Using serial dependence to predict confidence across observers and cognitive domains

Ning Mei ^a, Dobromir Rahnev ^b, David Soto ^{a,c}

a Basque Center on Cognition, Brain, and Language, San Sebastian, Spain

b School of Psychology, Georgia Institute of Technology, Atlanta, USA

c Ikerbasque, Basque Foundation for Science, Spain

Corresponding authors' email addresses:

Ning Mei: n.mei@bcbl.eu; David Soto: d.soto@bcbl.eu

Abstract

Our perceptual system appears hardwired to exploit regularities of input features across space and time in seemingly stable environments. This can lead to serial dependence effects whereby recent perceptual representations bias current perception. Serial dependence has also been demonstrated for more abstract representations such as perceptual confidence. Here we ask whether temporal patterns in the generation of confidence judgments across trials generalize across observers and different cognitive domains. Data from the Confidence Database across perceptual, memory, and cognitive paradigms was re-analyzed. Machine learning classifiers were used to predict the confidence on the current trial based on the history of confidence judgments on the previous trials. Cross-observer and cross-domain decoding results showed that a model trained to predict confidence in the perceptual domain generalized across observers to predict confidence across the different cognitive domains. The recent history of confidence was the most critical factor. The history of accuracy or type-1 reaction time alone, or in combination with confidence, did not improve the prediction of the current confidence. We also observed that confidence predictions generalized across correct and incorrect trials, indicating that serial dependence effects in confidence generation are uncoupled to metacognition (i.e. how we evaluate the precision of our own behavior). We discuss the ramifications of these findings for the ongoing debate on domain-general vs. specificity of metacognition.

Introduction

Perceptual judgments in the face of uncertainty are accompanied by metacognitive evaluations – a sense of confidence that tracks the probability of one's decision being correct. It remains, however, unclear how confidence is generated on a moment-to-moment basis.

Our perceptual systems appear hardwired to exploit the auto-correlation in perceptual input across space and time, leading to serial dependence effects in perception, such that recent prior perceptual representations bias current perception (Fischer and Whitney 2014). Serial dependence in perception reflects the influence that prior perceptual representations have on ongoing perception and is a widespread phenomenon demonstrated for low-level features such as oriented gratings (Fischer and Whitney 2014) to more complex objects such as faces (Lieberman, Fischer, and Whitney 2014) or ensemble visual representations (Manassi et al. 2017). It has also been argued that a mechanism based on serial dependence can account for different experimental results reported across perceptual, memory, and attention tasks (Kiyonaga et al. 2017).

Recent research also indicates the existence of serial dependence for perceptual confidence (Rahnev et al. 2015); see also (Mueller and Weidemann 2008). Serial dependence in confidence judgments across two different visual tasks has been demonstrated even when confidence is not always rated (Aguilar-Lleyda et al. 2021). The current study focuses on understanding the role of serial dependence in confidence generation beyond perception, across multiple cognitive domains. Here we test the hypothesis that serial dependence effects in confidence generation in perceptual tasks generalize across observers and cognitive domains.

We leverage the power of machine learning classifiers and out-of-sample predictions to investigate how predictable confidence is, whether the moment-to-moment predictability of confidence is task/domain-specific, or whether the patterns of serial dependence during confidence generation in the perceptual domain generalize to predict confidence judgments made by different observers across different cognitive domains. Machine learning classifiers were trained using confidence measures estimates from previous trials to predict the level of confidence in the current trial. The classifier was trained in one particular domain (i.e. perception) and then tested across different cognitive domains (i.e. memory) and generalization performance was quantified.

We predicted that if the dynamic representation of metacognitive confidence across trials is shared across different cognitive domains then the classifiers ought to generalize and hence predict the representation of metacognitive confidence in different observers and domains. Notably, the question addressed here is impossible to address using traditional statistical approaches in which all data is fit at once in the same statistical model. This is one of the advantages of using machine learning to study generalization and out-of-sample predictions regarding human decisions, including confidence ratings. For instance, a recent study showed evidence of serial dependence across domains in a paradigm in which recognition judgments were interleaved with perceptual judgments so that confidence in the perceptual judgment influenced subsequent memory confidence in the next trial (Kantner et al. 2019). However, this approach in which all data is fitted

at once within the statistical model does not enable to quantify and predict confidence in new examples, namely, across different observers and across cognitive domains.

The distribution of confidence judgments in visual tasks is highly stable within a particular observer when performance is measured across different days in the same visual task and also across visual tasks of similar structure (Ais et al. 2016). Such findings suggest that serial dependence may also be stable within individual observers when similar visual perceptual tasks are considered but raise questions on whether it is also similar across observers. Here we addressed whether the patterns of confidence across trials associated with serial dependence effects were similar across observers, by training and testing the classifiers in different observers. We also tested whether these confidence patterns are generalizable across correct and incorrect response trials (see Methods). A recent meta-analysis by (Rouault et al. 2018) found no behavioral evidence for the association of inter-individual measures of metacognitive performance (i.e. how confidence tracks accuracy) across perception and memory domains. This observation predicts that the pattern of moment-to-moment estimation of confidence might be different and not generalizable across cognitive domains.

