
InterpretCC: Intrinsic User-Centric Interpretability through Global Mixture of Experts

Vinitra Swamy
EPFL
vinitra.swamy@epfl.ch

Syrielle Montariol
EPFL
syrielle.montariol@epfl.ch

Julian Blackwell
EPFL
julian.blackwell@epfl.ch

Jibril Frej
EPFL
jibril.frej@epfl.ch

Martin Jaggi
EPFL
martin.jaggi@epfl.ch

Tanja Käser
EPFL
tanja.kaeser@epfl.ch

Abstract

Interpretability for neural networks is a trade-off between three key requirements: 1) faithfulness of the explanation (*i.e.*, how perfectly it explains the prediction), 2) understandability of the explanation by humans, and 3) model performance. Most existing methods compromise one or more of these requirements; *e.g.*, post-hoc approaches provide limited faithfulness, automatically identified feature masks compromise understandability, and intrinsically interpretable methods such as decision trees limit model performance. These shortcomings are unacceptable for sensitive applications such as education and healthcare, which require trustworthy explanations, actionable interpretations, and accurate predictions. In this work, we present **InterpretCC** (interpretable conditional computation), a family of interpretable-by-design neural networks that guarantee human-centric interpretability, while maintaining comparable performance to state-of-the-art models by adaptively and sparsely activating features before prediction. We extend this idea into an interpretable, global mixture-of-experts (MoE) model that allows humans to specify topics of interest, discretely separates the feature space for each data point into topical subnetworks, and adaptively and sparsely activates these topical subnetworks for prediction. We apply variations of the **InterpretCC** architecture for text, time series and tabular data across several real-world benchmarks, demonstrating comparable performance with non-interpretable baselines, outperforming interpretable-by-design baselines, and showing higher actionability and usefulness according to a user study.

1 Introduction

In recent years, the steep rise in popularity of neural networks has come with a severe weakness: the lack of interpretability of their predictions. Neural networks are considered as black-box models due to their high number of parameters and complex operations. Therefore, humans cannot understand how the features impact the network decisions under the hood; this is a crucial limitation in sensitive applications such as education or healthcare, where models' predictions might impact humans.

Interpretability is a relatively new field of research in the machine learning community. As a result of the increasing interest in providing explanations for black-box models, a range of popular methods have been proposed. Most of these are *post hoc*, *i.e.* applied after model training. Popular approaches include local instance-based methods (*e.g.* LIME [1], SHAP [2]), using adversarial examples for counterfactual explanations [3], or global, pattern-based explanations like PREMISE [4]. Post-hoc methods require the user to trust the explainer's approximation of the true explanation, which has

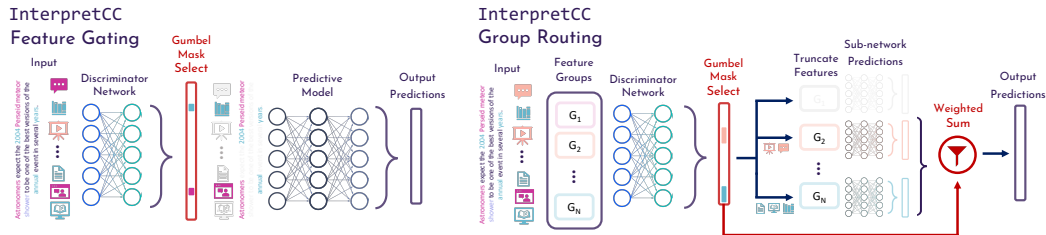


Figure 1: InterpretCC **Architectures**: *Feature Gating (left) with individual features*: (i) features are passed as input to a discriminator network which outputs a feature activation mask; (ii) features are passed individually as input to sub-network and sigmoid for prediction; (iii) activated features’ corresponding predictions are averaged for final output. *Group Routing (right) with pre-defined feature groups*: (i) features are passed through a discriminator network which outputs a route activation mask; (ii) features are split by group to be passed as input to corresponding expert sub-network for prediction; (iii) activated routes’ predictions are aggregated by a weighted sum for final output.

been shown to be systematically biased and inconsistent [5, 6]. Much fewer works provide *in-hoc* explanations, by interpreting model gradients at inference time or customizing the training protocol. Grad-CAM [7] for example assesses the contribution of a component to the model’s output. TCAV [8] and DTCAV [9] use user-defined concepts to interpret a neural network’s embeddings.

