2022, 2023). Specific to predictive processing, a recent meta-studied found a brain-spanning prediction network that encompassed functionally connected regions across the brain (Ficco et al., 2021). Furthermore, a recent preprint reported representations of prior information across all levels of processing in the mouse cortex (Findling et al., 2023). In line

with these works, we propose that the efficiency of a highly integrative system is best evaluated over a correspondingly large spatial extent.

Oxidative versus glycolytically driven energy metabolism

Lastly, a recent paper developed an integrative framework drawing a direct link between predictive processing and differential patterns of BOLD, CMR_{O2} and CMR_{Glc} (the cerebral metabolic rate of glucose) (Theriault et al., 2023). The authors argue that the ratio between ATP-yielding metabolites differs between bottom-up prediction errors and top-down predictions: Prediction errors rely on fast and flexible ATP generation via non-oxidative glycolysis, while prediction is based on the more efficient but less flexible oxidative phosphorylation. Importantly, the BOLD signal is driven by blood flow, which indicates oxygen availability but not necessarily oxygen consumption (Fox et al., 1988). As blood flow is more closely related to glucose use than oxygen use, BOLD is a better reflection of CMR_{Glc} than CMR_{O2} (Raichle & Mintun, 2006). It follows that, on the one hand, our CMR_{O2} data might be less sensitive to changes in error signaling but on the other hand, it could unveil prediction-related changes that BOLD does not capture. Future studies could dig deeper into these assumptions with the aim of providing a complete picture of the energy metabolism of the brain during predictive processing.

Methods

Participants

We recruited 44 participants including students, doctoral researchers, clinic staff and the general population of Munich. The sample size was based on a-priori analysis, aiming to detect at least a medium effect size (d>=0.5, alpha=0.05, beta=0.9). All participants took part in a familiarization MRI session (20 minutes), an online training phase over three days and the main MRI session (70 minutes) on the day after training completion. Two participants were excluded due to technical problems during data acquisition. One further subject was excluded from analysis because their data indicated that the training phase was not properly performed: The same button was pressed for every confidence rating prompt and performance was still at chance level after the full training phase. The remaining 41 participants (18 female, age [mean(std)] = 27(3.9)) were included for all analyses. The study was approved by the ethics board of the Technical University of Munich (TUM), and we acquired written informed consent from all participants.

Visual stimuli

We selected 224 full-color images of everyday objects from a larger image set (Brady et al., 2008). The stimuli were chosen to be maximally homogenous regarding salience, e.g., by excluding food items and bright colors. All allocations of stimuli to subjects and conditions were random. Figure 1a visualizes the creation of the visual streams. A pool of 80 images was generated for every subject, with half assigned to the predictable condition and the other half to the unpredictable condition. Based on these images, eight predictable sequences were created for every subject. These objects could only occur in the sequence and position determined during stimulus creation. For the unpredictable condition, a starting set of eight sequences was created. After all were presented during the training or scanning phase, eight completely new sequences were randomly created for every repetition. Consequently, these objects never occurred in the same sequence or position, but the total number of occurrences was the same as for predictable objects. Lastly, only for the scanning phase, half of the predictable sequences and half of the unpredictable objects were combined into surprising sequences. Predictable sequences formed the basis but had one to three objects between position two and five replaced with a random object from the unpredictable condition. The first object was always unchanged to trigger conditional predictions based on the learned transition probabilities.

Experimental design

Implementation. We used Psychopy (Peirce et al., 2019) to implement the design. The online training sessions were realized using pavlovia.org, where Javascript-translated Psychopy experiments can be run online with millisecond precision (Bridges et al., 2020; Sauter et al., 2020).

Main task. We presented participants with continuous visual streams based on the described object sequences. Each stimulus was presented for 800ms with no inter-stimulus interval within sequence and a 1100ms fixation cross between sequences. When all unique sequences of a condition were shown, their order was reshuffled for continued presentation with the constraint that objects (or sequences) could not appear twice in a row. This was done to minimize confounding effects of repetition suppression. To ensure fixation and concentration, participants were instructed to quickly react to occurrences of upside-down objects. These appeared with a probability of 10%, independent of condition.

Training phase. Over the course of three days prior to the main scanning session, participants followed an online implementation of the main design. This phase only included the

predictable und unpredictable sequences. We instructed every participant to perform the training in a quiet environment without distractions. Stimulation blocks lasted 20 minutes per day, with a break of one minute after 10 minutes. Instructions were shown on screen before the stimulation began and the first day included a five-minute familiarization block with on-screen feedback regarding cover-task button presses. The instructions stressed the importance of the cover task, but also noted that questions regarding the order of objects in a sequence would follow each training day. During this task, participants saw eight incomplete sequences (the first one to four objects were shown) for both conditions (to keep image familiarity the same across conditions). After every sequence, participants chose what they assumed would be the correct trailing object from five options and gave a confidence rating on a scale of one to seven. No feedback was given regarding performance and no information on the underlying conditions was disclosed.

Imaging phase. During scanning, stimuli were presented against a grey background and subtended 4° of visual angle. Each condition was presented for three long blocks, during the complete duration of a pcASL sequence (6 minutes), a T2* sequence (5.5 minutes) and a DSC sequence (2.5 minutes). The resulting CMR_{O2} values were calculated as described below (see *MRI data processing*). The order was randomized, although a condition could not occur three times in a row. We left breaks of one minute between each consecutive sequence. During the T2 sequence, there was no experimental stimulation. Per condition, a total of 169 object sequences over a stimulation time of approximately 14 minutes were presented. Before every scanning sequence, a reminder regarding the cover task was shown on screen. Additionally, participants got feedback on their mean reaction time after each sequence to promote attentiveness and motivation. We disclosed no information regarding the sequence patterns and the imaging session included no sequence completion test.

Object detection task. On the day of the imaging session, prior to entering the scanner room, participants performed an object detection task on a laptop. The presentation format followed the specifications of the main task, except for transitions between sequences. We presented eight unpredictable (reshuffled) and eight predictable (as learnt) sequences in random order. Prior to every sequence, a target object from the following sequence was shown. An onscreen prompt instructed participants to react as quickly as possible to the presentation of the target image by pressing the Arrow Up button. During the task, the target image was presented for

five seconds. To allow for anticipation, it could only match positions two to five of the following sequence.

MRI acquisition

For CMR_{O2} mapping, the following sequences were acquired:

- Multi-echo spin-echo T2 mapping: 3D gradient spin echo (GRASE) readout as described previously (Kaczmarz et al., 2020), 8 echoes, TE1 = ΔTE = 16ms, TR=251ms, α=90°, voxel size 2x2x3.3mm^3, 35 slices. T2 data was acquired once per subject, without any task.
- Multi-echo gradient-echo T2* mapping: As described previously (Hirsch et al., 2014; Kaczmarz et al., 2020), 12 echoes, TE1 = ΔTE = 5ms, TR=2229ms, α=30°, voxel size 2x2x3mm^3, gap 0.3mm, 35 slices. T2* was acquired for all conditions.
- Dynamic susceptibility imaging (DSC): As described previously (Hedderich et al., 2019). Injection of gadolinium-based contrast agent as a bolus after 5 dynamic scans, 0.1ml/kg (maximum: 8ml per injection, 16ml per session), flow rate: 4ml/s, plus 25ml NaCl. Single-shot GRE-EPI, EPI factor 49, 80 dynamic scans, TR = 2.0s, α=60°, acquisition voxel size 2x2x3.5mm^3, 35 slices. To stay within the limits of a full clinical dosage (16ml), we acquired DSC in two conditions only: Predictable and unpredictable. Processing of the data in the surprising condition used DSC from the predictable condition.
- Pseudo-continuous arterial spin labeling (pcASL): As described previously (Alsop et al., 2015), and implemented according to (Göttler et al., 2019; Kaczmarz et al., 2020). PLD 1800ms, label duration 1800ms, 4 background suppression pulses, 2D EPI readout, TE=11ms, TR=4500ms, α=90°, 20 slices, EPI factor 29, acquisition voxel size 3.28x3.5x6.0mm^3, gap 0.6mm, 30 dynamic scans including a proton density weighted M0 scan. ASL was acquired for all conditions.

Prior to data acquisition, a venous catheter was placed by a medical doctor through which blood samples were taken and sent to our in-house clinical chemistry laboratory. Creatinin values were analyzed as an indicator of healthy kidney function and contrast agent was only applied for subjects below a threshold of 1.3. No subject exceeded this value. Hemoglobin and hematocrit were requested and used in modelling of CMR_{O2}. Finally, arterial oxygen saturation was measured via a pulse oximeter (Nonin 7500FO, Nonin Medical B.V., The Netherlands).

MRI data processing

To calculate CMR_{O2}, the following parameters were integrated and derived as described below: The oxygen content of blood (O_2 saturation and Hematocrit), the flow of blood (CBF), and the relative oxygen extraction (OEF). Note that the resulting CMR_{O2} values represent a consumption rate for a given condition and are not time-resolved. The processing of the quantitative parameter maps was performed with in-house scripts in MATLAB and SPM12 (Wellcome Trust Centre for Neuroimaging, UCL, London, UK). T2* images were corrected for macroscopic magnetic background gradients with a standard sinc-Gauss excitation pulse (Baudrexel et al., 2009; Hirsch & Preibisch, 2013). Motion correction was performed using redundant acquisitions of k-space center (Nöth et al., 2014). R2' maps are derived from T2 and T2* images and yield the transverse, reversible relaxation rate that is dependent on the vascular dHb content (Blockley et al., 2013, 2015; Bright et al., 2019). However, confounds from uncorrectable strong magnetic field inhomogeneities at air-tissue boundaries, iron deposition in deep GM structures as well as white matter structure need to be considered (Hirsch & Preibisch, 2013; Kaczmarz et al., 2020). The cerebral blood volume (CBV) was derived from DSC MRI via full integration of leakage-corrected ΔR2*-curves (Boxermann, J.L., Schmainda, K.M., Weisskoff, R.M., 2006) and normalization to a white matter value of 2.5% (Leenders et al., 1990) as described previously (Hedderich et al., 2019; Kluge et al., 2016). From R2' and CBV parameter maps, the oxygen extraction fraction (OEF) was calculated (Christen et al., 2012; Hirsch et al., 2014; Yablonskiy & Haacke, 1994). CBF maps were calculated from pcASL data, based on average pairwise differences of motion-corrected label and control images and a proton-density weighted image.

For each subject and condition, we calculated CMR_{O2} in a voxel-wise manner by combining all parameter maps via *Fick's principle*

$$\mathbf{CMR_{O2} = OEF \cdot CBF \cdot Ca_{O2}}$$

where Ca_{O2} is the oxygen carrying capacity of hemoglobin and was calculated as $Ca_{O2} = 0.334 \cdot Hct \cdot 55.6 \cdot O_2sat$, with O_2sat being the oxygen saturation measured by the pulse oximeter (Bright et al., 2019; Y. Ma et al., 2020) and Hct representing Hematocrit as measured by blood sampling prior to scanning. CBF was upscaled by 25% to account for systematic CBF underestimation due to four background-suppression pulses (Garcia et al., 2005; Mutsaerts et al., 2014). All parameter maps of each individual subject were registered to the first echo of their multi echo T2 data.

For statistical analysis of CMR_{O2}, we only included voxels with a grey matter probability of > 0.5. Furthermore, the images were masked using an intersection mask to exclude voxels with excessive susceptibility, indicative of artefacts (T2 and T2* > 120ms, R2' > 9ms) and voxels with biologically unlikely CBF (> 90 ml/min/100g) or OEF (> 90%).

Statistical analysis

Mixed models. We used robust linear mixed models as implemented in the R package robustlmer (Koller, 2016) to minimize the effect of outliers. This method uses the Huber loss function, which is quadratic for small differences, but linear for large differences. The random model term was used to specify conditions as repeated (nested) measurements within subjects. Consequently, our model estimated a subject-wise CMR_{O2} baseline for each condition. The fixed effects included a three-way interaction of condition, confidence and network as well as a two-way interaction of condition and reaction time in the cover task. Confidence values were based on the average condition-wise ratings after the last training day. Since surprising sequences were based on predictable ones, we assigned the same confidence ratings.

The input to our model were condition-wise regional CMR_{O2} values, obtained by the median voxel value within 400 functional areas as defined by the Schaefer parcellation (Schaefer et al., 2018). This parcellation was chosen for two reasons: First, regions are defined by coactivation patterns, meaning that averaging over the respective voxels upholds some functional homogeneity. Secondly, it defines seven distributed functional networks that have been found to vary in their tendency to perform sensory or higher cognitive processing. These include, loosely ordered from higher cognitive to sensory (Margulies et al., 2016): Limbic network, default mode network (DMN), control network (fronto-parietal network), ventral attention network (salience network), dorsal attention network, visual network and somato-motor network. 26 of the 400 areas, mainly in the temporal pole and orbitofrontal cortex, had to be excluded from the analysis. This was due to susceptibility artefacts in the R2' maps, leading to signal dropout in these regions. The limbic network is therefore not covered by our model. Across subjects and areas, a total of 45507 observations entered the model estimation. We z-standardized the behavioral predictors and left the outcome CMR_{O2} values unchanged. For model estimation, we chose the predictable condition and the visual network as reference categories. A significant reference effect in the absence of significant parameter estimates for other categories means that the effect of the predictor does not vary across categorical levels.

This also means that a significant effect for a non-reference category level has to be interpreted in relation to the reference effect. Due to the quantitative nature of CMRO2, the reference effect is itself interpretable as an effect for the predictable condition in the visual network. Finally, we compared our model to a null model without the condition predictor using the performance package (Lüdecke et al., 2021) for standard mixed models as implemented in lme4. Since the model outcome represents absolute CMR_{O2}, parameter estimates can be interpreted as the predicted change in energy cost for a given change in standardized predictors.