Methods

This study re-analyzed data from the Confidence Database (Rahnev et al. 2020). We selected all datasets that employed a 4-point confidence scale because this scale was most common in the perceptual domain thereby providing a rich dataset to train the classifier and assess the degree of generalization across other domains (i.e., cognitive, memory, and mixed). This resulted in a total of 32 datasets: 16 from the perceptual domain, four from the cognitive domain, six from the memory domain, and six from a mixed domain (where the data came from tasks from multiple domains). Details about each dataset – including the name of the dataset, number of subjects, and total number of trials – are included in Supplementary Table 1. In total, there were 3,077 subjects and 925,091 trials across all 32 datasets. The data were downloaded in April 2020 and may not reflect datasets added to the Confidence Database after that date.

Machine learning models

We used machine learning models: a random forest (RF) classifier and a linear vector machine (SVM) as a baseline model (see Table 1).

We chose the linear SVM due to its simplicity and its ability to perform a large number of inferences during prediction, thus, the linear SVM served as a baseline classification model. We chose the random forest classifier to perform the 4-classes classification task because the random forest classifier allowed us to compute the feature importance and make interpretable relationships between the features and the target. This was the same strategy as in our prior study (Mei et al., 2021).

Model name	Random Forest classifier	Linear SVM
Advantages	<ol style="list-style-type: none"> 1. An ensemble model 2. One of the best models for multiclass classification 3. Interpretation of the feature importance is straightforward 4. Overfitting can be regularized by increasing the number ensemble tree-based models 	<ol style="list-style-type: none"> 1. A simple linear model 2. Powerful for scaling up the inference of a trained model 3. Interpretation of the feature weights is simple and directional 4. Overfitting can be regularized by L2 penalty
Disadvantages	<ol style="list-style-type: none"> 1. It is difficult to control for the complexity of the ensemble treed-based models due to the large number of hyperparameters, such as the depth, resample sizes for training the individual tree-based models 2. Training and inference time is much longer than SVM for big datasets 	<ol style="list-style-type: none"> 1. Not suitable for multiclass classification 2. Feature weights are distributed in multiple trained models, making it difficult to integrate

Table 1. Summary of relevant properties of each the RF and SVM classifier.

An RF classifier is an ensemble of multiple simpler classifiers that are trained on a subset of the training data, and then the final decision is made by a voting algorithm of these classifiers (Breiman, 2001). The RF classifier was implemented using the Scikit-learn Python library using the 0.24 version which included some modifications of implementations of Breiman (2001) to avoid overfitting. We used an RF classifier that contained 500 decision tree classifiers and the entropy loss objective function. The rest of the hyperparameters were the default. The output probabilities from the RF model were normalized within each ensemble tree-based model. The predicted class probabilities of an input sample were computed as the mean predicted class probabilities of the tree-based model in the forest. The class probability of a single tree is the fraction of samples of the same class in a leaf. The predicted class probabilities of an input sample within the individual tree were summed up to one. This was implemented in the scikit-learn generic “predict_proba” function.

We used the SVM classifier as a baseline model. SVM is often used for binary classification tasks. In the current study, we used the Scikit-learn implemented SVM with the liblinear kernel (Pedregosa et al., 2011; Fan et al., 2008). L2 regularization was also added to control for overfitting. L2 regularization restricts the sum of the parameters in a regression model so that the model is forced to use simpler functions to model the data. This leads to the model having lower variance and makes it less likely to overfit (Hoerl and Kennard, 1970). Because the classification task in the current study is multiclass, the strategy to handle the multiclass support was a one-vs-rest scheme. In short, the data was reformulated into one of the four classes being class “1” and the rest of the classes being class “0”. Thus, the SVM classifier learned four one-vs-rest binarized discriminative patterns. During the testing phase, for a given matrix of one-hot coded confidence, the SVM classifier predicted the probability of being one of the four classes, giving a vector of four values. Before comparing to the one-hot coded true label, the predicted probabilities were passed to a softmax function to make the sum of the four values equal to one.

Cross-validation procedure

For within-perceptual-domain classifications, the classifier was trained on data from all subjects but one in a particular study, and the classifier was tested on the remaining subject. For across-domain generalization, the classifier was trained on one of the studies in the perceptual domain and then tested on each subject of each study of a different domain. The data was split into 'correct' and 'incorrect' examples; specifically, 'correct' and 'incorrect' here referred to whether the target trial (in which confidence was predicted) was correct or not, but the training examples (based on the previous 7 trials) could contain correct and incorrect trials. Accordingly, we run the classification analysis four times: (1) Training the classifier with the instances in which the current (to be predicted) trial was correct and testing on similar left-out instances; (2) Training the classifier with the instances in which the current trial was incorrect and testing on similar left-out instances; (3) Training the classifier with the instances in which the current trial was correct and testing on the instances in which the current was incorrect; and (4) Training the classifier with the instances in which the current trial was incorrect and testing on the instances in which the current was correct.

To measure the classification performance, the classifiers predicted the probabilities of the levels of confidence for the test set. Each column of the prediction matrix corresponded to each level of confidence. Thus, each column of the prediction matrix was compared against each column of the one-hot coded label matrix using the area under the receiver operating curve (ROC AUC). The average of the four ROC AUC measures represented the cross-validation performance of a given fold of cross-validation. The range of ROC AUC score is between zero and one, where 0.5 is referred to as the theoretical chance level and 1 is the perfect accuracy, while values below 0.5 mean that performance is worse than guessing.