Finally, intrinsically interpretable models have mainly focused on example-based approaches, overwhelmingly for the image modality (*e.g.* B-cos networks [10], PIP-Net [11], ProtoPNet [12]) and rarely in time-series, tabular, or text modalities [13]. Very recent interpretable-by-design literature in mixture-of-experts models has highlighted a hierarchical neural network structure with subnetworks, combining interpretable experts (*i.e.* decision trees) with DNNs for partially interpretable points [14], selectively activating experts [15], or extracting automated concepts over the input space for routing [16]. However, the information used for the explanation is often overly complex and not tailored for the use case, limiting understandability and actionability.

In this paper, we therefore present a **user-centered interpretable-by-design** model that achieves guaranteed interpretability (faithfulness) and provides human-centric explanations, while maintaining comparable predictive performance to its black-box counterparts. To achieve these goals, we use conditional computation in neural networks to craft interpretable neural pathways using two different architectures (see Figure 1) based on routing through individual features or entire feature groups.

Our models enable a statement of the form: “This entry was predicted to be X because and only because it was assigned to human-interpretable categories A and B”. We refer to interpretability from the users’ perspective, focusing on the model’s local reasoning for a decision on a specific data point, as opposed to a global understanding of the model’s internals. Our models are characterized by sparse explicit routing, truncated feature spaces, and adaptivity per data point. These traits are important for human-centric trustworthiness as they provide clear and concise instance-level explanations [17, 18].

With our family of InterpretCC models, we provide the following contributions:

- [1] **InterpretCC Feature Gating: A simple, interpretable-by-design NN architecture** using a gating mechanism to sparsely activate specific features.
- [2] **InterpretCC Group Routing: An interpretable mixture-of-experts architecture** that uses human-specified group routing to separate the feature space and sparsely activate specific experts.
- [3] **An extension to real-world, human-centric modalities and domains** that are often ignored by interpretable-by-design deep learning approaches: time-series (education domain), tabular (health, synthetic data), and text (sentiment, news classification).¹
- [4] **A novel XAI user study** comparing human preferences of interpretable-by-design model explanations towards designing educational interventions for struggling students.

Across experiments on eight datasets from various domains, we show that InterpretCC models perform comparatively to non-interpretable baselines, while outperforming interpretable-by-design baselines. Moreover, participants of the user study preferred InterpretCC explanations over

¹We do not use vision datasets since extracting concepts from vision has been well-studied by existing, modality-specific interpretability approaches, *e.g.* [10, 16, 19, 20].

| Method | Explanation | | | | |
|--|-------------|--|-----------------------|------------|----------|
| | Granularity | Basis | Faithfulness | Sparsity | Coverage |
| LIME, SHAP (Post-Hoc) | Feature | Use all input features | Approximation | Not sparse | Partial |
| TCAV (In-Hoc) | Concept | User defines concepts through examples | Aligned with concepts | Sparse | Partial |
| | Feature | Use all input features | Guaranteed | Not sparse | Full |
| SENN | Concept | Automated concept selection integrated | Aligned with concepts | Not sparse | Partial |
| | Feature | Use all input features | Guaranteed | Not sparse | Full |
| InterpretCC (Feature Gating, Group Routing) | Feature | Use all input features | Guaranteed | Sparse | Full |
| | Concept | External definition of group of features as concepts | Guaranteed | Sparse | Full |

Table 1: **Method Comparison:** InterpretCC models are at the unique intersection of explanation faithfulness, sparsity, coverage, flexible granularity, and user-defined basis.

baselines in terms of actionability, usefulness, conciseness and trustworthiness. We provide our code and experiments open source: <https://github.com/epfl-ml4ed/interpretcc/>.

2 Background

Architecture Foundations. *Conditional Computation (CC)* has become widely used to improve the computationally expensive training and inference of large neural networks by activating only parts of the network [21, 22]. Bengio et al. [23] outline how block dropout conditional computation policies can be optimized using reinforcement learning. Inspired by the foundations laid out by CC, mixture-of-expert models have rapidly gained popularity for improving the efficiency of neural networks. BASE layers [24] direct each token to a designated expert and Switch Transformers [25] use CC to select one out of 4 feedforward networks across each transformer layer, optimizing computational resources. Mixtral [26] is a recent LLM using a mixture of experts to select 2 out of 8 expert networks at each layer, reducing the numbers of active parameters by a factor of 4 compared to training, while allowing each token to have access to all the parameters. With the InterpretCC models, we extend a similar routing idea with instance-dependent gating decisions towards an interpretability objective as opposed to only an efficiency or performance objective.

Interpretability Foundations. Explainability can be integrated into different stages of the modeling pipeline: post-hoc (after model training), in-hoc (e.g., interpreting model gradients), and intrinsic (interpretable by design). Our positioning of InterpretCC in comparison to the most popular approaches is described in Table 1. Specifically, we categorize the approaches using their basis and the stage they are applied to in the pipeline as well as four key requirements for human-centric explanations [17, 18]: *faithfulness* (the explanation reflects the model behavior with certainty), *sparsity* (the model uses a minimal amount of features to ease actionability and understandability), *coverage* (the explanation covers the entirety of what the model uses for prediction), and *flexible granularity*, where the key features or the feature grouping can be *user-defined*.