Model-based CMR_{O2} cost predictions. For model-based predictions, we added the corresponding parameter estimates as described in the results. Total CMR_{O2} was therefore calculated as the sum of the intercept (the sample baseline for predictable input in the visual network), the network effect (the difference of the network energy consumption from the visual network) and the reference effect of confidence. We calculated the predicted cortical CMR_{O2} for weakly confident (5^{th} percentile, confidence(z)=-1.96), average (50^{th} percentile, confidence(z)=0) and highly confident (95^{th} percentile, confidence(z)=1.96) subjects. The outcome was then scaled to the grey matter mass of each network in MNI space, approximated by grey matter voxel count multiplied by a tissue mass of 0.0014 gram per cubic millimeter (Barber et al., 1970; IT'IS Foundation, 2022). Finally, predicted CMR_{O2} was summed over networks to obtain cortical values. We repeated this procedure for the 5^{th} and 95^{th} percentile of the confidence parameter estimate (Table 1) to obtain an upper and lower bound of the predicted energy cost.

Acknowledgements

The authors want to thank Christine Preibisch for providing MRI protocols and template code for acquisition and processing of mqBOLD MRI data.

References

Acerbi, L., Vijayakumar, S., & Wolpert, D. M. (2014). On the Origins of Suboptimality in Human Probabilistic Inference. *PLoS Computational Biology*, *10*(6), e1003661. https://doi.org/10.1371/journal.pcbi.1003661

Adler, W. T., & Ma, W. J. (2018). Comparing Bayesian and non-Bayesian accounts of human confidence reports. *PLOS Computational Biology*, *14*(11), e1006572. https://doi.org/10.1371/journal.pcbi.1006572

Aizenstein, H. J. (2004). Regional Brain Activation during Concurrent Implicit and Explicit Sequence Learning. *Cerebral Cortex*, *14*(2), 199–208. https://doi.org/10.1093/cercor/bhg119

Alamia, A., Orban de Xivry, J.-J., San Anton, E., Olivier, E., Cleeremans, A., & Zenon, A. (2016). Unconscious associative learning with conscious cues. *Neuroscience of Consciousness*, *2016*(1). https://doi.org/10.1093/nc/niw016

Ali, A., Ahmad, N., De Groot, E., Johannes Van Gerven, M. A., & Kietzmann, T. C. (2022). Predictive coding is a consequence of energy efficiency in recurrent neural networks. *Patterns*, *3*(12), 100639. https://doi.org/10.1016/j.patter.2022.100639

Alsop, D. C., Detre, J. A., Golay, X., Günther, M., Hendrikse, J., Hernandez-Garcia, L., Lu, H., MacIntosh, B. J., Parkes, L. M., Smits, M., van Osch, M. J. P., Wang, D. J. J., Wong, E. C., & Zaharchuk, G. (2015). Recommended implementation of arterial spin-labeled perfusion MRI for clinical applications: A consensus of the ISMRM perfusion study group and the European consortium for ASL in dementia: Recommended Implementation of ASL for Clinical Applications. *Magnetic Resonance in Medicine*, *73*(1), 102–116. https://doi.org/10.1002/mrm.25197

Bach, D. R., & Dolan, R. J. (2012). Knowing how much you don't know: A neural organization of uncertainty estimates. *Nature Reviews Neuroscience*, *13*(8), 572–586. https://doi.org/10.1038/nrn3289

Barber, T. W., Brockway, J. A., & Higgins, L. S. (1970). THE DENSITY OF TISSUES IN AND ABOUT THE HEAD. *Acta Neurologica Scandinavica*, *46*(1), 85–92. https://doi.org/10.1111/j.1600-0404.1970.tb05606.x

Bastos, A. M., Usrey, W. M., Adams, R. A., Mangun, G. R., Fries, P., & Friston, K. (2012). Canonical Microcircuits for Predictive Coding. *Neuron*, *76*(4), 695–711. https://doi.org/10.1016/j.neuron.2012.10.038

Batterink, L. J., Reber, P. J., Neville, H. J., & Paller, K. A. (2015). Implicit and explicit contributions to statistical learning. *Journal of Memory and Language*, *83*, 62–78. https://doi.org/10.1016/j.jml.2015.04.004

Baudrexel, S., Volz, S., Preibisch, C., Klein, J. C., Steinmetz, H., Hilker, R., & Deichmann, R. (2009). Rapid single-scan T2*-mapping using exponential excitation pulses and image-based correction for linear background gradients: T2* Mapping With Field

Gradient Correction. *Magnetic Resonance in Medicine*, *62*(1), 263–268. https://doi.org/10.1002/mrm.21971

Beck, J. M., Ma, W. J., Pitkow, X., Latham, P. E., & Pouget, A. (2012). Not Noisy, Just Wrong: The Role of Suboptimal Inference in Behavioral Variability. *Neuron*, *74*(1), 30–39. https://doi.org/10.1016/j.neuron.2012.03.016

Beierholm, U., Rohe, T., Ferrari, A., Stegle, O., & Noppeney, U. (2020). Using the past to estimate sensory uncertainty. *eLife*, *9*, e54172. https://doi.org/10.7554/eLife.54172

Blockley, N. P., Griffeth, V. E. M., Simon, A. B., & Buxton, R. B. (2013). A review of calibrated blood oxygenation level-dependent (BOLD) methods for the measurement of task-induced changes in brain oxygen metabolism: A REVIEW OF CALIBRATED BOLD METHODS. *NMR in Biomedicine*, *26*(8), 987–1003. https://doi.org/10.1002/nbm.2847

Blockley, N. P., Griffeth, V. E. M., Simon, A. B., Dubowitz, D. J., & Buxton, R. B. (2015). Calibrating the BOLD response without administering gases: Comparison of hypercapnia calibration with calibration using an asymmetric spin echo. *NeuroImage*, *104*, 423–429. https://doi.org/10.1016/j.neuroimage.2014.09.061

Bounmy, T., Eger, E., & Meyniel, F. (2023). A characterization of the neural representation of confidence during probabilistic learning. *NeuroImage*, *268*, 119849. https://doi.org/10.1016/j.neuroimage.2022.119849

Boxermann, J.L., Schmainda, K.M., Weisskoff, R.M. (2006). Relative Cerebral Blood Volume Maps Corrected for Contrast Agent Extravasation Significantly Correlate with Glioma Tumor Grade, Whereas Uncorrected Maps Do Not. *Amercian Journal of Neuroradiology*, *27*(4), 859–867.

Brady, T., Konkle, T., Alvarez, G. A., & Oliva, A. (2008). Visual long-term memory has a massive storage capacity for object details. *Proceedings of the National Academy of Sciences*, *105*(38), 14325–14329. https://doi.org/10.1073/pnas.0803390105

Bridges, D., Pitiot, A., MacAskill, M. R., & Peirce, J. W. (2020). The timing mega-study: Comparing a range of experiment generators, both lab-based and online. *PeerJ*, *8*, e9414. https://doi.org/10.7717/peerj.9414

Bright, M. G., Croal, P. L., Blockley, N. P., & Bulte, D. P. (2019). Multiparametric measurement of cerebral physiology using calibrated fMRI. *NeuroImage*, *187*, 128–144. https://doi.org/10.1016/j.neuroimage.2017.12.049

Chalk, M., Marre, O., & Tkačik, G. (2018). Toward a unified theory of efficient, predictive, and sparse coding. *Proceedings of the National Academy of Sciences*, *115*(1), 186–191. https://doi.org/10.1073/pnas.1711114115

Christen, T., Schmiedeskamp, H., Straka, M., Bammer, R., & Zaharchuk, G. (2012). Measuring brain oxygenation in humans using a multiparametric quantitative blood oxygenation level dependent MRI approach. *Magnetic Resonance in Medicine*, *68*(3), 905–911. https://doi.org/10.1002/mrm.23283

Christie, S. T., & Schrater, P. (2015). Cognitive cost as dynamic allocation of energetic resources. *Frontiers in Neuroscience*, *9*. https://doi.org/10.3389/fnins.2015.00289

Clark, A. (2013). Whatever next? Predictive brains, situated agents, and the future of cognitive science. *Behavioral and Brain Sciences*, *36*(3), 181–204. https://doi.org/10.1017/S0140525X12000477

Commins, S. (2018). Efficiency: An underlying principle of learning? *Reviews in the Neurosciences*, *29*(2), 183–197. https://doi.org/10.1515/revneuro-2017-0050

Conway, C. M. (2020). How does the brain learn environmental structure? Ten core principles for understanding the neurocognitive mechanisms of statistical learning. *Neuroscience & Biobehavioral Reviews*, *112*, 279–299. https://doi.org/10.1016/j.neubiorev.2020.01.032

Dale, R., Duran, N., & Morehead, R. (2012). Prediction during statistical learning, and implications for the implicit/explicit divide. *Advances in Cognitive Psychology*, *8*(2), 196–209. https://doi.org/10.5709/acp-0115-z

de Lange, F. P., Heilbron, M., & Kok, P. (2018). How Do Expectations Shape Perception? *Trends in Cognitive Sciences*, *22*(9), 764–779. https://doi.org/10.1016/j.tics.2018.06.002

Devonshire, I. M., Papadakis, N. G., Port, M., Berwick, J., Kennerley, A. J., Mayhew, J. E. W., & Overton, P. G. (2012). Neurovascular coupling is brain region-dependent. *NeuroImage*, *59*(3), 1997–2006. https://doi.org/10.1016/j.neuroimage.2011.09.050

Dienel, G. A. (2014). Energy Metabolism of the Brain. In *From Molecules to Networks: An Introduction to Cellular and Molecular Neuroscience* (Third Edition, S. 53–117). Elsevier/Academic Press.

Drew, P. J. (2019). Vascular and neural basis of the BOLD signal. *Current Opinion in Neurobiology*, *58*, 61–69. https://doi.org/10.1016/j.conb.2019.06.004

Drew, P. J. (2022). Neurovascular coupling: Motive unknown. *Trends in Neurosciences*, *45*(11), 809–819. https://doi.org/10.1016/j.tins.2022.08.004

Feuerriegel, D., Vogels, R., & Kovács, G. (2021). Evaluating the evidence for expectation suppression in the visual system. *Neuroscience & Biobehavioral Reviews*, *126*, 368–381. https://doi.org/10.1016/j.neubiorev.2021.04.002

Feuerriegel, D., Yook, J., Quek, G. L., Hogendoorn, H., & Bode, S. (2021). Visual mismatch responses index surprise signalling but not expectation suppression. *Cortex*, *134*, 16–29. https://doi.org/10.1016/j.cortex.2020.10.006

Ficco, L., Mancuso, L., Manuello, J., Teneggi, A., Liloia, D., Duca, S., Costa, T., Kovacs, G. Z., & Cauda, F. (2021). Disentangling predictive processing in the brain: A meta-analytic study in favour of a predictive network. *Scientific Reports*, *11*(1), 16258. https://doi.org/10.1038/s41598-021-95603-5

Findling, C., Hubert, F., International Brain Laboratory, Acerbi, L., Benson, B., Benson, J., Birman, D., Bonacchi, N., Carandini, M., Catarino, J. A., Chapuis, G. A., Churchland, A. K., Dan, Y., DeWitt, E. E., Engel, T. A., Fabbri, M., Faulkner, M., Fiete, I. R., Freitas-Silva, L., … Pouget, A. (2023). *Brain-wide representations of prior information in mouse decision-making* [Preprint]. Neuroscience. https://doi.org/10.1101/2023.07.04.547684

Fiser, J., & Lengyel, G. (2022). Statistical Learning in Vision. *Annual Review of Vision Science*, *8*(1), 265–290. https://doi.org/10.1146/annurev-vision-100720-103343

Fox, P., Raichle, M., Mintun, M., & Dence, C. (1988). Nonoxidative glucose consumption during focal physiologic neural activity. *Science*, *241*(4864), 462–464. https://doi.org/10.1126/science.3260686

Friston, K. (2010). The free-energy principle: A unified brain theory? *Nature Reviews Neuroscience*, *11*(2), 127–138. https://doi.org/10.1038/nrn2787

Garcia, D. M., Duhamel, G., & Alsop, D. C. (2005). Efficiency of inversion pulses for background suppressed arterial spin labeling. *Magnetic Resonance in Medicine*, *54*(2), 366–372. https://doi.org/10.1002/mrm.20556

Geurts, L. S., Cooke, J. R. H., Van Bergen, R. S., & Jehee, J. F. M. (2022). Subjective confidence reflects representation of Bayesian probability in cortex. *Nature Human Behaviour*, *6*(2), 294–305. https://doi.org/10.1038/s41562-021-01247-w

Göttler, J., Kaczmarz, S., Kallmayer, M., Wustrow, I., Eckstein, H.-H., Zimmer, C., Sorg, C., Preibisch, C., & Hyder, F. (2019). Flow-metabolism uncoupling in patients with asymptomatic unilateral carotid artery stenosis assessed by multi-modal magnetic resonance imaging. *Journal of Cerebral Blood Flow & Metabolism*, *39*(11), 2132–2143. https://doi.org/10.1177/0271678X18783369

Griffeth, V. E. M., Simon, A. B., & Buxton, R. B. (2015). The coupling of cerebral blood flow and oxygen metabolism with brain activation is similar for simple and complex stimuli in human primary visual cortex. *NeuroImage*, *104*, 156–162. https://doi.org/10.1016/j.neuroimage.2014.10.003

Harris, J. J., Jolivet, R., & Attwell, D. (2012). Synaptic Energy Use and Supply. *Neuron*, *75*(5), 762–777. https://doi.org/10.1016/j.neuron.2012.08.019

Hedderich, D., Kluge, A., Pyka, T., Zimmer, C., Kirschke, J. S., Wiestler, B., & Preibisch, C. (2019). Consistency of normalized cerebral blood volume values in glioblastoma using different leakage correction algorithms on dynamic susceptibility contrast magnetic resonance imaging data without and with preload. *Journal of Neuroradiology*, *46*(1), 44–51. https://doi.org/10.1016/j.neurad.2018.04.006

Hirsch, N. M., & Preibisch, C. (2013). T2* Mapping with Background Gradient Correction Using Different Excitation Pulse Shapes. *American Journal of Neuroradiology*, *34*(6), E65–E68. https://doi.org/10.3174/ajnr.A3021

Hirsch, N. M., Toth, V., Förschler, A., Kooijman, H., Zimmer, C., & Preibisch, C. (2014). Technical considerations on the validity of blood oxygenation level-dependent-based MR assessment of vascular deoxygenation: BOLD-BASED ASSESSMENT OF VASCULAR DEOXYGENATION. *NMR in Biomedicine*, *27*(7), 853–862. https://doi.org/10.1002/nbm.3131