Feature importance

In order to measure how the confidence ratings of the previous trials contribute to classifying the confidence rating of the current trial, we looked into the RF classifiers, computing

feature importances using permutation tests (Altmann et al., 2010). For each fold of cross-validation, after an RF classifier was trained on the training data, a permutation test algorithm was applied to both the RF classifier and the testing data. During the permutation, for the testing data, the order of the trials of confidence ratings of one of the previous trials was shuffled while the rest remained unchanged. Predictions of the testing data were made using the shuffled data. A new ROC AUC score was computed by comparing the true labels and the predictions. The difference between the new ROC AUC score and the true ROC AUC score represented the contribution of the confidence ratings of the particular previous trial that were shuffled. A positive difference meant that the feature was good for the classification, while zero meant that the feature was not important. Particularly, compared to the early proposed feature importance extraction algorithm (Breiman, 2001), the permutation feature importance algorithm could return negative feature importance estimates. Negative feature importance meant the inclusion of the feature would make the classification less effective.

Second-level statistics

After cross-validation of the classifiers for within- and cross-domain predictions, we performed second-level statistics on the results. For the classification results within the perceptual domain, we averaged the ROC AUC scores over the cross-validation folds for each study before the second-level statistics. For the cross-domain classification results, the classifier was trained in the perceptual domain and then we averaged the ROC AUC scores within each testing fold for each domain (i.e. cognitive, memory, and mixed). Therefore, each average of decoding performance contained 16 independent ROC AUC scores, one per study of the perceptual domain of the Confidence Database. Classification performance within this domain, and for each level of correctness if applicable, was assessed by comparing the ROC AUC scores of each study to 0.5 (theoretical chance level for the ROC AUC metric) using a one-sample permutation test as implemented in the EnvStats R library (<https://cran.r-project.org/web/packages/EnvStats/index.html>) but using custom Python scripts.

Initially, we computed the average of the ROC AUC. Then we subtracted the average from each experimental ROC AUC score and added 0.5 to get a vector that had a mean of 0.5 and the same variance as the original vector. We drew 16 observations with replacement from this vector, hence creating 16 ‘fake’ data points which were then averaged in order to get an estimate of the mean of the population when the ROC AUC was assumed to be 0.5. We repeated these steps (sampling and average) 10,000 times to estimate the chance level distribution (i.e. the distribution of the null hypothesis). The probability of the chance level being greater or equal to the experimental score was used as the p-value to determine the significance of the experimental score compared to 0.5. If the p-value was lower than the critical level, we determined the ROC AUC scores were significantly greater than 0.5. For cross-domain classification and each level of correctness, we performed the same statistical analysis to the ROC AUC scores. For within-perceptual domain decoding, the p-values of the permutation tests were corrected by the Bonferroni procedure in total, while for cross-domain decoding, the p-values were corrected by the Bonferroni procedure for each of the domains.

For the analysis of the feature importance estimates of the random forest classifier, we performed permutation tests similar to those described above to compare the feature importance estimates against zero. Given a set of feature importance estimates, its mean was subtracted from the estimates, hence creating ‘fake’ estimates centered at zero. We then drew samples with replacement from the ‘fake’ estimates and we saved the average of the samples. This was repeated 10,000 times to get the chance level distribution. The probability of the chance level being greater or equal to the average of the feature importance estimates was used as the p-value to determine the significance of the experimental score compared to zero. For the within-perceptual domain decoding analyses, the p-values of the permutation tests were corrected by the Bonferroni procedure. For the cross-domain decoding analyses, the p-values were Bonferroni corrected on each of the domains.

In order to analyze the linear trend of the feature importance values as a function of the trial indices, we fitted linear regression models using the trial indices as the independent variable and the feature importance values as the dependent variable. The regression models were cross-validated with a 20-fold cross-validation procedure to estimate the performance of the regression. Then, we performed permutation tests to assess the statistical significance of the regression models. We shuffled the correspondence between the feature importance values and the trial indices and then we fitted a new regression model and cross-validated the model using the same 20-fold cross-validation procedure. We repeated this for 1,000 times to estimate the empirical chance level of the regression model. The probability of the chance level being greater or equal to the average of the original regression performance was used as the p-value to determine the significance of the regression coefficients compared to zero.

Past history v.s. recent history

To further investigate how the recency of confidence ratings in the previous trials influenced the prediction of confidence rating in the current trial, we split the confidence ratings in the previous trial into “past” and “recent”, where “past” included confidence ratings in trials of T-5, T-6, and T-7, while “recent” included confidence ratings in trials of T-1, T-2, and T-3. Confidence rating in trial T-4 was not used in this analysis to equate the number of features. We applied the same feature and label preparation, machine learning models, and cross-validation procedures to decode the confidence rating in the current trial using either the confidence on “past” or “recent” trials. After cross-validation of the decoding, we applied the same second-level statistics to the ROC AUC results, comparing the ROC AUC against 0.5.

Results

We analyzed studies on metacognition across four different domains from the Confidence Database. We investigated the generation of confidence judgments by training machine learning classifiers using measures of confidence from previous trials to predict the level of confidence in the current trial. We quantified the extent to which the representation of confidence generalizes across

different observers across different cognitive domains. Two classifiers were used: a random forest (RF) classifier and a linear vector machine (SVM) as a baseline model. Testing the two classifiers allowed us to test the robustness of the predictions across the different models.

Initially, each classifier was trained in the perceptual domain and then tested within the same domain using a leave-one-subject-out cross-validation procedure for each study in the perceptual domain (see Methods). Importantly, the cross-validation procedure was performed separately for correct and incorrect trials when decoding the confidence level.

Subsequently, the classifier trained in the perceptual domain was tested across domains, i.e. cognitive, memory, and mixed domains from the Confidence Database (Rahnev et al. 2020).

Classification of confidence within the perceptual domain:

We found that the level of confidence in the current trial could be predicted based on the history of previous trials. Figure 1 illustrates that classifier predictions were successful regardless of whether the current trial was correct or incorrect. Averaging across each of the subplots in Figure 1, the decoding scores were 0.64 ± 0.048 (mean (M) $\pm$ standard deviation (SD)) for SVM and 0.64 ± 0.003 for RF (detailed statistics are shown in Supplementary Table 2). All the cross-validation with the split of data into correct and incorrect examples were significantly greater than 0.5 ($p < .001$, corrected by Bonferroni procedure for multiple comparisons, as indicated in each of the Figures). The results suggested that both linear and nonlinear models could learn the patterns of history of confidence ratings to predict the confidence rating of the current trial. Similar results were observed using the SVM classifier (Figure S1).

*Figure 1. Decoding confidence levels within the perceptual domain. The random forest classifiers were trained and tested within each study dataset using a leave-one-subject-out cross-validation procedure. The decoding scores were averaged across the cross-validation folds for each study and then the distribution of these averaged scores is plotted. The prediction of confidence level in the current trial based on the previous confidence levels was clearly above chance levels. Statistical significance was estimated by means of resampled one-sample tests against 0.5. Error Bars show the standard errors across cross-validation folds. The estimated significant levels were corrected by Bonferroni correction procedure (n.s.: not significant, *: $p < 0.05$, **: $p < 0.01$, ***: $p < 0.001$).*

We hypothesized that trials that were closer to the current trial may be more important for the prediction of confidence in the current trial than those that were further away from the current trial. To test this hypothesis, we conducted the same decoding analyses based on recent trials (-1, -2, -3 trials back) and based on past trials (-5, -6, -7 trials back). The results showed that confidence in the current trial could be predictive from both recent and also past trials, however, decoding from the most recent trials was better (see Figure 2). Similar results were observed using the SVM classifier (Figure S2).

*Figure 2. Decoding of confidence within the perceptual domain based on recent vs. past trials back. The random forest classifiers were trained and tested within each study dataset using a leave-one-subject-out cross-validation procedure for trials of T-7, T-6, T-5 and T-3, T-2, T-1 separately. The decoding scores were averaged across the cross-validation folds for each study and then the distribution of these averaged scores is plotted. Statistical significance of each unique condition was estimated by means of resampled one-sample tests against 0.5. Statistical significance of the difference between ‘past’ and ‘recent’ conditions was estimated by comparing the corresponding pair of ‘past’ and ‘recent’ conditions. Error bars show the standard errors across cross-validation folds. The estimated significant levels were corrected by Bonferroni correction procedure (n.s.: not significant, *: $p < 0.05$, **: $p < 0.01$, ***: $p < 0.001$).*

We also analyzed the feature importance of the random forest classifier (see Methods) associated with each of the trials back. The results showed that the vast majority of the trials back had feature importances that were significantly different from zero, though the most recent trial was most important for predicting confidence in the current trial (see Figure 3). We then fit a linear regression model, using the trials as the independent variable and the feature importance as the dependent variable. The model fitting was conducted for: 1) results using correct trials as the training set and tested with correct trials, 2) results using correct trials as the training set and tested with incorrect trials, 3) results using incorrect trials as the training set and tested with correct trials, and 4) results using incorrect trials as the training set and tested with incorrect trials. After cross-validating the linear regression model using leave-one-study-out procedure as described above, we derived the statistical significance of the slope of the regression model using a permutation test.

During permutation, we shuffled the correspondence between the independent variable and the dependent variable, and fitted a linear regression model using the shuffled data. This procedure was repeated 1,000 times to estimate the empirical chance level of the slope of the regression model we fitted with the unshuffled data. The significant level of the slope was the probability of the fitted slope being greater or equal than the empirical chance levels. The results of permutation tests for four conditions were corrected using Bonferroni correction. The slopes were all positive over the four conditions and they were all statistically significant (see Figure 3). Therefore, the results suggested that the more recent trials contributed more to the prediction of the current confidence.

*Figure 3. Illustration of the feature importance estimates of the RF classifiers within the perceptual domain. Permutation tests were performed on the estimates associated with each trial back assessing its significance against zero importance. The vast majority of the trials back had feature importances that were statistically significant after multiple comparison correction. The only feature importance that was not significant was the feature importance of the seventh trial back when training and testing on incorrect trials ($p = 0.059$). Error bars show the standard errors across cross-validation folds. The rest of the feature importance estimates were all significantly greater than zero ($p < 0.001$). The shaded areas associated with the regression lines were resampled standard errors of the fitted regression line. The resampling method was random sampling, differing from the cross-validation procedure, which was implemented with the algorithm of the Seaborn library. n.s.: not significant, *: $p < 0.05$, **: $p < 0.01$, ***: $p < 0.001$*

Testing the generalization of confidence across domains.