Howarth, C., Gleeson, P., & Attwell, D. (2012). Updated Energy Budgets for Neural Computation in the Neocortex and Cerebellum. *Journal of Cerebral Blood Flow & Metabolism*, *32*(7), 1222–1232. https://doi.org/10.1038/jcbfm.2012.35

Huettel, S. A., Mack, P. B., & McCarthy, G. (2002). Perceiving patterns in random series: Dynamic processing of sequence in prefrontal cortex. *Nature Neuroscience*, *5*(5), 485–490. https://doi.org/10.1038/nn841

Hyder, D. S. F. (2010). Neurovascular and neurometabolic couplings in dynamic calibrated fMRI: transient oxidative neuroenergetics for block-design and event-related

paradigms. *Frontiers in Neuroenergetics*, *2*. https://doi.org/10.3389/fnene.2010.00018

IT'IS Foundation. (2022). *Tissue Properties Database V4.1* [..Db, .xls, .txt, .pdf]. IT'IS Foundation. https://doi.org/10.13099/VIP21000-04-1

Kaczmarz, S., Hyder, F., & Preibisch, C. (2020). Oxygen extraction fraction mapping with multi-parametric quantitative BOLD MRI: Reduced transverse relaxation bias using 3D-GraSE imaging. *NeuroImage*, *220*, 117095. https://doi.org/10.1016/j.neuroimage.2020.117095

Kluge, A., Lukas, M., Toth, V., Pyka, T., Zimmer, C., & Preibisch, C. (2016). Analysis of three leakage-correction methods for DSC-based measurement of relative cerebral blood volume with respect to heterogeneity in human gliomas. *Magnetic Resonance Imaging*, *34*(4), 410–421. https://doi.org/10.1016/j.mri.2015.12.015

Knill, D. C., & Pouget, A. (2004). The Bayesian brain: The role of uncertainty in neural coding and computation. *Trends in Neurosciences*, *27*(12), 712–719. https://doi.org/10.1016/j.tins.2004.10.007

Koblinger, Á., Fiser, J., & Lengyel, M. (2021). Representations of uncertainty: Where art thou? *Current Opinion in Behavioral Sciences*, *38*, 150–162. https://doi.org/10.1016/j.cobeha.2021.03.009

Koller, M. (2016). robustlmm: An R Package for Robust Estimation of Linear Mixed-Effects Models. *Journal of Statistical Software*, *75*(6). https://doi.org/10.18637/jss.v075.i06

Lee, P., Chandel, N. S., & Simon, M. C. (2020). Cellular adaptation to hypoxia through hypoxia inducible factors and beyond. *Nature Reviews Molecular Cell Biology*, *21*(5), 268–283. https://doi.org/10.1038/s41580-020-0227-y

Leenders, K. L., Perani, D., Lammertsma, A. A., Heather, J. D., Buckingham, P., Jones, T., Healy, M. J. R., Gibbs, J. M., Wise, R. J. S., Hatazawa, J., Herold, S., Beaney, R. P., Brooks, D. J., Spinks, T., Rhodes, C., & Frackowiak, R. S. J. (1990). CEREBRAL BLOOD FLOW, BLOOD VOLUME AND OXYGEN UTILIZATION: NORMAL VALUES AND EFFECT OF AGE. *Brain*, *113*(1), 27–47. https://doi.org/10.1093/brain/113.1.27

Lin, A.-L., Fox, P. T., Hardies, J., Duong, T. Q., & Gao, J.-H. (2010). Nonlinear coupling between cerebral blood flow, oxygen consumption, and ATP production in human visual cortex. *Proceedings of the National Academy of Sciences*, *107*(18), 8446–8451. https://doi.org/10.1073/pnas.0909711107

Luczak, A., McNaughton, B. L., & Kubo, Y. (2022). Neurons learn by predicting future activity. *Nature Machine Intelligence*, *4*(1), 62–72. https://doi.org/10.1038/s42256-021-00430-y

Lüdecke, D., Ben-Shachar, M., Patil, I., Waggoner, P., & Makowski, D. (2021). performance: An R Package for Assessment, Comparison and Testing of Statistical Models. *Journal of Open Source Software*, *6*(60), 3139. https://doi.org/10.21105/joss.03139

Ma, W. J., & Jazayeri, M. (2014). Neural Coding of Uncertainty and Probability. *Annual Review of Neuroscience*, *37*(1), 205–220. https://doi.org/10.1146/annurev-neuro-071013-014017

Ma, Y., Sun, H., Cho, J., Mazerolle, E. L., Wang, Y., & Pike, G. B. (2020). Cerebral OEF quantification: A comparison study between quantitative susceptibility mapping and dual-gas calibrated BOLD imaging. *Magnetic Resonance in Medicine*, *83*(1), 68–82. https://doi.org/10.1002/mrm.27907

Maheu, M., Meyniel, F., & Dehaene, S. (2022). Rational arbitration between statistics and rules in human sequence processing. *Nature Human Behaviour*, *6*(8), 1087–1103. https://doi.org/10.1038/s41562-021-01259-6

Manahova, M. E., Mostert, P., Kok, P., Schoffelen, J.-M., & de Lange, F. P. (2018). Stimulus Familiarity and Expectation Jointly Modulate Neural Activity in the Visual Ventral Stream. *Journal of Cognitive Neuroscience*, *30*(9), 1366–1377. https://doi.org/10.1162/jocn_a_01281

Margulies, D. S., Ghosh, S. S., Goulas, A., Falkiewicz, M., Huntenburg, J. M., Langs, G., Bezgin, G., Eickhoff, S. B., Castellanos, F. X., Petrides, M., Jefferies, E., & Smallwood, J. (2016). Situating the default-mode network along a principal gradient of macroscale cortical organization. *Proceedings of the National Academy of Sciences*, *113*(44), 12574–12579. https://doi.org/10.1073/pnas.1608282113

Meyniel, F., & Dehaene, S. (2017). Brain networks for confidence weighting and hierarchical inference during probabilistic learning. *Proceedings of the National Academy of Sciences*, *114*(19). https://doi.org/10.1073/pnas.1615773114

Meyniel, F., Schlunegger, D., & Dehaene, S. (2015). The Sense of Confidence during Probabilistic Learning: A Normative Account. *PLOS Computational Biology*, *11*(6), e1004305. https://doi.org/10.1371/journal.pcbi.1004305

Meyniel, F., Sigman, M., & Mainen, Z. F. (2015). Confidence as Bayesian Probability: From Neural Origins to Behavior. *Neuron*, *88*(1), 78–92. https://doi.org/10.1016/j.neuron.2015.09.039

Moradi, F., Buračas, G. T., & Buxton, R. B. (2012). Attention strongly increases oxygen metabolic response to stimulus in primary visual cortex. *NeuroImage*, *59*(1), 601–607. https://doi.org/10.1016/j.neuroimage.2011.07.078

Moradi, F., & Buxton, R. B. (2013). Adaptation of cerebral oxygen metabolism and blood flow and modulation of neurovascular coupling with prolonged stimulation in human visual cortex. *NeuroImage*, *82*, 182–189. https://doi.org/10.1016/j.neuroimage.2013.05.110

Mutsaerts, H. J. M. M., Steketee, R. M. E., Heijtel, D. F. R., Kuijer, J. P. A., van Osch, M. J. P., Majoie, C. B. L. M., Smits, M., & Nederveen, A. J. (2014). Inter-Vendor Reproducibility of Pseudo-Continuous Arterial Spin Labeling at 3 Tesla. *PLoS ONE*, *9*(8), e104108. https://doi.org/10.1371/journal.pone.0104108

Noble, S., Curtiss, J., Pessoa, L., & Scheinost, D. (2023). *The tip of the iceberg: A call to embrace anti-localizationism in human neuroscience research*. PsyArXiv. https://doi.org/10.31234/osf.io/9eqh6

Noble, S., Mejia, A. F., Zalesky, A., & Scheinost, D. (2022). Improving power in functional magnetic resonance imaging by moving beyond cluster-level inference. *Proceedings of the National Academy of Sciences*, *119*(32), e2203020119. https://doi.org/10.1073/pnas.2203020119

Peirce, J., Gray, J. R., Simpson, S., MacAskill, M., Höchenberger, R., Sogo, H., Kastman, E., & Lindeløv, J. K. (2019). PsychoPy2: Experiments in behavior made easy. *Behavior Research Methods*, *51*(1), 195–203. https://doi.org/10.3758/s13428-018-01193-y

Pessoa, L. (2023). The Entangled Brain. *Journal of Cognitive Neuroscience*, 1–12. https://doi.org/10.1162/jocn_a_01908

Pouget, A., Drugowitsch, J., & Kepecs, A. (2016). Confidence and certainty: Distinct probabilistic quantities for different goals. *Nature Neuroscience*, *19*(3), 366–374. https://doi.org/10.1038/nn.4240

Quintela-López, T., Shiina, H., & Attwell, D. (2022). Neuronal energy use and brain evolution. *Current Biology*, *32*(12), R650–R655. https://doi.org/10.1016/j.cub.2022.02.005

Raichle, M. E., & Mintun, M. A. (2006). BRAIN WORK AND BRAIN IMAGING. *Annual Review of Neuroscience*, *29*(1), 449–476. https://doi.org/10.1146/annurev.neuro.29.051605.112819

Rao, R. P. N., & Ballard, D. H. (1999). Predictive coding in the visual cortex: A functional interpretation of some extra-classical receptive-field effects. *Nature Neuroscience*, *2*(1), 79–87. https://doi.org/10.1038/4580

Richter, D., Ekman, M., & de Lange, F. P. (2018). Suppressed Sensory Response to Predictable Object Stimuli throughout the Ventral Visual Stream. *The Journal of Neuroscience*, *38*(34), 7452–7461. https://doi.org/10.1523/JNEUROSCI.3421-17.2018

Roberts, J. A., Iyer, K. K., Vanhatalo, S., & Breakspear, M. (2014). Critical role for resource constraints in neural models. *Frontiers in Systems Neuroscience*, *8*. https://doi.org/10.3389/fnsys.2014.00154

Sanders, J. I., Hangya, B., & Kepecs, A. (2016). Signatures of a Statistical Computation in the Human Sense of Confidence. *Neuron*, *90*(3), 499–506. https://doi.org/10.1016/j.neuron.2016.03.025

Sauter, M., Draschkow, D., & Mack, W. (2020). Building, Hosting and Recruiting: A Brief Introduction to Running Behavioral Experiments Online. *Brain Sciences*, *10*(4), 251. https://doi.org/10.3390/brainsci10040251

Schaefer, A., Kong, R., Gordon, E. M., Laumann, T. O., Zuo, X.-N., Holmes, A. J., Eickhoff, S. B., & Yeo, B. T. T. (2018). Local-Global Parcellation of the Human Cerebral Cortex from Intrinsic Functional Connectivity MRI. *Cerebral Cortex*, *28*(9), 3095–3114. https://doi.org/10.1093/cercor/bhx179

Schapiro, A., & Turk-Browne, N. (2015). Statistical Learning. In *Brain Mapping* (S. 501–506). Elsevier. https://doi.org/10.1016/B978-0-12-397025-1.00276-1

Sengupta, B., Stemmler, M. B., & Friston, K. (2013). Information and Efficiency in the Nervous System—A Synthesis. *PLoS Computational Biology*, *9*(7), e1003157. https://doi.org/10.1371/journal.pcbi.1003157

Stefanics, G., Kimura, M., & Czigler, I. (2011). Visual Mismatch Negativity Reveals Automatic Detection of Sequential Regularity Violation. *Frontiers in Human Neuroscience*, *5*. https://doi.org/10.3389/fnhum.2011.00046

Teufel, C., & Fletcher, P. C. (2020). Forms of prediction in the nervous system. *Nature Reviews Neuroscience*, *21*(4), 231–242. https://doi.org/10.1038/s41583-020-0275-5

Theriault, J. E., Shaffer, C., Dienel, G. A., Sander, C. Y., Hooker, J. M., Dickerson, B. C., Barrett, L. F., & Quigley, K. S. (2023). A Functional Account of Stimulation-based Aerobic Glycolysis and its Role in Interpreting BOLD Signal Intensity Increases in Neuroimaging Experiments. *Neuroscience & Biobehavioral Reviews*, 105373. https://doi.org/10.1016/j.neubiorev.2023.105373

Turk-Browne, N. B., Jungé, J. A., & Scholl, B. J. (2005). The Automaticity of Visual Statistical Learning. *Journal of Experimental Psychology: General*, *134*(4), 552–564. https://doi.org/10.1037/0096-3445.134.4.552

Turk-Browne, N. B., Scholl, B. J., Chun, M. M., & Johnson, M. K. (2009). Neural Evidence of Statistical Learning: Efficient Detection of Visual Regularities Without Awareness. *Journal of Cognitive Neuroscience*, *21*(10), 1934–1945. https://doi.org/10.1162/jocn.2009.21131

Turk-Browne, N. B., Scholl, B. J., Johnson, M. K., & Chun, M. M. (2010). Implicit Perceptual Anticipation Triggered by Statistical Learning. *Journal of Neuroscience*, *30*(33), 11177–11187. https://doi.org/10.1523/JNEUROSCI.0858-10.2010

Vadillo, M. A., Konstantinidis, E., & Shanks, D. R. (2016). Underpowered samples, false negatives, and unconscious learning. *Psychonomic Bulletin & Review*, *23*(1), 87–102. https://doi.org/10.3758/s13423-015-0892-6

Walsh, K. S., McGovern, D. P., Clark, A., & O'Connell, R. G. (2020). Evaluating the neurophysiological evidence for predictive processing as a model of perception. *Annals of the New York Academy of Sciences*, *1464*(1), 242–268. https://doi.org/10.1111/nyas.14321

Warren, R. E., & Frier, B. M. (2005). Hypoglycaemia and cognitive function. *Diabetes, Obesity and Metabolism*, *7*(5), 493–503. https://doi.org/10.1111/j.1463-1326.2004.00421.x

Xu, F., Ge, Y., & Lu, H. (2009). Noninvasive quantification of whole-brain cerebral metabolic rate of oxygen ($CMRO_2$) by MRI: Quantification of $CMRO_2$. *Magnetic Resonance in Medicine*, *62*(1), 141–148. https://doi.org/10.1002/mrm.21994