Having demonstrated that current estimates of perceptual confidence are affected by prior metacognitive decisions, we then examined whether the construction of metacognitive confidence in the perceptual domain generalizes across different other domains. Here, the classifier was trained on one of the studies in the perceptual domain (see Methods) and then tested on all the data of each experiment in a different other domain (i.e., cognitive, memory, and mixed domains; see Figure 4 and Supplementary Table 3). We observed that regardless of the training data, testing data, type of models, or domains of generalizing to, the confidence ratings of the previous seven trials were able to predict the confidence rating in the current trial (all scores were compared against to 0.5 level; all p 's $< .001$ corrected for multiple comparisons). Similar results were observed using the SVM classifier (Figure S3).

Figure 4. Cross-domain decoding scores of confidence predictions. The random forest classifiers were trained on the perceptual domain and then tested on the other three domains. The statistical significance was estimated by resampled one-sample tests against 0.5. Each subplot shows the performance of cognitive, memory and mixed domains respectively. Error bars show the standard error across cross-validation folds. The estimated significant levels were corrected by Bonferroni correction procedure (n.s.: not significant, *: $p < 0.05$, **: $p < 0.01$, ***: $p < 0.001$).

The above results demonstrate that a classifier trained in the perceptual domain generalises to predict confidence across different cognitive domains.

We then tested whether the trials that were closer to the current trial were more important for predicting confidence. The same decoding analyses were conducted but now using just the most recent trials (-1, -2, -3 trials back) or the past trials (-5, -6, -7 trials back). Again, we observed that confidence in the current trial could be predictive from both recent and also past trials, however, decoding from the most recent trials was better (see Figure 5). Similar results were observed using the SVM classifier (Figure S4).

Figure 5. *Decoding confidence across domains based on recent vs. past trials. The random forest classifiers were trained on the perceptual domain and then tested on the other three domains for trials of T-7, T-6, T-5 and T-3, T-2, T1 separately. The statistical significance of each unique condition was estimated by means of resampled one-sample tests against 0.5. Statistical significance of the difference between ‘past’ and ‘recent’ conditions was estimated by comparing the corresponding pair of ‘past’ and ‘recent’ conditions. Error bars show the standard errors across cross-validation folds. The estimated significant levels were corrected by Bonferroni correction procedure (n.s.: not significant, *: $p < 0.05$, **: $p < 0.01$, ***: $p < 0.001$).*

As before, we determined the relative importance of each of the trials back at predicting metacognitive confidence in the current trial. The results showed that the vast majority of the previous trials had feature importances that were significantly different from zero, though again the most recent trial was most important for generalizing confidence in the current trial across a different domain. Additionally, we fitted linear regression models to measure the feature importance as a function of trial indices, separately for different domains, training data, and testing data. There was a positive linear trend among all of these regression models. The results suggested that the confidence from the more recent trials were the most informative to predict the confidence ratings at the current trial. Figure 6 illustrates these results.

Figure 6. *Illustration of the feature importance estimates of the RF classifiers trained in the perceptual domain for cross-domain decoding. Error bars show the standard errors across cross-validation folds. Permutation tests were conducted on the feature importance estimates,*

*comparing against zero for each of the trials back. The shaded areas associated with the regression lines were resampled standard errors of the fitted regression line. The resampling methods used the random sampling implemented with the algorithm of the Seaborn library: n.s.: not significant, *: $p < 0.05$, **: $p < 0.01$, ***: $p < 0.001$, corrected for multiple comparisons within each domain.*

We note here that we elected to train the classifiers in the perceptual domain because the amount of training examples was much bigger in this domain. However, it may be asked whether the pattern of results generalise when the classifiers are trained in the memory or cognitive domains.

To quantify the full generalizability of confidence decoding across all domains, we tested how the decoders performed when they were trained and tested within the same domain (i.e., memory or memory) and when they were tested the decoders in a different domain (i.e., from memory to perception or cognitive). In doing so, we also tested how different attributes beyond the history of confidence ratings, namely, the history of type-1 reaction time, and the history of accuracy, contributed to predict the confidence in the current trial.

We run the classification for each combination of the different attributes: confidence only, reaction time only, accuracy only, confidence and reaction time, confidence and accuracy, and all three of them. In each combination, the features of each attribute were concatenated horizontally. Note that for this analysis, the data was not split into “correct” and “incorrect” examples. When the classifier was cross-validated within the same domain, we followed the leave-one-subject-out cross-validation procedure, so that the classifier was trained in all the data except one subject, and then tested in the left-out subject. When the classifier was cross-validated across different domains, we first trained the classifier in a given domain (i.e., Memory) and then tested the classifier in another domain (i.e., Cognitive) for each subject.

Figure 7 shows the decoding results demonstrating that history of confidence alone was the most critical attribute for predicting confidence in the current trial, and that the recent history of accuracy or type-1 reaction times did not contribute to predicting the current confidence. Figure 8 displays the feature importance of the confidence ratings in the previous trials. The pattern of results was similar to that reported in the above analyses. Across the different domains, the confidence rating from the previous one trial was generally the most important feature compared to the confidence ratings between T-2 and T-7 trials. This pattern did not change when including or adding information related to the history of accuracy and/or reaction time (Figure 8). Similar results were observed using the SVM classifier (Figure S5). We then analyzed the slopes of the regression lines of the feature importance as a function of previous trials individually for different attributes (confidence, accuracy, and RT). We concatenated the features of all attributes from the previous trials to predict the confidence in the current trial. The results are presented in Figure 8 (see also Table S1 for the statistical significance measures of the models). The results show that confidences from the previous trials contributed the most to the prediction of confidence in the current trial, and the most recent trials contributed more to the prediction than past trials. This linear trend was statistically significant within and across domains. On the other hand, relative to confidence, the accuracy and RT from the previous trials contributed little to the prediction of the current confidence.