Yablonskiy, D. A., & Haacke, E. M. (1994). Theory of NMR signal behavior in magnetically inhomogeneous tissues: The static dephasing regime. *Magnetic Resonance in Medicine*, *32*(6), 749–763. https://doi.org/10.1002/mrm.1910320610

Yon, D., & Frith, C. D. (2021). Precision and the Bayesian brain. *Current Biology*, *31*(17), R1026–R1032. https://doi.org/10.1016/j.cub.2021.07.044

Yu, L., & Yu, Y. (2017). Energy-efficient neural information processing in individual neurons and neuronal networks: Energy Efficiency in Neural Systems. *Journal of Neuroscience Research*, *95*(11), 2253–2266. https://doi.org/10.1002/jnr.24131

Zénon, A., Solopchuk, O., & Pezzulo, G. (2019). An information-theoretic perspective on the costs of cognition. *Neuropsychologia*, *123*, 5–18. https://doi.org/10.1016/j.neuropsychologia.2018.09.013

Zhou, D., Lynn, C. W., Cui, Z., Ciric, R., Baum, G. L., Moore, T. M., Roalf, D. R., Detre, J. A., Gur, R. C., Gur, R. E., Satterthwaite, T. D., & Bassett, D. S. (2022). Efficient coding in the economics of human brain connectomics. *Network Neuroscience*, *6*(1), 234–274. https://doi.org/10.1162/netn_a_00223

Supplementary material

	CMRO2	
Predictors	*Estimates*	*p*
(Intercept)	136.33 (128.84 – 143.83)	**<0.001**
condition [surp]	-4.49 (-9.81 – 0.84)	0.099
condition [unpred]	-4.03 (-10.66 – 2.59)	0.233
confidence	-5.39 (-9.88 – -0.90)	**0.019**
network [Cont]	10.91 (7.28 – 14.54)	**<0.001**
network [Default]	8.26 (5.08 – 11.43)	**<0.001**
network [DorsAttn]	1.04 (-2.68 – 4.77)	0.583
network [SalVentAttn]	-17.67 (-21.37 – -13.96)	**<0.001**
network [SomMot]	-11.81 (-15.08 – -8.54)	**<0.001**
covertask RT	-4.84 (-7.73 – -1.95)	**0.001**
condition [surp] × confidence	5.67 (0.68 – 10.67)	**0.026**
condition [unpred] × confidence	5.10 (-2.26 – 12.45)	0.174
condition [surp] × network [Cont]	-2.99 (-8.12 – 2.14)	0.254
condition [unpred] × network [Cont]	6.51 (0.80 – 12.21)	**0.025**
condition [surp] × network [Default]	-2.89 (-7.38 – 1.61)	0.208
condition [unpred] × network [Default]	-0.77 (-5.77 – 4.22)	0.761
condition [surp] × network [DorsAttn]	-1.76 (-7.04 – 3.51)	0.512
condition [unpred] × network [DorsAttn]	-1.13 (-6.99 – 4.73)	0.706
condition [surp] × network [SalVentAttn]	-0.45 (-5.69 – 4.80)	0.868
condition [unpred] × network [SalVentAttn]	3.22 (-2.60 – 9.05)	0.278
condition [surp] × network [SomMot]	-1.30 (-5.93 – 3.33)	0.582
condition [unpred] × network [SomMot]	0.58 (-4.57 – 5.72)	0.826
confidence × network [Cont]	5.03 (1.70 – 8.35)	**0.003**
confidence × network [Default]	0.72 (-2.19 – 3.63)	0.627
confidence × network [DorsAttn]	-2.77 (-6.19 – 0.65)	0.112
confidence × network [SalVentAttn]	2.28 (-1.12 – 5.68)	0.190
confidence × network [SomMot]	1.16 (-1.84 – 4.17)	0.448
condition [surp] × covertask RT	-0.89 (-4.05 – 2.27)	0.580
condition [unpred] × covertask RT	3.12 (0.16 – 6.08)	**0.039**
(condition [surp] × confidence) × network [Cont]	0.30 (-4.40 – 5.00)	0.899
(condition [unpred] × confidence) × network [Cont]	-0.27 (-6.33 – 5.78)	0.930
(condition [surp] × confidence) × network [Default]	1.19 (-2.93 – 5.31)	0.572
(condition [unpred] × confidence) × network [Default]	-1.63 (-6.93 – 3.68)	0.548
(condition [surp] × confidence) × network [DorsAttn]	0.04 (-4.79 – 4.88)	0.986
(condition [unpred] × confidence) × network [DorsAttn]	4.75 (-1.48 – 10.99)	0.135
(condition [surp] × confidence) × network [SalVentAttn]	-0.13 (-4.94 – 4.68)	0.957
(condition [unpred] × confidence) × network [SalVentAttn]	-0.24 (-6.44 – 5.95)	0.938
(condition [surp] × confidence) × network [SomMot]	0.79 (-3.45 – 5.04)	0.715
(condition [unpred] × confidence) × network [SomMot]	-0.99 (-6.46 – 4.49)	0.723
Random Effects		
σ^2	1611.66	
$\tau_{00\ condition:subject}$	32.06	
$\tau_{00\ subject}$	381.16	
ICC	0.20	
$N_{condition}$	3	
$N_{subject}$	41	
Observations	45507	
Marginal R^2 / Conditional R^2	0.057 / 0.250	

Table S1. Full results of the robust linear mixed model with the following formula (following R convention): condition*confidence*network + condition:reaction_time + (1|subject/condition)

4 General discussion

In the present work, I combined current theories of cognitive function with metabolic imaging to investigate the energetic impact of predictive processing in the human brain. I focused on predictive processing in the temporal domain, as this allowed me to hold low-level attributes of the stimuli constant. Across a training period of three days, participants saw both deterministic and random sequences of visual objects. This led to both explicit and implicit learning effects regarding deterministic sequences. First, participants improved in reproducing their composition. Second, they detected these objects more quickly than objects from random sequences. The latter is indicative of anticipation, a hallmark of predictive processing.

While previous studies relied on defining the predictability of visual stimulation purely by its objective structure, I expected a modulating effect of subjective uncertainty. According to Bayesian formulations of cognition, humans constantly infer the most likely current and future state of their environment in a probabilistic manner. I showed that participants indeed varied in their subjective uncertainty when tasked to reproduce the sequences. While most were highly confident after three days of training, others remained insecure. I also observed considerable variance in confidence ratings for random sequences, although these stayed constant over the training phase. Relating explicit and implicit markers of learning, I showed that the improved detection speed of predictable objects scaled with confidence ratings.

In addition to the behavioral validation of my design, I used fMRI to validate its effects on BOLD activity. Unlike previous work, stimuli were presented in a block design that better corresponded to the long acquisition times of the subsequent metabolic imaging sequences. I replicated both modality-specific prediction error activity in the visual stream and modality-general clusters of prediction in the posterior cortex. Analyzing the effect of confidence, I provided the first evidence that the magnitude of predictive BOLD activity depends on the level of confidence regarding the presented patterns.

Building on these findings, I acquired multiparametric quantitative BOLD data and calculated the regional rate of oxygen consumption for the same design. Instead of using mass-univariate statistics to localize areas of maximum effect, my focus was on the net effect of energy consumption on the level of cortical networks. First, echoing my BOLD results, the effect of sensory predictability on energy consumption depended on the confidence level. Predictable input strongly decreased in cost for increasing confidence levels, while surprising input was not affected. Studies using fMRI or EEG consistently find surprising input to elicit

stronger brain activity than predictable input. My results suggest that an impact on large-scale energy use only emerges at high levels of confidence. Intriguingly, the cost-decreasing effect of confidence was not different for unpredictable sequences. This points to a possible independence of the cost-saving effect of confident predictions from their accuracy. Finally, higher confidence was linked to quicker reaction times in the cover task, which was independent of sequence predictability. As quicker reaction times led to higher energy usage, this partially offset the decreases in energy consumption for predictable input. One interpretation is that participants with high confidence dynamically reallocated their energetic resources.

I framed these projects in the context of efficiency in the brain – predictive processing is expected to result in perception that is maximally informative while staying minimally costly. To operationalize the cost aspect of efficiency, I quantified a biologically valid measure of energy metabolism. The increased informative content of perception was reflected in more accurate and quicker behavior with regards to predictable sequences. However, information is a well-defined mathematical quantity. A complete image of efficient processing would thus benefit from rigorous quantification of both energy consumption and information content. In the following paragraphs, I will therefore use information theoretic concepts to expand upon my results.

4.1 Efficiency and information theory

Over half a century ago, Barlow (1961) proposed an influential theory of neuronal function: The brain reduces the redundancy of sensory information, essentially stripping away all aspects that do not meaningfully add to its identification of interpretation. This is now known as the efficient coding hypothesis (Loh & Bartulovic, 2014; Simoncelli, 2003) and is formulated in the language of information theory (an introduction is given in Timme & Lapish, 2018). Neuronal firing carries information about the external world, reducing the organism's uncertainty about the state of the environment. Information is measured in *bits*, with each bit adjudicating between two competing interpretations. A bitrate, then, refers to the amount of information transmitted in a given timeframe. To be efficient, neurons maximize their bitrate relative to their energy consumption (Niven & Laughlin, 2008; Sengupta & Stemmler, 2014; Yu & Yu, 2017). Zooming out, the architecture of macro-scale neuronal networks also follows the principle of efficiency: On this level, connections between distant neurons are the major cost factor, leading to an optimization of *wiring cost* (Bullmore & Sporns, 2012). Efficiency on both

levels, from single neurons to brain networks, is crucial given the highly integrated nature of brain function (Pessoa, 2023).

The language of information theory can be applied to the cognitive level in the context of active inference, the guiding principle of perception and action under the FEP (Friston et al., 2016). Active inference is a variant of Bayesian inference, but since it explicitly includes the notion of efficiency, I will focus on this special formulation here. Under active inference, the amount of information in a sensory event is relative to our prior beliefs. A central concept is *surprise*, which is maximal when we observe an event that we believed to be very unlikely. Surprising events have a high bitrate, since they carry information that was, by definition, not accurately accounted for by our beliefs (Baldi & Itti, 2010). While surprise is specific to a single event (or state), *entropy* extends this notion to a system (or process) that can be described by a range of states. Entropy is maximal if the predictability of the current (or future) states is minimal. In terms of probabilities, the different states of a high-entropy process are all equally likely. Crucially, under active inference, entropy is both an attribute of the physical world and our internal model of it. With respect to the latter, high entropy means holding maximally unspecific beliefs regarding our observations. As suggested by Barlow (1961), active inference posits that humans reduce the redundancy of sensory information by only processing aspects of high surprise. Events correctly predicted by the internal model are already available to the mind and don't need to be processed in their entirety.

However, surprising information does not lead to model updating at all costs. In line with concerns that full Bayesian inference might not be realistically implemented in the brain (Koblinger et al., 2021), an approximation is calculated. Through an iterative process, the model converges to the simplest set of beliefs that can still account for the evidence. Changes to the model are only made if the gain in accuracy outweighs the cost of changing the model. This also shields the model against overfitting on the basis of rare events. The maximum entropy principle (Jaynes, 1957) applies here: In the absence of sufficiently strong evidence, the internal model stays maximally agnostic with respect to its beliefs.

4.2 Bayesian inference and imaging markers of predictive processing

These concepts can directly be applied to my design. Predictable sequences have minimal entropy since the transition between two objects is deterministic. During the participants' inference process, their internal model converges to the generative process. Accordingly, the entropy of the model and the bitrate of sensory input decrease. This is

evidenced by the increased confidence ratings which have been shown to map on the certainty of beliefs (Meyniel, Sigman, et al., 2015; Sanders et al., 2016). For any given predictable object, the trailing object is expected with higher certainty and elicits lower surprise. At the same time, if an object occurs that deviates from specific expectations, it elicits higher surprise. My BOLD results support this line of reasoning: With increasing confidence, surprise decreases for expected objects and increases for deviations. This is reflected by the increasing magnitude of relative PE activity in the visual stream.

However, while these findings map well onto the conceptual understanding of surprise (see also Egner et al., 2010; Richter et al., 2018), an interpretation in terms of the mathematical concept is more complicated. This has been approached by performing the inference process computationally, based on the mathematical rules of Bayesian inference. Calculations of surprise over time can then be correlated with the time course of BOLD or EEG activity. In a recent fMRI study, Bounmy et al. (2023) reported that the activity across multiple visual areas tracks Bayesian surprise. This is in contrast to a similar study that found no effects of surprise in sensory areas (Meyniel & Dehaene, 2017). However, activity in a range of posterior (e.g. intraparietal sulcus) and frontal (e.g. superior frontal gyrus) regions covaried with surprise in both studies. This can be explained by the hierarchical nature of predictive coding in the brain (Alexander & Brown, 2018; Jiang & Rao, 2022). PEs are assumed to be calculated at every level of the cortical hierarchy, corresponding to different levels of complexity. In my BOLD data, these effects were present at the uncorrected level (see Project I, Figure 3c) but did not survive cluster correction. Contrasting predictable and unpredictable sequences resulted in clusters which are in line with previous studies. Meyniel and Dehaene (2017) found that highly similar bilateral clusters in the parietal cortex covary with the extent of confidence-weighted updating. This is supported by my behavioral data, according to which participants were still learning the transitional probabilities when they entered the scanner. Finally, it should be noted that dissociations in localization are not necessarily a sign of competing findings. A recent meta-study found that predictive processing is subserved by a brain-spanning network (Ficco et al., 2021). Variations in experimental design and statistical method can lead to different, equally valid clusters surviving the analysis (Gonzalez-Castillo et al., 2012).

4.3 Precision-weighting and attention

Whether attention is a confounder, or an aspect of predictive processing is still debated in the literature. While the present study controlled for spatial attention by use of a cover task,

the differences in reaction time point to varying levels of attentional resources across conditions. Some authors maintain that attention is a separable phenomenon from prediction that only rests on task-relevancy and the salience of external stimuli (Alink & Blank, 2021). The authors argue that both attention-based and prediction-based processes could result in the same experimental observations. However, attention and prediction are not necessarily competing explanations.

Formulations of predictive coding under the FEP posit that attention reflects the uncertainty of perceptual beliefs (Feldman & Friston, 2010). If predictions regarding current or future states of the world are imprecise, this uncertainty might be resolved by new sensory information. This leads to an increased allocation of resources towards sensory input (framed as an increase in the precision of sensory evidence). Consequently, attention boosts the processing of sensory input because our internal model aims for the reduction of uncertainty. However, a range of studies show that the detection of expected stimuli is quicker and more accurate (de Lange et al., 2018). Accepting the previous line of reasoning, this can't follow from uncertain beliefs. However, in the same way that sensory evidence is subject to precision-weighting, so are prior beliefs. Instead of resulting in facilitated processing of sensory input, this leads to a stronger effect of predictions. This effect can be observed in studies providing prior information for inherently ambiguous stimuli. Various studies used so-called two-tone ("Mooney") images, which are greyscale images that have been binarized so that only black and white pixels remain. While they are initially difficult to categorize, participants become significantly quicker and more accurate after a familiarization phase with the original images (Dolan et al., 1997; Teufel et al., 2018). As the sensory input remained the same, the changes in reaction time likely stem from changes to prior beliefs. Although speculative, the two described sources of reaction time decreases could both stem from a reduced complexity of belief updating. When either prior beliefs or sensory evidence has high precision, and the other factor low precision, little competition arises between expectations and input. This might allow for quicker reactions.

4.4 The efficiency of perception based on internal models

The surprise-focused account of perception provides a convincing explanation for lower brain activity during expected stimulation. However, if perception is not driven by sensory evidence in these cases, it must be driven by prior knowledge. Predictable information generally facilitates perception (de Lange et al., 2018) meaning that the same (if not more)

information is available to the brain than sensory input alone provides. The anatomical implementation of predictive coding rests on separate prediction and prediction-error encoding neuronal populations (Bastos et al., 2012). Consequently, neural activity at each level of the hierarchy is needed during the perception of sensory input, irrespective of the weighting of prior versus evidence. Why would perception driven by prediction be more energetically efficient?

The straightforward explanation is that the cost of perception is simply the addition of the cost of prediction and the cost of surprise (and subsequent model updating) (for a related argument see Press et al., 2020). However, there is evidence that representing predictable information is in itself cost-efficient. From an information theoretic standpoint, processes of lower entropy allow for data compression. Assuming that the amount of bits a given neuronal population can represent is limited, a finite number of short codes (e.g. 010 in binary) are available. If frequently encountered (predictable) stimuli are represented with short codes and long codes are reserved for rare stimuli, then any predictable environment will trigger short codes more often. The dominance of short codes in the brain has been studied under the term "sparse coding" (Olshausen & Field, 2004). Chariker et al. (2016) showed that the relay of visual information in the thalamus is based on a surprisingly low number of neurons, highly suggestive of a sparse code for visual perception. Despite possibly reduced energetic cost, sparsity comes with downsides: It has been argued that they impede generalization (Spanne & Jörntell, 2015) and limit the accuracy of predictions in time (Chalk et al., 2018). Sparsity can thus not be the only means of efficient coding for intelligent lifeforms.

In higher stages of processing, a different mechanism for efficient codes has been suggested: Dimensionality reduction. While dedicated statistical methods exist, I here use this term for the representation of multivariate information with fewer variables based on the underlying similarities. A simple case is chunking (Cowan, 2012), a psychological phenomenon where stimuli that occur together in space or time are held in working memory as a single entity. Brady, Konkle & Alvarez (2009) specifically argued that this includes statistical regularities. With respect to my design, this would suggest that confident participants represented predictable sequences as a single chunk, necessitating fewer bits for their representation. Related neuroimaging evidence comes from studies showing that stimuli co-occurring in time or space converge in their representation (Pudhiyidath et al., 2022; Schapiro et al., 2012). Similar to the drawbacks of sparse coding, using representations of lower

complexity is not always beneficial. If differences between similar stimuli are task-relevant, more specific ("sharpened") representations might be preferable (Kok, Jehee, et al., 2012), promoting higher-dimensional representations (Tang et al., 2019). A methodological caveat applies here: Representations as measured by fMRI refer to BOLD activation patterns on the level of voxels. Interpreting these signals as a neural code is problematic: Apart from the dissociations between the BOLD signal and neural activity, a voxel includes hundreds of thousands of neurons (see Kriegeskorte & Kievit, 2013 for a discussion of these issues).

Finally, research on prediction and efficient representations can be unified within the cognitive map framework (Behrens et al., 2018; Bellmund et al., 2018). Cognitive maps are a generalization of spatial maps that are used for navigation in physical space by animals and humans alike. Spatial maps have a well-established neuronal foundation focused on hippocampal cells. Physical space is represented by *grid cells*, who fire according to a consistent interval of distance travelled. Locations in space are represented by *place cells*, whose activity corresponds to specific points in space. It has since been argued that the brain represents relations between concepts in the same way it treats relations between places in physical space: As a map in abstract (cognitive) space. These maps may afford relational learning, generalization and prediction all with the same underlying neural architecture (Whittington et al., 2020). Supporting the applicability of this concept to my design, Stiso et al., (2022) found evidence of cognitive map formation during sequence learning. It should be noted that while predictive coding and cognitive maps use similar conceptual language, there is little overlap in research (but see Stachenfeld et al., 2017). Recently, predictive coding and cognitive maps have been contrasted as models of early visual cortex activity (Linton, 2021). However, computational work suggests that predictive coding algorithms can serve as a means efficient cognitive map construction (Gornet & Thomson, 2023).

In summary, the efficiency of representing environments of low entropy can be expressed by information theory. In the brain, sparse neural codes, low-dimensional representations, and cognitive spaces might be means of implementing efficient function. Predictive coding provides a neural architecture and algorithmic motif that aims to minimize associated neuronal activity. On a cognitive and behavioral level, this results in efficient Bayesian inference as suggested by active inference. This neuro-cognitive architecture could ultimately serve the minimization of energy expenditure, which is favored by evolution. While

intriguing, future studies are needed to link all these concepts. The present work provides a first piece of evidence in line with this claim.

4.5 From biological to artificial intelligence

Research on the efficiency of human perception has implications beyond the understanding of human intelligence. Recent large language models (LLMs) seem to mimic human intelligence with respect to written language and have impacted science and the general public alike (see Clusmann et al., 2023; Demszky et al., 2023 for perspectives in psychology and medicine). A review of the similarities between artificial and biological intelligence is beyond the scope of this work (but see Xu & Poo, 2023) and here I will focus on the aspect of efficiency and energy consumption. Strikingly, estimations of electrical energy consumption of LLMs show a power draw on the order of megawatt hours (training phase) to kilowatt hours (application phase) (Luccioni et al., 2022, 2023). In comparison, ATP consumption in the brain amounts to the equivalent of only 17 watts (Levy & Calvert, 2021).

As predictive coding is a general algorithmic motif, it can be implemented in machine learning models. In this context, it has been discussed as a biologically realistic and potentially more efficient alternative to the dominant backpropagation algorithm (Salvatori et al., 2023; Zahid et al., 2023). While direct comparisons are not available, a recent study provides evidence for a link between predictive processing and energy consumption. Instead of testing whether predictive coding results in energy efficiency, Ali et al., (2022) asked what kind of architecture a neural network develops if its energy consumption is constrained. To this end, the authors trained a recurrent neural network to predict sequences of numbers. Here, a crucial similarity between artificial and biological neural networks was exploited: Both are comprised of layers of neurons that each receive and send signals from other layers. As the input to neurons (postsynaptic activations) is the main cost factor in the brain, their network was optimized to minimize its input weights. The results confirmed that a network trained to minimize input uses lower average activity than a classical network. Strikingly, when the authors examined the activity patterns of individual network neurons, they found that some predicted upcoming pixels while others were sensitive to deviations from these predictions. As this architecture evolved without any explicit implementation, this suggests that predictive coding is a consequence of limited energetic resources. While the differences between artificial and biological networks are only just being explored (see e.g. Lillicrap et al., 2020),

this shows that the coding principles and energetics constraints of the brain can be used to improve the efficiency of machine learning applications.

4.6 Limitations

In isolating the aspect of predictability, the present design controlled for a range of cognitive and behavioral variables. Many of these are crucial aspects of natural human behavior, like movement, action, curiosity, and self-guided attention. This heavily reduces the complexity of natural experience. For this reason, overly controlled experimental designs have been criticized as it is difficult to ascertain whether the resulting patterns are ecologically valid (reviewed in Sonkusare et al., 2019). Furthermore, the context of a MR scan is highly unnatural, since it heavily restricts movement and exposes the participants to constant noise. Recent work in mice suggests that complex, natural behavior elicits equally complex and distributed activity in the brain that is not captured by controlled designs (Benson et al., 2023). Under active inference, perception is only one part of the way humans interact with their environment: During pure perception, optimization of the internal model is limited to the extent of updating. Action, however, endows us with the alternative to choose (or change) our sensory input instead. Consequently, my results are a partial account of the effect of predictive processing on energy efficiency.

For the present projects, I ensured that measured brain activity exclusively results from the perception of single objects and internal prediction. To that end, stimuli were only presented in the center of the screen and the cover task necessitated constant fixation and attention. In the natural world however, objects are nearly always in context, which has been shown to elicit predictive processing (Bar, 2004). While my stimuli were more naturalistic than traditionally used gratings or simple shapes, they still limited the visual stream to relatively low complexity. In recent years, movies have been used in research to approximate the complexity of natural experience. Movies elicit complex viewing behavior that varies between closer examination of details and exploration of novel aspects (Ramos Gameiro et al., 2017). This is akin to the action aspect of active inference. Currently, whole-brain metabolic imaging methods do not provide the temporal resolution to account for changes in sensory input based on self-guided attention on a second or even millisecond level.

A range of studies suggests that predictability and structure play a crucial role in free movie viewing. Hasson et al. (2008) reported a gradient of temporal windows in the brain based on BOLD signals – while sensory cortices were sensitive to input on a frame-by-frame

basis, higher cognitive networks represented accumulated information over scenes. The authors also explored the importance of predictability by showing movie scenes in reverse. While sensory areas showed high similarity in activity between forward and backward scenes, higher cognitive areas differed more strongly in their response. A recent study by Lee, Aly and Baldassano (2021) confirmed the hierarchy of temporal windows and explicitly linked movie viewing to anticipation. The authors found that fMRI patterns corresponding to specific scenes shifted to a timeframe before these scenes started on repeated viewing. While these findings are in line with the hierarchical nature of predictive coding, they suggest that the full range of predictive processing can only be elicited by complex input.

Regarding the employed imaging methods, mqBOLD improves upon BOLD in terms of quantification and biological validity. At the same time, the method comes with some unique drawbacks. As noted before, the temporal resolution is limited to blocks of multiple minutes, precluding research on highly dynamic processes. However, this is largely a result of the signal to noise ratio of the perfusion sequences, which might be improved in future versions. Furthermore, measures of oxygen extraction should ideally only depend on venous blood since this is indicative of the amount of consumed oxygen (Hua et al., 2019). The presently used method cannot differentiate between venous and arterial CBV, resulting in a mixed signal that can bias CMR_{O2} estimates. Finally, while oxidative energy metabolism accounts for up to 98% of ATP production in the visual cortex (Lin et al., 2010), the ratio between oxidative and non-oxidative metabolism is an active area of research. Consequently, a holistic measurement of energy consumption should include both glucose and oxygen consumption.

5 Conclusion and future directions

Efficient function is a crucial aspect of any system with limited resource availability. Without an accurate sense of vision, our ability to survive and make purposeful decisions would be heavily compromised. At the same time, the incurred energetic cost must not exceed the advantages gained by our visual system. In the present work, I examined whether this balance can be struck by reducing the redundancy of visual information. Building on previous findings showing that expected visual input elicits weaker brain activity, I measured oxygen consumption in the brain under varying levels of stimulus predictability. The results show that the impact of predictability depends on the confidence we associate with our expectations. Energy metabolic cost was minimized when stimulation was deterministic and predictions

highly confident. I found that this decrease was present throughout most of the cortex, leading to a system-wide increase in energetic efficiency. This is in line with recent formulations of the brain as a highly integrated inference system and provides the first evidence of the large-scale energetic efficiency of precise expectations.

My results suggest multiple avenues for future research. First, the precision and accuracy of predictions could be manipulated independently. This would help understanding to which extent the efficiency of visual processing rests on objective accuracy. Second, studies on the mechanisms behind efficient coding would benefit from metabolic imaging data. An understanding of the contribution of different coding mechanisms has implications for the development of more environmentally friendly AI. Lastly, metabolic efficiency should be studied under more naturalistic conditions. Combining new developments in perfusion imaging and functional PET with free movie viewing or video game playing would inform us about the interaction of efficient processing and free attention allocation.