Figure 7. Within- and cross-domain decoding ROC scores of confidence predictions. The random forest classifiers were trained on a given domain and then tested on the other three domains. The statistical significance was estimated by resampled one-sample tests against 0.5. Error bars show the standard error across cross-validation folds. The estimated significant levels were corrected by Bonferroni correction procedure (*n.s.*: not significant, *: $p < 0.05$, **: $p < 0.01$).

Figure 8. Feature importance of the attributes of the previous trials, confidence, accuracy, and reaction time, in predicting the confidence in the current trial. The random forest classifiers were cross-validated with confidence, accuracy, and reaction time from the previous seven trials as concatenated features and with confidence in the current trial as the target (see Methods).

Finally, we note that the one-versus-rest classification procedure that we used does not take into account the ordinal nature of confidence. We elected to use the random forest classifier to perform classification because it allowed us to compute the importance of each feature and make interpretable relationships between the features and the target (see Mei et al., 2020). Nevertheless, we re-run the analyses using a random forest regression procedure, which should take into account the ordinal nature of confidence. The results showed that the random forest regression models were not as sensitive as the classification models, but the regression models performed better than chance level when the feature attributes included confidence from the previous trials (i.e., confidence only, confidence + accuracy, confidence + RT, and all three attributes) (Supplementary Figures 7 and 8, and Supplementary Tables S2 illustrate the results). One potential reason that the regression models did not perform as good as the classification models may relate to the measure of the model performance, namely, variance explained (i.e. how much variance of the true confidence ratings were explained by the predicted confidence ratings). This measure may not be sensitive enough for the goal of prediction. Additionally, variance explained is more sensitive to extreme values in the predicted confidence ratings, which was not the case for ROC AUC.

Discussion

The goal of the present study was to determine whether serial dependence effects in confidence generation generalize across different observers and cognitive domains. In the perceptual domain, classification results across observers showed that the current level of confidence can be predicted by the prior confidence estimates. This result is in keeping with previous studies on confidence leak (Rahnev et al. 2015, Mueller and Weidemann 2008, Aguilar-Lleyda et al. 2021). We also observed that classifiers trained to predict confidence in the perceptual domain generalized to other cognitive domains. The extent to which confidence can be predicted is best addressed by using cross-validation with out-of-sample generalization. For instance, a recent study showed evidence of serial dependence across domains in a paradigm in which recognition judgments were interleaved with perceptual judgments so that confidence in the perceptual judgment influenced subsequent memory confidence in the next trial (Kantner et al. 2019). However, this study used standard statistical approaches in which all data are fitted at once within the same statistical model, thereby not allowing quantification and prediction of confidence in new examples and experimental contexts, namely, across different observers and across cognitive domains, as observed in the current study. Intriguingly, confidence predictions generalized across correct and incorrect trials suggesting that serial dependence effects in confidence generation are uncoupled to metacognition (i.e. how confidence tracks accuracy). In line with this, the results

showed that the recent history of confidence was the most critical factor and that accuracy or type-1 reaction time alone, or in combination with confidence, did not improve the prediction of the current confidence. Relatedly, previous observation indicated that inter-individual differences in serial dependence are negatively associated with metacognitive sensitivity (Rahnev et al. 2015). This observation highlights that serial dependence may not always confer adaptive value to perception (Cicchini, Mikellidou, and Burr 2018) and cognitive performance, but can under certain circumstances provide sub-optimal or maladaptive outcomes (Kiyonaga et al. 2017). This pervasiveness of serial dependence may reflect the habit of human observers to experience the world in a stable, auto-correlated manner (Fischer and Whitney 2014). The generalization of serial dependence in confidence generation across correct and incorrect trials is difficult to explain according to strict normative models of decision confidence as reflecting the read-out of the perceptual signal (Macmillan and Douglas Creelman 2004; Kiani and Shadlen 2009; Hebart et al. 2016), and it is in line with prior studies that observed dissociations between perceptual performance and metacognitive confidence (Samaha, Switzky, and Postle 2019; Desender, Boldt, and Yeung 2018; Samaha et al. 2016; Koizumi, Maniscalco, and Lau 2015).

The current results have important implications for the ongoing debate on the domain-generality/specificity of metacognition. Evidence in favor of domain-general metacognitive mechanisms comes from studies showing that observers can successfully assign confidence to two discriminations in different modalities (i.e. visual and auditory), notably, as efficiently as when the two discriminations involve the same sensory domain (de Gardelle, Le Corre, and Mamassian 2016). This question is normally addressed in behavioral studies by computing the correlation of individual measures of metacognition (i.e. meta- d' or M_{ratio}) across different domains (i.e. perception and memory). Studies using this approach have found mixed results with some reporting significant correlations (Fitzgerald, Arvaneh, and Dockree 2017; Ruby, Giles, and Lau, n.d.; McCurdy et al. 2013) but others failing to find such correlations (Morales, Lau, and Fleming 2018; Baird et al. 2013, 2015). A recent meta-analysis (Rouault et al. 2018) of behavioral studies observed mixed evidence for the association of inter-individual measures of metacognitive performance across perception and memory domains (but see Mazancieux et al., n.d.). Different factors can contribute to these mixed results: lack of statistical power, differences in metacognition measures and differences in task requirement across studies, uncertainty in the estimation of the model parameters across different indices of metacognitive performance such as meta- d' (Rouault et al. 2018), and also different sources of metacognitive inefficiency dominating different tasks (Shekhar and Rahnev 2021). The current results contribute to this debate by showing that classifier predictions of confidence judgments generalize across cognitive domains. However, we note that the domain-generality issue in metacognitive research typically relates to metacognitive sensitivity or efficiency, rather than confidence per se. The present results indicate that serial dependence in confidence ratings is a robust phenomenon and likely not domain-specific.