References

Acerbi, L., Vijayakumar, S., & Wolpert, D. M. (2014). On the Origins of Suboptimality in Human Probabilistic Inference. *PLoS Computational Biology*, *10*(6), e1003661. https://doi.org/10.1371/journal.pcbi.1003661

Adler, W. T., & Ma, W. J. (2018). Comparing Bayesian and non-Bayesian accounts of human confidence reports. *PLOS Computational Biology*, *14*(11), e1006572. https://doi.org/10.1371/journal.pcbi.1006572

Ais, J., Zylberberg, A., Barttfeld, P., & Sigman, M. (2016). Individual consistency in the accuracy and distribution of confidence judgments. *Cognition*, *146*, 377–386. https://doi.org/10.1016/j.cognition.2015.10.006

Aizenstein, H. J. (2004). Regional Brain Activation during Concurrent Implicit and Explicit Sequence Learning. *Cerebral Cortex*, *14*(2), 199–208. https://doi.org/10.1093/cercor/bhg119

Alamia, A., Orban de Xivry, J.-J., San Anton, E., Olivier, E., Cleeremans, A., & Zenon, A. (2016). Unconscious associative learning with conscious cues. *Neuroscience of Consciousness*, *2016*(1). https://doi.org/10.1093/nc/niw016

Alamia, A., & Zénon, A. (2016). Statistical Regularities Attract Attention when Task-Relevant. *Frontiers in Human Neuroscience*, *10*. https://doi.org/10.3389/fnhum.2016.00042

Alexander, W. H., & Brown, J. W. (2018). Frontal cortex function as derived from hierarchical predictive coding. *Scientific Reports*, *8*(1). https://doi.org/10.1038/s41598-018-21407-9

Ali, A., Ahmad, N., De Groot, E., Johannes Van Gerven, M. A., & Kietzmann, T. C. (2022). Predictive coding is a consequence of energy efficiency in recurrent neural networks. *Patterns*, *3*(12), 100639. https://doi.org/10.1016/j.patter.2022.100639

Alink, A., & Blank, H. (2021). Can expectation suppression be explained by reduced attention to predictable stimuli? *NeuroImage*, *231*, 117824. https://doi.org/10.1016/j.neuroimage.2021.117824

Alink, A., Schwiedrzik, C. M., Kohler, A., Singer, W., & Muckli, L. (2010). Stimulus Predictability Reduces Responses in Primary Visual Cortex. *Journal of Neuroscience*, *30*(8), 2960–2966. https://doi.org/10.1523/JNEUROSCI.3730-10.2010

Ances, B. M., Leontiev, O., Perthen, J. E., Liang, C., Lansing, A. E., & Buxton, R. B. (2008). Regional differences in the coupling of cerebral blood flow and oxygen metabolism

changes in response to activation: Implications for BOLD-fMRI. *NeuroImage*, *39*(4), 1510–1521. https://doi.org/10.1016/j.neuroimage.2007.11.015

Baldi, P., & Itti, L. (2010). Of bits and wows: A Bayesian theory of surprise with applications to attention. *Neural Networks*, *23*(5), 649–666. https://doi.org/10.1016/j.neunet.2009.12.007

Bang, D., & Fleming, S. M. (2018). Distinct encoding of decision confidence in human medial prefrontal cortex. *Proceedings of the National Academy of Sciences*, *115*(23), 6082–6087. https://doi.org/10.1073/pnas.1800795115

Bar, M. (2004). Visual objects in context. *Nature Reviews Neuroscience*, *5*(8), 617–629. https://doi.org/10.1038/nrn1476

Barlow, H. (1961). Possible Principles Underlying the Transformations of Sensory Messages. In *Sensory Communication* (S. 217–234). MIT Press.

Barlow, H. (1990). Conditions for versatile learning, Helmholtz's unconscious inference, and the task of perception. *Vision Research*, *30*(11), 1561–1571. https://doi.org/10.1016/0042-6989(90)90144-A

Bastos, A. M., Lundqvist, M., Waite, A. S., Kopell, N., & Miller, E. K. (2020). Layer and rhythm specificity for predictive routing. *Proceedings of the National Academy of Sciences*, *117*(49), 31459–31469. https://doi.org/10.1073/pnas.2014868117

Bastos, A. M., Usrey, W. M., Adams, R. A., Mangun, G. R., Fries, P., & Friston, K. (2012). Canonical Microcircuits for Predictive Coding. *Neuron*, *76*(4), 695–711. https://doi.org/10.1016/j.neuron.2012.10.038

Batterink, L. J., Paller, K. A., & Reber, P. J. (2019). Understanding the Neural Bases of Implicit and Statistical Learning. *Topics in Cognitive Science*, *11*(3), 482–503. https://doi.org/10.1111/tops.12420

Batterink, L. J., Reber, P. J., Neville, H. J., & Paller, K. A. (2015). Implicit and explicit contributions to statistical learning. *Journal of Memory and Language*, *83*, 62–78. https://doi.org/10.1016/j.jml.2015.04.004

Beck, J. M., Ma, W. J., Pitkow, X., Latham, P. E., & Pouget, A. (2012). Not Noisy, Just Wrong: The Role of Suboptimal Inference in Behavioral Variability. *Neuron*, *74*(1), 30–39. https://doi.org/10.1016/j.neuron.2012.03.016

Behrens, T. E. J., Muller, T. H., Whittington, J. C. R., Mark, S., Baram, A. B., Stachenfeld, K. L., & Kurth-Nelson, Z. (2018). What Is a Cognitive Map? Organizing Knowledge for

Flexible Behavior. *Neuron*, *100*(2), 490–509. https://doi.org/10.1016/j.neuron.2018.10.002

Bellmund, J. L. S., Gärdenfors, P., Moser, E. I., & Doeller, C. F. (2018). Navigating cognition: Spatial codes for human thinking. *Science*, *362*(6415), eaat6766. https://doi.org/10.1126/science.aat6766

Benson, B., Benson, J., Birman, D., Bonacchi, N., Carandini, M., Catarino, J. A., Chapuis, G. A., Churchland, A. K., Dan, Y., Dayan, P., DeWitt, E. E., Engel, T. A., Fabbri, M., Faulkner, M., Fiete, I. R., Findling, C., Freitas-Silva, L., Gerçek, B., Harris, K. D., … Witten, I. B. (2023). *A Brain-Wide Map of Neural Activity during Complex Behaviour* [Preprint]. Neuroscience. https://doi.org/10.1101/2023.07.04.547681

Boundy-Singer, Z. M., Ziemba, C. M., & Goris, R. L. T. (2022). Confidence reflects a noisy decision reliability estimate. *Nature Human Behaviour*, *7*(1), 142–154. https://doi.org/10.1038/s41562-022-01464-x

Bounmy, T., Eger, E., & Meyniel, F. (2023). A characterization of the neural representation of confidence during probabilistic learning. *NeuroImage*, *268*, 119849. https://doi.org/10.1016/j.neuroimage.2022.119849

Brady, T., Konkle, T., & Alvarez, G. A. (2009). Compression in visual working memory: Using statistical regularities to form more efficient memory representations. *Journal of Experimental Psychology: General*, *138*(4), 487–502. https://doi.org/10.1037/a0016797

Bullmore, E., & Sporns, O. (2012). The economy of brain network organization. *Nature Reviews Neuroscience*, *13*(5), 336–349. https://doi.org/10.1038/nrn3214

Buxton, R. B., Griffeth, V. E. M., Simon, A. B., & Moradi, F. (2014). Variability of the coupling of blood flow and oxygen metabolism responses in the brain: A problem for interpreting BOLD studies but potentially a new window on the underlying neural activity. *Frontiers in Neuroscience*, *8*. https://doi.org/10.3389/fnins.2014.00139

Buxton, R. B., Uludağ, K., Dubowitz, D. J., & Liu, T. T. (2004). Modeling the hemodynamic response to brain activation. *NeuroImage*, *23*, S220–S233. https://doi.org/10.1016/j.neuroimage.2004.07.013

Chalk, M., Marre, O., & Tkačik, G. (2018). Toward a unified theory of efficient, predictive, and sparse coding. *Proceedings of the National Academy of Sciences*, *115*(1), 186–191. https://doi.org/10.1073/pnas.1711114115

Chariker, L., Shapley, R., & Young, L.-S. (2016). Orientation Selectivity from Very Sparse LGN Inputs in a Comprehensive Model of Macaque V1 Cortex. *The Journal of Neuroscience*, *36*(49), 12368–12384. https://doi.org/10.1523/JNEUROSCI.2603-16.2016

Christen, T., Bolar, D. S., & Zaharchuk, G. (2013). Imaging Brain Oxygenation with MRI Using Blood Oxygenation Approaches: Methods, Validation, and Clinical Applications. *American Journal of Neuroradiology*, *34*(6), 1113–1123. https://doi.org/10.3174/ajnr.A3070

Christen, T., Schmiedeskamp, H., Straka, M., Bammer, R., & Zaharchuk, G. (2012). Measuring brain oxygenation in humans using a multiparametric quantitative blood oxygenation level dependent MRI approach. *Magnetic Resonance in Medicine*, *68*(3), 905–911. https://doi.org/10.1002/mrm.23283

Clark, A. (2013). Whatever next? Predictive brains, situated agents, and the future of cognitive science. *Behavioral and Brain Sciences*, *36*(3), 181–204. https://doi.org/10.1017/S0140525X12000477

Clusmann, J., Kolbinger, F. R., Muti, H. S., Carrero, Z. I., Eckardt, J.-N., Laleh, N. G., Löffler, C. M. L., Schwarzkopf, S.-C., Unger, M., Veldhuizen, G. P., Wagner, S. J., & Kather, J. N. (2023). The future landscape of large language models in medicine. *Communications Medicine*, *3*(1), 141. https://doi.org/10.1038/s43856-023-00370-1

Colombo, M., Elkin, L., & Hartmann, S. (2021). Being Realist about Bayes, and the Predictive Processing Theory of Mind. *The British Journal for the Philosophy of Science*, *72*(1), 185–220. https://doi.org/10.1093/bjps/axy059

Conway, C. M. (2020). How does the brain learn environmental structure? Ten core principles for understanding the neurocognitive mechanisms of statistical learning. *Neuroscience & Biobehavioral Reviews*, *112*, 279–299. https://doi.org/10.1016/j.neubiorev.2020.01.032

Cowan, N. (2012). *Working Memory Capacity* (0 Aufl.). Psychology Press. https://doi.org/10.4324/9780203342398

Dale, R., Duran, N., & Morehead, R. (2012). Prediction during statistical learning, and implications for the implicit/explicit divide. *Advances in Cognitive Psychology*, *8*(2), 196–209. https://doi.org/10.5709/acp-0115-z

Davis, B., & Hasson, U. (2018). Predictability of what or where reduces brain activity, but a bottleneck occurs when both are predictable. *NeuroImage*, *167*, 224–236. https://doi.org/10.1016/j.neuroimage.2016.06.001

Dayan, P., Hinton, G. E., Neal, R. M., & Zemel, R. S. (1995). The Helmholtz Machine. *Neural Computation*, *7*(5), 889–904. https://doi.org/10.1162/neco.1995.7.5.889

de Lange, F. P., Heilbron, M., & Kok, P. (2018). How Do Expectations Shape Perception? *Trends in Cognitive Sciences*, *22*(9), 764–779. https://doi.org/10.1016/j.tics.2018.06.002

Demszky, D., Yang, D., Yeager, D. S., Bryan, C. J., Clapper, M., Chandhok, S., Eichstaedt, J. C., Hecht, C., Jamieson, J., Johnson, M., Jones, M., Krettek-Cobb, D., Lai, L., JonesMitchell, N., Ong, D. C., Dweck, C. S., Gross, J. J., & Pennebaker, J. W. (2023). Using large language models in psychology. *Nature Reviews Psychology*, *2*(11), 688–701. https://doi.org/10.1038/s44159-023-00241-5

Destrebecqz, A., & Cleeremans, A. (2001). Can sequence learning be implicit? New evidence with the process dissociation procedure. *Psychonomic Bulletin & Review*, *8*(2), 343–350. https://doi.org/10.3758/BF03196171

Devonshire, I. M., Papadakis, N. G., Port, M., Berwick, J., Kennerley, A. J., Mayhew, J. E. W., & Overton, P. G. (2012). Neurovascular coupling is brain region-dependent. *NeuroImage*, *59*(3), 1997–2006. https://doi.org/10.1016/j.neuroimage.2011.09.050

Dolan, R. J., Fink, G. R., Rolls, E., Booth, M., Holmes, A., Frackowiak, R. S. J., & Friston, K. (1997). How the brain learns to see objects and faces in an impoverished context. *Nature*, *389*(6651), 596–599. https://doi.org/10.1038/39309

Drew, P. J. (2019). Vascular and neural basis of the BOLD signal. *Current Opinion in Neurobiology*, *58*, 61–69. https://doi.org/10.1016/j.conb.2019.06.004

Egner, T., Monti, J. M., & Summerfield, C. (2010). Expectation and Surprise Determine Neural Population Responses in the Ventral Visual Stream. *Journal of Neuroscience*, *30*(49), 16601–16608. https://doi.org/10.1523/JNEUROSCI.2770-10.2010

Ekman, M., Kok, P., & de Lange, F. P. (2017). Time-compressed preplay of anticipated events in human primary visual cortex. *Nature Communications*, *8*(1). https://doi.org/10.1038/ncomms15276

Eltzschig, H. K., & Carmeliet, P. (2011). Hypoxia and Inflammation. *New England Journal of Medicine*, *364*(7), 656–665. https://doi.org/10.1056/NEJMra0910283

Fan, H., Burke, T., Sambrano, D. C., Dial, E., Phelps, E. A., & Gershman, S. J. (2023). Pupil Size Encodes Uncertainty during Exploration. *Journal of Cognitive Neuroscience*, *35*(9), 1508–1520. https://doi.org/10.1162/jocn_a_02025

Feldman, H., & Friston, K. J. (2010). Attention, Uncertainty, and Free-Energy. *Frontiers in Human Neuroscience*, *4*. https://doi.org/10.3389/fnhum.2010.00215

Feuerriegel, D., Vogels, R., & Kovács, G. (2021). Evaluating the evidence for expectation suppression in the visual system. *Neuroscience & Biobehavioral Reviews*, *126*, 368–381. https://doi.org/10.1016/j.neubiorev.2021.04.002

Ficco, L., Mancuso, L., Manuello, J., Teneggi, A., Liloia, D., Duca, S., Costa, T., Kovacs, G. Z., & Cauda, F. (2021). Disentangling predictive processing in the brain: A meta-analytic study in favour of a predictive network. *Scientific Reports*, *11*(1), 16258. https://doi.org/10.1038/s41598-021-95603-5

Fick, A. (1870). Ueber die Messung des Blutquantums in den Herzventrikeln. *Sitzungsberichte der Physikalisch-Medizinischen Gesellschaft zu Würzburg*, 16.

Fiser, J., & Aslin, R. N. (2002). Statistical learning of higher-order temporal structure from visual shape sequences. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *28*(3), 458–467. https://doi.org/10.1037//0278-7393.28.3.458

Fiser, J., & Lengyel, G. (2022). Statistical Learning in Vision. *Annual Review of Vision Science*, *8*(1), 265–290. https://doi.org/10.1146/annurev-vision-100720-103343

Fleming, S. M. (2024). Metacognition and Confidence: A Review and Synthesis. *Annual Review of Psychology*, *75*(1), annurev-psych-022423-032425. https://doi.org/10.1146/annurev-psych-022423-032425

Fodor, J. A. (1983). *The modularity of mind: An essay on faculty psychology*. MIT Press.