Psychology findings have been subject to recent scrutiny due to failures to replicate in new samples or replications that do not hold up to the size of the published effects (Open Science Collaboration 2015). One of the factors contributing to replication failures may be related to misuse and appropriateness of statistical tests, beyond the so-called p-hacking or the selective sub-sampling of the data. Many studies fit a statistical model with all experimental data at once,

which can lead to overfitting and poor generalization of the model with new observations that are similar but come from a different sample. Because of this, it has been argued that psychological science can greatly benefit from the field of machine learning in which pattern classifiers are tested in their ability to predict new data coming from the same or different participants (Yarkoni and Westfall 2017). These authors have made the strong argument that psychological theories contribute little to predict future human behavior with a respectable level of precision. In the same vein, there has been a recent emphasis on changing current statistical practices and encourage the adoption of “estimation thinking”, namely, to provide a quantitative model of the effect under investigation, rather than the standard “dichotomous thinking” based on the traditional null hypothesis testing framework to reject the null (Cumming 2014). The present approach using pattern classification with out-of-sample generalization thereby provides evidence that confidence generation can be reliably quantified and predicted across new samples and experimental contexts. The use of machine learning can prove very useful towards developing predictive models of confidence across different populations and experimental contexts (see (Fleming et al. 2016; Mei et al. 2020), for a similar approach to predict prospective beliefs of self-performance). It is difficult to make conclusions regarding the level of ROC prediction scores obtained here without a prior context of similar studies using different measures of predictive performance. However, ROC is arguable the best measure of predictive performance compared to other measures of effect size (Rice and Harris 2005), and the current ROC confidence prediction scores of ~ 0.65 for within-domain and slightly lower for cross-domain generalization, indicate that confidence prediction for unobserved data is a robust phenomenon.

Acknowledgments

This project was funded by project grant PID2019-105494GB-I00 from the Spanish State Research Agency (DS). We thank Megan Peters and Alan Lee for the helpful feedback on a previous version of the manuscript.

Open Practices Statements

The data for all experiments and analysis scripts are available at <https://osf.io/s46pr/> and https://github.com/nmningmei/decoding_confidence_dataset.

References

- Aguilar-Lleyda, David, Mahiko Konishi, Jérôme Sackur, and Vincent de Gardelle. 2021. “Confidence Can Be Automatically Integrated across Two Visual Decisions.” *Journal of Experimental Psychology. Human Perception and Performance* 47 (2): 161–71.
- Ais, Joaquín, Ariel Zylberberg, Pablo Barttfeld, and Mariano Sigman. 2016. “Individual Consistency in the Accuracy and Distribution of Confidence Judgments.” *Cognition* 146 (January): 377–86.
- Baird, Benjamin, Matthew Cieslak, Jonathan Smallwood, Scott T. Grafton, and Jonathan W.

- Schooler. 2015. "Regional White Matter Variation Associated with Domain-Specific Metacognitive Accuracy." *Journal of Cognitive Neuroscience* 27 (3): 440–52.
- Baird, Benjamin, Jonathan Smallwood, Krzysztof J. Gorgolewski, and Daniel S. Margulies. 2013. "Medial and Lateral Networks in Anterior Prefrontal Cortex Support Metacognitive Ability for Memory and Perception." *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience* 33 (42): 16657–65.
- Cicchini, Guido Marco, Kyriaki Mikellidou, and David C. Burr. 2018. "The Functional Role of Serial Dependence." *Proceedings. Biological Sciences / The Royal Society* 285 (1890). <https://doi.org/10.1098/rspb.2018.1722>.
- Cumming, Geoff. 2014. "The New Statistics: Why and How." *Psychological Science* 25 (1): 7–29.
- Desender, Kobe, Annika Boldt, and Nick Yeung. 2018. "Subjective Confidence Predicts Information Seeking in Decision Making." *Psychological Science* 29 (5): 761–78.
- Fischer, Jason, and David Whitney. 2014. "Serial Dependence in Visual Perception." *Nature Neuroscience* 17 (5): 738–43.
- Fitzgerald, Lisa M., Mahnaz Arvanah, and Paul M. Dockree. 2017. "Domain-Specific and Domain-General Processes Underlying Metacognitive Judgments." *Consciousness and Cognition* 49 (March): 264–77.
- Fleming, Stephen M., Sébastien Massoni, Thibault Gajdos, and Jean-Christophe Vergnaud. 2016. "Metacognition about the Past and Future: Quantifying Common and Distinct Influences on Prospective and Retrospective Judgments of Self-Performance." *Neuroscience of Consciousness* 2016 (1): niw018.
- Gardelle, Vincent de, François Le Corre, and Pascal Mamassian. 2016. "Confidence as a Common Currency between Vision and Audition." *PloS One* 11 (1): e0147901.
- Hebart, Martin N., Yoren Schriever, Tobias H. Donner, and John-Dylan Haynes. 2016. "The Relationship between Perceptual Decision Variables and Confidence in the Human Brain." *Cerebral Cortex* 26 (1): 118–30.
- Kantner, Justin, Lisa A. Solinger, David Grybinas, and Ian G. Dobbins. 2019. "Confidence Carryover during Interleaved Memory and Perception Judgments." *Memory & Cognition* 47 (2): 195–211.
- Kiani, Roozbeh, and Michael N. Shadlen. 2009. "Representation of Confidence Associated with a Decision by Neurons in the Parietal Cortex." *Science* 324 (5928): 759–64.
- Kiyonaga, Anastasia, Jason M. Scimeca, Daniel P. Bliss, and David Whitney. 2017. "Serial Dependence across Perception, Attention, and Memory." *Trends in Cognitive Sciences* 21 (7): 493–97.
- Koizumi, Ai, Brian Maniscalco, and Hakwan Lau. 2015. "Does Perceptual Confidence Facilitate Cognitive Control?" *Attention, Perception, & Psychophysics*. <https://doi.org/10.3758/s13414-015-0843-3>.
- Liberman, Alina, Jason Fischer, and David Whitney. 2014. "Serial Dependence in the Perception of Faces." *Current Biology: CB* 24 (21): 2569–74.
- Macmillan, Neil A., and C. Douglas Creelman. 2004. *Detection Theory: A User's Guide*. Psychology Press.
- Manassi, Mauro, Alina Liberman, Wesley Chaney, and David Whitney. 2017. "The Perceived Stability of Scenes: Serial Dependence in Ensemble Representations." *Scientific Reports*. <https://doi.org/10.1038/s41598-017-02201-5>.
- Mazancieux, Audrey, Stephen M. Fleming, Céline Souchay, and Chris Moulin. n.d. "Retrospective Confidence Judgments across Tasks: Domain-General Processes Underlying Metacognitive Accuracy." <https://doi.org/10.31234/osf.io/dr7ba>.
- McCurdy, Li Yan, Brian Maniscalco, Janet Metcalfe, Ka Yuet Liu, Floris P. de Lange, and Hakwan

- Lau. 2013. "Anatomical Coupling between Distinct Metacognitive Systems for Memory and Visual Perception." *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience* 33 (5): 1897–1906.
- Mei, Ning, Sean Rankine, Einar Olafsson, and David Soto. 2020. "Similar History Biases for Distinct Prospective Decisions of Self-Performance." *Scientific Reports* 10 (1): 5854.
- Morales, Jorge, Hakwan Lau, and Stephen M. Fleming. 2018. "Domain-General and Domain-Specific Patterns of Activity Supporting Metacognition in Human Prefrontal Cortex." *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience* 38 (14): 3534–46.
- Mueller, Shane T., and Christoph T. Weidemann. 2008. "Decision Noise: An Explanation for Observed Violations of Signal Detection Theory." *Psychonomic Bulletin & Review* 15 (3): 465–94.
- Open Science Collaboration. 2015. "PSYCHOLOGY. Estimating the Reproducibility of Psychological Science." *Science* 349 (6251): aac4716.
- Rahnev, Dobromir, Kobe Desender, Alan L. F. Lee, William T. Adler, David Aguilar-Lleyda, Başak Akdoğan, Polina Arbuzova, et al. 2020. "The Confidence Database." *Nature Human Behaviour* 4 (3): 317–25.
- Rahnev, Dobromir, Ai Koizumi, Li Yan McCurdy, Mark D'Esposito, and Hakwan Lau. 2015. "Confidence Leak in Perceptual Decision Making." *Psychological Science* 26 (11): 1664–80.
- Rice, Marnie E., and Grant T. Harris. 2005. "Comparing Effect Sizes in Follow-up Studies: ROC Area, Cohen's D, and R." *Law and Human Behavior*.
<https://doi.org/10.1007/s10979-005-6832-7>.
- Rouault, Marion, Andrew McWilliams, Micah G. Allen, and Stephen M. Fleming. 2018. "Human Metacognition across Domains: Insights from Individual Differences and Neuroimaging." *Personality Neuroscience* 1 (August). <https://doi.org/10.1017/pen.2018.16>.
- Ruby, Eugene, Nathan Giles, and Hakwan Lau. n.d. "Finding Domain-General Metacognitive Mechanisms Requires Using Appropriate Tasks." <https://doi.org/10.1101/211805>.
- Samaha, Jason, John J. Barrett, Andrew D. Sheldon, Joshua J. LaRocque, and Bradley R. Postle. 2016. "Dissociating Perceptual Confidence from Discrimination Accuracy Reveals No Influence of Metacognitive Awareness on Working Memory." *Frontiers in Psychology* 7 (June): 851.
- Samaha, Jason, Missy Switzky, and Bradley R. Postle. 2019. "Confidence Boosts Serial Dependence in Orientation Estimation." *Journal of Vision* 19 (4): 25.
- Shekhar, Medha, and Dobromir Rahnev. 2021. "Sources of Metacognitive Inefficiency." *Trends in Cognitive Sciences* 25 (1): 12–23.
- Yarkoni, Tal, and Jacob Westfall. 2017. "Choosing Prediction Over Explanation in Psychology: Lessons From Machine Learning." *Perspectives on Psychological Science: A Journal of the Association for Psychological Science* 12 (6): 1100–1122.