Fox, P., & Raichle, M. E. (1986). Focal physiological uncoupling of cerebral blood flow and oxidative metabolism during somatosensory stimulation in human subjects. *Proceedings of the National Academy of Sciences*, *83*(4), 1140–1144. https://doi.org/10.1073/pnas.83.4.1140

Fox, P., Raichle, M., Mintun, M., & Dence, C. (1988). Nonoxidative glucose consumption during focal physiologic neural activity. *Science*, *241*(4864), 462–464. https://doi.org/10.1126/science.3260686

Friston, K. (2010). The free-energy principle: A unified brain theory? *Nature Reviews Neuroscience*, *11*(2), 127–138. https://doi.org/10.1038/nrn2787

Friston, K. (2018). Does predictive coding have a future? *Nature Neuroscience*, *21*(8), 1019–1021. https://doi.org/10.1038/s41593-018-0200-7

Friston, K., FitzGerald, T., Rigoli, F., Schwartenbeck, P., O□Doherty, J., & Pezzulo, G. (2016). Active inference and learning. *Neuroscience & Biobehavioral Reviews*, *68*, 862–879. https://doi.org/10.1016/j.neubiorev.2016.06.022

Friston, K. J., Sajid, N., Quiroga-Martinez, D. R., Parr, T., Price, C. J., & Holmes, E. (2021). Active listening. *Hearing Research*, *399*, 107998. https://doi.org/10.1016/j.heares.2020.107998

Friston, K., Moran, R. J., Nagai, Y., Taniguchi, T., Gomi, H., & Tenenbaum, J. (2021). World model learning and inference. *Neural Networks*, *144*, 573–590. https://doi.org/10.1016/j.neunet.2021.09.011

Geurts, L. S., Cooke, J. R. H., Van Bergen, R. S., & Jehee, J. F. M. (2022). Subjective confidence reflects representation of Bayesian probability in cortex. *Nature Human Behaviour*, *6*(2), 294–305. https://doi.org/10.1038/s41562-021-01247-w

Gheysen, F., & Fias, W. (2012). Dissociable neural systems of sequence learning. *Advances in Cognitive Psychology*, *8*(2), 73–82. https://doi.org/10.5709/acp-0105-1

Gonzalez-Castillo, J., Saad, Z. S., Handwerker, D. A., Inati, S. J., Brenowitz, N., & Bandettini, P. A. (2012). Whole-brain, time-locked activation with simple tasks revealed using massive averaging and model-free analysis. *Proceedings of the National Academy of Sciences*, *109*(14), 5487–5492. https://doi.org/10.1073/pnas.1121049109

Gornet, J. A., & Thomson, M. (2023). *Automated construction of cognitive maps with predictive coding* [Preprint]. Neuroscience. https://doi.org/10.1101/2023.09.18.558369

Göttler, J., Kaczmarz, S., Kallmayer, M., Wustrow, I., Eckstein, H.-H., Zimmer, C., Sorg, C., Preibisch, C., & Hyder, F. (2019). Flow-metabolism uncoupling in patients with asymptomatic unilateral carotid artery stenosis assessed by multi-modal magnetic resonance imaging. *Journal of Cerebral Blood Flow & Metabolism*, *39*(11), 2132–2143. https://doi.org/10.1177/0271678X18783369

Hangya, B., Sanders, J. I., & Kepecs, A. (2016). A Mathematical Framework for Statistical Decision Confidence. *Neural Computation*, *28*(9), 1840–1858. https://doi.org/10.1162/NECO_a_00864

Hasson, U., Yang, E., Vallines, I., Heeger, D. J., & Rubin, N. (2008). A Hierarchy of Temporal Receptive Windows in Human Cortex. *Journal of Neuroscience*, *28*(10), 2539–2550. https://doi.org/10.1523/JNEUROSCI.5487-07.2008

Hebart, M. N., Schriever, Y., Donner, T. H., & Haynes, J.-D. (2016). The Relationship between Perceptual Decision Variables and Confidence in the Human Brain. *Cerebral Cortex*, *26*(1), 118–130. https://doi.org/10.1093/cercor/bhu181

Helmholtz, H. von. (1867). *Handbuch der physiologischen Optik*.

Hirsch, N. M., Toth, V., Förschler, A., Kooijman, H., Zimmer, C., & Preibisch, C. (2014). Technical considerations on the validity of blood oxygenation level-dependent-based MR assessment of vascular deoxygenation: BOLD-BASED ASSESSMENT OF VASCULAR DEOXYGENATION. *NMR in Biomedicine*, *27*(7), 853–862. https://doi.org/10.1002/nbm.3131

Hohwy, J. (2012). Attention and Conscious Perception in the Hypothesis Testing Brain. *Frontiers in Psychology*, *3*. https://doi.org/10.3389/fpsyg.2012.00096

Hohwy, J. (2013). *The Predictive Mind*. Oxford University Press. https://doi.org/10.1093/acprof:oso/9780199682737.001.0001

Howarth, C., Gleeson, P., & Attwell, D. (2012). Updated Energy Budgets for Neural Computation in the Neocortex and Cerebellum. *Journal of Cerebral Blood Flow & Metabolism*, *32*(7), 1222–1232. https://doi.org/10.1038/jcbfm.2012.35

Hua, J., Liu, P., Kim, T., Donahue, M., Rane, S., Chen, J. J., Qin, Q., & Kim, S.-G. (2019). MRI techniques to measure arterial and venous cerebral blood volume. *NeuroImage*, *187*, 17–31. https://doi.org/10.1016/j.neuroimage.2018.02.027

Jaynes, E. T. (1957). Information Theory and Statistical Mechanics. *Physical Review*, *106*(4), 620–630. https://doi.org/10.1103/PhysRev.106.620

Jiang, L. P., & Rao, R. P. N. (2022). Predictive Coding Theories of Cortical Function. In L. P. Jiang & R. P. N. Rao, *Oxford Research Encyclopedia of Neuroscience*. Oxford University Press. https://doi.org/10.1093/acrefore/9780190264086.013.328

Kaczmarz, S., Hyder, F., & Preibisch, C. (2020). Oxygen extraction fraction mapping with multi-parametric quantitative BOLD MRI: Reduced transverse relaxation bias using 3D-GraSE imaging. *NeuroImage*, *220*, 117095. https://doi.org/10.1016/j.neuroimage.2020.117095

Kaposvari, P., Kumar, S., & Vogels, R. (2018). Statistical Learning Signals in Macaque Inferior Temporal Cortex. *Cerebral Cortex*, *28*(1), 250–266. https://doi.org/10.1093/cercor/bhw374

Khalvati, K., Kiani, R., & Rao, R. P. N. (2021). Bayesian inference with incomplete knowledge explains perceptual confidence and its deviations from accuracy. *Nature Communications*, *12*(1), 5704. https://doi.org/10.1038/s41467-021-25419-4

Kim, S.-G., & Ogawa, S. (2012). Biophysical and Physiological Origins of Blood Oxygenation Level-Dependent fMRI Signals. *Journal of Cerebral Blood Flow & Metabolism*, *32*(7), 1188–1206. https://doi.org/10.1038/jcbfm.2012.23

Knill, D. C., & Pouget, A. (2004). The Bayesian brain: The role of uncertainty in neural coding and computation. *Trends in Neurosciences*, *27*(12), 712–719. https://doi.org/10.1016/j.tins.2004.10.007

Koblinger, Á., Fiser, J., & Lengyel, M. (2021). Representations of uncertainty: Where art thou? *Current Opinion in Behavioral Sciences*, *38*, 150–162. https://doi.org/10.1016/j.cobeha.2021.03.009

Kogo, N., & Trengove, C. (2015). Is predictive coding theory articulated enough to be testable? *Frontiers in Computational Neuroscience*, *9*. https://doi.org/10.3389/fncom.2015.00111

Kok, P., Bains, L. J., van Mourik, T., Norris, D. G., & de Lange, F. P. (2016). Selective Activation of the Deep Layers of the Human Primary Visual Cortex by Top-Down Feedback. *Current Biology*, *26*(3), 371–376. https://doi.org/10.1016/j.cub.2015.12.038

Kok, P., Failing, M. F., & de Lange, F. P. (2014). Prior Expectations Evoke Stimulus Templates in the Primary Visual Cortex. *Journal of Cognitive Neuroscience*, *26*(7), 1546–1554. https://doi.org/10.1162/jocn_a_00562

Kok, P., Jehee, J. F. M., & de Lange, F. P. (2012). Less Is More: Expectation Sharpens Representations in the Primary Visual Cortex. *Neuron*, *75*(2), 265–270. https://doi.org/10.1016/j.neuron.2012.04.034

Kok, P., Rahnev, D., Jehee, J. F. M., Lau, H. C., & de Lange, F. P. (2012). Attention Reverses the Effect of Prediction in Silencing Sensory Signals. *Cerebral Cortex*, *22*(9), 2197–2206. https://doi.org/10.1093/cercor/bhr310

Kriegeskorte, N., & Kievit, R. A. (2013). Representational geometry: Integrating cognition, computation, and the brain. *Trends in Cognitive Sciences*, *17*(8), 401–412. https://doi.org/10.1016/j.tics.2013.06.007

Lavín, C., San Martín, R., & Rosales Jubal, E. (2014). Pupil dilation signals uncertainty and surprise in a learning gambling task. *Frontiers in Behavioral Neuroscience*, *7*. https://doi.org/10.3389/fnbeh.2013.00218

Lee, C. S., Aly, M., & Baldassano, C. (2021). Anticipation of temporally structured events in the brain. *eLife*, *10*, e64972. https://doi.org/10.7554/eLife.64972

Levy, W. B., & Calvert, V. G. (2021). Communication consumes 35 times more energy than computation in the human cortex, but both costs are needed to predict synapse number. *Proceedings of the National Academy of Sciences*, *118*(18), e2008173118. https://doi.org/10.1073/pnas.2008173118

Lillicrap, T. P., Santoro, A., Marris, L., Akerman, C. J., & Hinton, G. (2020). Backpropagation and the brain. *Nature Reviews Neuroscience*, *21*(6), 335–346. https://doi.org/10.1038/s41583-020-0277-3

Lin, A.-L., Fox, P. T., Hardies, J., Duong, T. Q., & Gao, J.-H. (2010). Nonlinear coupling between cerebral blood flow, oxygen consumption, and ATP production in human visual cortex. *Proceedings of the National Academy of Sciences*, *107*(18), 8446–8451. https://doi.org/10.1073/pnas.0909711107

Linton, P. (2021). V1 as an egocentric cognitive map. *Neuroscience of Consciousness*, *2021*(2), niab017. https://doi.org/10.1093/nc/niab017

Logothetis, N. K., Pauls, J., Augath, M., Trinath, T., & Oeltermann, A. (2001). Neurophysiological investigation of the basis of the fMRI signal. *Nature*, *412*(6843), Article 6843. https://doi.org/10.1038/35084005

Loh, L. K., & Bartulovic, M. (2014). *Efficient Coding Hypothesis and an Introduction to Information Theory*. 8.

Luccioni, A. S., Jernite, Y., & Strubell, E. (2023). *Power Hungry Processing: Watts Driving the Cost of AI Deployment?* (arXiv:2311.16863). arXiv. http://arxiv.org/abs/2311.16863

Luccioni, A. S., Viguier, S., & Ligozat, A.-L. (2022). *Estimating the Carbon Footprint of BLOOM, a 176B Parameter Language Model* (arXiv:2211.02001). arXiv. http://arxiv.org/abs/2211.02001

Lupyan, G. (2015). Cognitive Penetrability of Perception in the Age of Prediction: Predictive Systems are Penetrable Systems. *Review of Philosophy and Psychology*, *6*(4), 547–569. https://doi.org/10.1007/s13164-015-0253-4

Ma, W. J., Beck, J. M., Latham, P. E., & Pouget, A. (2006). Bayesian inference with probabilistic population codes. *Nature Neuroscience*, *9*(11), 1432–1438. https://doi.org/10.1038/nn1790

Maheu, M., Meyniel, F., & Dehaene, S. (2022). Rational arbitration between statistics and rules in human sequence processing. *Nature Human Behaviour*, *6*(8), 1087–1103. https://doi.org/10.1038/s41562-021-01259-6

Meyniel, F., & Dehaene, S. (2017). Brain networks for confidence weighting and hierarchical inference during probabilistic learning. *Proceedings of the National Academy of Sciences*, *114*(19). https://doi.org/10.1073/pnas.1615773114

Meyniel, F., Schlunegger, D., & Dehaene, S. (2015). The Sense of Confidence during Probabilistic Learning: A Normative Account. *PLOS Computational Biology*, *11*(6), e1004305. https://doi.org/10.1371/journal.pcbi.1004305

Meyniel, F., Sigman, M., & Mainen, Z. F. (2015). Confidence as Bayesian Probability: From Neural Origins to Behavior. *Neuron*, *88*(1), 78–92. https://doi.org/10.1016/j.neuron.2015.09.039

Miłkowski, M., & Litwin, P. (2022). Testable or bust: Theoretical lessons for predictive processing. *Synthese*, *200*(6), 462. https://doi.org/10.1007/s11229-022-03891-9

Moradi, F., Buračas, G. T., & Buxton, R. B. (2012). Attention strongly increases oxygen metabolic response to stimulus in primary visual cortex. *NeuroImage*, *59*(1), 601–607. https://doi.org/10.1016/j.neuroimage.2011.07.078

Moradi, F., & Buxton, R. B. (2013). Adaptation of cerebral oxygen metabolism and blood flow and modulation of neurovascular coupling with prolonged stimulation in human visual cortex. *NeuroImage*, *82*, 182–189. https://doi.org/10.1016/j.neuroimage.2013.05.110

Niven, J. E., & Laughlin, S. B. (2008). Energy limitation as a selective pressure on the evolution of sensory systems. *Journal of Experimental Biology*, *211*(11), 1792–1804. https://doi.org/10.1242/jeb.017574

Ogawa, S., Lee, T. M., Kay, A. R., & Tank, D. W. (1990). Brain magnetic resonance imaging with contrast dependent on blood oxygenation. *Proceedings of the National Academy of Sciences*, *87*(24), 9868–9872. https://doi.org/10.1073/pnas.87.24.9868

Olshausen, B., & Field, D. (2004). Sparse coding of sensory inputs. *Current Opinion in Neurobiology*, *14*(4), 481–487. https://doi.org/10.1016/j.conb.2004.07.007

Paulson, O. B., Hasselbalch, S. G., Rostrup, E., Knudsen, G. M., & Pelligrino, D. (2010). Cerebral Blood Flow Response to Functional Activation. *Journal of Cerebral Blood Flow & Metabolism*, *30*(1), 2–14. https://doi.org/10.1038/jcbfm.2009.188

Pessoa, L. (2023). The Entangled Brain. *Journal of Cognitive Neuroscience*, 1–12. https://doi.org/10.1162/jocn_a_01908

Pleskac, T. J., & Busemeyer, J. R. (2010). Two-stage dynamic signal detection: A theory of choice, decision time, and confidence. *Psychological Review*, *117*(3), 864–901. https://doi.org/10.1037/a0019737

Pouget, A., Beck, J. M., Ma, W. J., & Latham, P. E. (2013). Probabilistic brains: Knowns and unknowns. *Nature Neuroscience*, *16*(9), 1170–1178. https://doi.org/10.1038/nn.3495

Pouget, A., Drugowitsch, J., & Kepecs, A. (2016). Confidence and certainty: Distinct probabilistic quantities for different goals. *Nature Neuroscience*, *19*(3), 366–374. https://doi.org/10.1038/nn.4240

Press, C., Kok, P., & Yon, D. (2020). The Perceptual Prediction Paradox. *Trends in Cognitive Sciences*, *24*(1), 13–24. https://doi.org/10.1016/j.tics.2019.11.003

Pudhiyidath, A., Morton, N. W., Viveros Duran, R., Schapiro, A. C., Momennejad, I., Hinojosa-Rowland, D. M., Molitor, R. J., & Preston, A. R. (2022). Representations of Temporal Community Structure in Hippocampus and Precuneus Predict Inductive Reasoning Decisions. *Journal of Cognitive Neuroscience*, *34*(10), 1736–1760. https://doi.org/10.1162/jocn_a_01864

Pylyshyn, Z. (1999). Is vision continuous with cognition?: The case for cognitive impenetrability of visual perception. *Behavioral and Brain Sciences*, *22*(3), 341–365. https://doi.org/10.1017/S0140525X99002022

Quintela-López, T., Shiina, H., & Attwell, D. (2022). Neuronal energy use and brain evolution. *Current Biology*, *32*(12), R650–R655. https://doi.org/10.1016/j.cub.2022.02.005

Rahnev, D. (2019). The Bayesian brain: What is it and do humans have it? *Behavioral and Brain Sciences*, *42*, e238. https://doi.org/10.1017/S0140525X19001377

Ramachandran, S., Meyer, T., & Olson, C. R. (2017). Prediction suppression and surprise enhancement in monkey inferotemporal cortex. *Journal of Neurophysiology*, *118*(1), 374–382. https://doi.org/10.1152/jn.00136.2017

Ramos Gameiro, R., Kaspar, K., König, S. U., Nordholt, S., & König, P. (2017). Exploration and Exploitation in Natural Viewing Behavior. *Scientific Reports*, *7*(1). https://doi.org/10.1038/s41598-017-02526-1

Rao, R. P. N., & Ballard, D. H. (1999). Predictive coding in the visual cortex: A functional interpretation of some extra-classical receptive-field effects. *Nature Neuroscience*, *2*(1), 79–87. https://doi.org/10.1038/4580

Rescorla, M. (2021). Bayesian modeling of the mind: From norms to neurons. *WIREs Cognitive Science*, *12*(1), e1540. https://doi.org/10.1002/wcs.1540

Richter, D., & de Lange, F. P. (2019). Statistical learning attenuates visual activity only for attended stimuli. *eLife*, *8*, e47869. https://doi.org/10.7554/eLife.47869

Richter, D., Ekman, M., & de Lange, F. P. (2018). Suppressed Sensory Response to Predictable Object Stimuli throughout the Ventral Visual Stream. *The Journal of Neuroscience*, *38*(34), 7452–7461. https://doi.org/10.1523/JNEUROSCI.3421-17.2018

Rolfe, D. F., & Brown, G. C. (1997). Cellular energy utilization and molecular origin of standard metabolic rate in mammals. *Physiological Reviews*, *77*(3), 731–758. https://doi.org/10.1152/physrev.1997.77.3.731

Salinas, E., & Abbott, L. F. (1994). Vector reconstruction from firing rates. *Journal of Computational Neuroscience*, *1*(1–2), 89–107. https://doi.org/10.1007/BF00962720

Salvatori, T., Mali, A., Buckley, C. L., Lukasiewicz, T., Rao, R. P. N., Friston, K., & Ororbia, A. (2023). *Brain-Inspired Computational Intelligence via Predictive Coding* (arXiv:2308.07870). arXiv. http://arxiv.org/abs/2308.07870

Sanborn, A. N., & Chater, N. (2016). Bayesian Brains without Probabilities. *Trends in Cognitive Sciences*, *20*(12), 883–893. https://doi.org/10.1016/j.tics.2016.10.003

Sanders, J. I., Hangya, B., & Kepecs, A. (2016). Signatures of a Statistical Computation in the Human Sense of Confidence. *Neuron*, *90*(3), 499–506. https://doi.org/10.1016/j.neuron.2016.03.025

Schapiro, A. C., Kustner, L. V., & Turk-Browne, N. B. (2012). Shaping of Object Representations in the Human Medial Temporal Lobe Based on Temporal

Regularities. *Current Biology*, *22*(17), 1622–1627. https://doi.org/10.1016/j.cub.2012.06.056

Schapiro, A., & Turk-Browne, N. (2015). Statistical Learning. In *Brain Mapping* (S. 501–506). Elsevier. https://doi.org/10.1016/B978-0-12-397025-1.00276-1

Sengupta, B., & Stemmler, M. B. (2014). Power Consumption During Neuronal Computation. *Proceedings of the IEEE*, *102*(5), 738–750. https://doi.org/10.1109/JPROC.2014.2307755

Sengupta, B., Stemmler, M. B., & Friston, K. (2013). Information and Efficiency in the Nervous System—A Synthesis. *PLoS Computational Biology*, *9*(7), e1003157. https://doi.org/10.1371/journal.pcbi.1003157

Shekhar, M., & Rahnev, D. (2021). Sources of Metacognitive Inefficiency. *Trends in Cognitive Sciences*, *25*(1), 12–23. https://doi.org/10.1016/j.tics.2020.10.007

Shobatake, R., Ota, H., Takahashi, N., Ueno, S., Sugie, K., & Takasawa, S. (2022). The Impact of Intermittent Hypoxia on Metabolism and Cognition. *International Journal of Molecular Sciences*, *23*(21), 12957. https://doi.org/10.3390/ijms232112957

Simoncelli, E. P. (2003). Vision and the statistics of the visual environment. *Current Opinion in Neurobiology*, *13*(2), 144–149. https://doi.org/10.1016/S0959-4388(03)00047-3

Sonkusare, S., Breakspear, M., & Guo, C. (2019). Naturalistic Stimuli in Neuroscience: Critically Acclaimed. *Trends in Cognitive Sciences*, *23*(8), 699–714. https://doi.org/10.1016/j.tics.2019.05.004

Spanne, A., & Jörntell, H. (2015). Questioning the role of sparse coding in the brain. *Trends in Neurosciences*, *38*(7), 417–427. https://doi.org/10.1016/j.tins.2015.05.005

Stachenfeld, K. L., Botvinick, M. M., & Gershman, S. J. (2017). The hippocampus as a predictive map. *Nature Neuroscience*, *20*(11), 1643–1653. https://doi.org/10.1038/nn.4650

Stiso, J., Lynn, C. W., Kahn, A. E., Rangarajan, V., Szymula, K. P., Archer, R., Revell, A., Stein, J. M., Litt, B., Davis, K. A., Lucas, T. H., & Bassett, D. S. (2022). Neurophysiological Evidence for Cognitive Map Formation during Sequence Learning. *eneuro*, *9*(2), ENEURO.0361-21.2022. https://doi.org/10.1523/ENEURO.0361-21.2022

Summerfield, C., Egner, T., Greene, M., Koechlin, E., Mangels, J., & Hirsch, J. (2006). Predictive Codes for Forthcoming Perception in the Frontal Cortex. *Science*, *314*(5803), 1311–1314. https://doi.org/10.1126/science.1132028

Tang, E., Mattar, M. G., Giusti, C., Lydon-Staley, D. M., Thompson-Schill, S. L., & Bassett, D. S. (2019). Effective learning is accompanied by high-dimensional and efficient representations of neural activity. *Nature Neuroscience*, *22*(6), 1000–1009. https://doi.org/10.1038/s41593-019-0400-9

Teufel, C., Dakin, S. C., & Fletcher, P. C. (2018). Prior object-knowledge sharpens properties of early visual feature-detectors. *Scientific Reports*, *8*(1). https://doi.org/10.1038/s41598-018-28845-5

Timme, N. M., & Lapish, C. (2018). A Tutorial for Information Theory in Neuroscience. *Eneuro*, *5*(3), ENEURO.0052-18.2018. https://doi.org/10.1523/ENEURO.0052-18.2018

Turk-Browne, N. B., Jungé, J. A., & Scholl, B. J. (2005). The Automaticity of Visual Statistical Learning. *Journal of Experimental Psychology: General*, *134*(4), 552–564. https://doi.org/10.1037/0096-3445.134.4.552

Vadillo, M. A., Konstantinidis, E., & Shanks, D. R. (2016). Underpowered samples, false negatives, and unconscious learning. *Psychonomic Bulletin & Review*, *23*(1), 87–102. https://doi.org/10.3758/s13423-015-0892-6

Van Bergen, R. S., Ji Ma, W., Pratte, M. S., & Jehee, J. F. M. (2015). Sensory uncertainty decoded from visual cortex predicts behavior. *Nature Neuroscience*, *18*(12), 1728–1730. https://doi.org/10.1038/nn.4150

Vincent, P., Parr, T., Benrimoh, D., & Friston, K. J. (2019). With an eye on uncertainty: Modelling pupillary responses to environmental volatility. *PLOS Computational Biology*, *15*(7), e1007126. https://doi.org/10.1371/journal.pcbi.1007126

Walsh, K. S., McGovern, D. P., Clark, A., & O'Connell, R. G. (2020). Evaluating the neurophysiological evidence for predictive processing as a model of perception. *Annals of the New York Academy of Sciences*, *1464*(1), 242–268. https://doi.org/10.1111/nyas.14321

Warren, R. E., & Frier, B. M. (2005). Hypoglycaemia and cognitive function. *Diabetes, Obesity and Metabolism*, *7*(5), 493–503. https://doi.org/10.1111/j.1463-1326.2004.00421.x

Whittington, J. C. R., Muller, T. H., Mark, S., Chen, G., Barry, C., Burgess, N., & Behrens, T. E. J. (2020). The Tolman-Eichenbaum Machine: Unifying Space and Relational Memory through Generalization in the Hippocampal Formation. *Cell*, *183*(5), 1249-1263.e23. https://doi.org/10.1016/j.cell.2020.10.024

Xu, B., & Poo, M. (2023). Large language models and brain-inspired general intelligence. *National Science Review*, *10*(10), nwad267. https://doi.org/10.1093/nsr/nwad267

Xu, F., Ge, Y., & Lu, H. (2009). Noninvasive quantification of whole-brain cerebral metabolic rate of oxygen (CMRO $_2$) by MRI. *Magnetic Resonance in Medicine*, *62*(1), 141–148. https://doi.org/10.1002/mrm.21994

Yon, D., & Frith, C. D. (2021). Precision and the Bayesian brain. *Current Biology*, *31*(17), R1026–R1032. https://doi.org/10.1016/j.cub.2021.07.044

Yon, D., Heyes, C., & Press, C. (2020). Beliefs and desires in the predictive brain. *Nature Communications*, *11*(1), 4404. https://doi.org/10.1038/s41467-020-18332-9

Yu, L., & Yu, Y. (2017). Energy-efficient neural information processing in individual neurons and neuronal networks: Energy Efficiency in Neural Systems. *Journal of Neuroscience Research*, *95*(11), 2253–2266. https://doi.org/10.1002/jnr.24131

Zahid, U., Guo, Q., & Fountas, Z. (2023). Predictive Coding as a Neuromorphic Alternative to Backpropagation: A Critical Evaluation. *Neural Computation*, *35*(12), 1881–1909. https://doi.org/10.1162/neco_a_01620

Zhou, D., Lynn, C. W., Cui, Z., Ciric, R., Baum, G. L., Moore, T. M., Roalf, D. R., Detre, J. A., Gur, R. C., Gur, R. E., Satterthwaite, T. D., & Bassett, D. S. (2022). Efficient coding in the economics of human brain connectomics. *Network Neuroscience*, *6*(1), 234–274. https://doi.org/10.1162/netn_a_00223

DECLARATION OF AUTHOR CONTRIBUTIONS

Project I: Subjective confidence modulates individual BOLD patterns of predictive processing

Hechler A., de Lange F., Riedl V.

AH performed all design implementation, investigation, formal analysis, visualization, and writing. FdL contributed to experimental design conceptualization. VR provided supervision, project conceptualization, funding and writing review.

My contributions to this manuscript in detail

For this project, I designed and programmed both the online and in-person version of the experimental design. I performed all data acquisition, processing, analysis, visualization, and writing.

Project II: The energy metabolic footprint of predictive processing in the human brain

Hechler A., de Lange F., Riedl V.

AH performed all design implementation, investigation, formal analysis, visualization, and writing. FdL contributed to experimental design conceptualization and provided writing commentary. VR provided supervision, project conceptualization, funding and writing review.

My contributions to this manuscript in detail

For this project, I designed and programmed both the online and in-person version of the experimental design. I performed all data acquisition, processing, analysis, visualization, and writing.

München, 10.06.2024

Date André Hechler Prof. Dr. med. Valentin Riedl

AFFIDAVIT / EIDESSTATTLICHE VERSICHERUNG

Hiermit versichere ich an Eides statt, dass ich die vorliegende Dissertation „The energy metabolic footprint of predictive processing in the human brain“ selbstständig angefertigt habe, mich außer der angegebenen keiner weiteren Hilfsmittel bedient und alle Erkenntnisse, die aus dem Schrifttum ganz oder annähernd übernommen sind, als solche kenntlich gemacht und nach ihrer Herkunft unter Bezeichnung der Fundstelle einzeln nachgewiesen habe.

I hereby confirm that the dissertation „The energy metabolic footprint of predictive processing in the human brain“ is the result of my own work and that I have only used sources or materials listed and specified in the dissertation.

München, 10.06.2024

Date/Datum André Hechler