Post-hoc methods such as LIME [1] or SHAP [2] approximate what the model finds important, and therefore cannot be faithful or complete. Other in-hoc interpretability approaches use human-understandable concepts (e.g. TCAV [8] and DTCAV [9]), or hybrid methods using both human-defined and unsupervised concepts [13]. Such example-based approaches (e.g., TCAV) provide only partial coverage and limited faithfulness. Inspired by these approaches, InterpretCC allows users to specify interpretable concepts that are directly useful to them. However, we do not use examples, but instead allow users to specify groupings over the feature space, achieving guaranteed interpretability and full coverage.

Initial approaches have also explored expert models for intrinsic interpretability. The Interpretable Mixture of Experts (IME) framework [14] proposes different tradeoffs to using a mixture of experts for interpretability, showcasing how linear models can be used as experts to have faithful explanations.

LIMoE [27] specializes visual experts to identify concepts such as textures and faces, enriching interpretability for vision. Also for vision, interpretability based on groups is explored by the Sum-of-Parts (SOP) model [16], where the prediction process involves making sparse groups of features, highlighting the model’s dependence on subsets of features for its decisions. The most relevant intrinsic models to our approach are Self-Explaining Neural Networks (SENN) [28] and Neural Additive Models (NAM) [29]. SENN automatically extracts concepts described by prototypical examples alongside concept relevances to be combined for a prediction, again limiting faithfulness and coverage. NAM employs models for each feature in the dataset which outputs a score that is linearly combined for the prediction, making it non-sparse. InterpretCC differentiates from these models by filtering the feature space (sparsity), using user-specified concepts (basis), and not supporting overlapping in feature groups (coverage).

3 Methodology

Given an input x , the objective of our approach is to select a sparse subset of x that will be used to predict the output to solve the classification task. We propose two architectures:

Feature Gating: The approach only processes a subset of the features by applying a sparse mask \mathcal{M} on the input x before processing it by a model f . The output is given by: $f(\mathcal{M}(x))$.

Group Routing: A sparse mixture of models [30] applied on human-interpretable groups of features where each expert is assigned to a group of features:

$$\sum_{i=1}^K G(x)_i \cdot f_i(\mathcal{M}_G(x)_i), \tag{1}$$

where $\mathcal{M}_G(\cdot)_i$ is a sparse mask selecting only the features of group i , f_i is the expert model associated with the i -th group, and $G(x)_i$ is the output of the gating network for group i . If $G(x)_i = 0$, the entire group of features is ignored.

3.1 Feature Gating

The Feature Gating architecture shown in Fig. 1 is the first step towards using CC paths for interpretability. The features are first passed through a discriminator network D to select which ones to use for computing the output. The Gumbel Softmax trick [31] is applied on each dimension of $D(x)$ to select features in a differentiable way (see Appendix F for more details). A feature j is activated (the associated value in the mask is non-zero) if the Gumbel Softmax output exceeds a threshold τ , a hyperparameter. This allows the model to adaptively select the number of features based on each instance, using fewer features for simpler cases and more for complex ones.

The output is computed using a model f on the masked input $\mathcal{M}(x)$. Since the explainability is at the feature level, using a black box model for f does not detract from the interpretability.

3.2 Group Routing

We build upon the instance-dependent gating architecture with feature groups. As displayed in Figure 1, instead of selecting features individually, the mask is applied to human interpretable groups of features. Doing so encourages cross-feature interactions while maintaining a meaningful grouping for human users. To select the features belonging to group i , we use a binary mask $\mathcal{M}_G(x)_i$ that is computed using human-specified rules. In section 4, we detail our approach to compute $\mathcal{M}_G(x)_i$ for each dataset used in our experiments.

Group Routing is a sparse mixture of experts utilizing a gating network to assign a weight $G(x)_i$ to each group. This process mirrors that of Feature Gating, starting with a discriminator network D_G whose output has K dimensions (K being the number of groups). It then applies the Gumbel Softmax and a threshold τ_G to each group. As indicated in equation 1, the output of the model is a weighted sum of the output of each expert f_i that only uses the features from the i -th group as input. Using our sparsity criteria, we ensure that few groups are used to compute the output, making the Group Routing intrinsically interpretable at the group level, regardless of the types of models used as experts. Group Routing also exhibits several traits in efficient inference without compromising the number of parameters the model can use at training.

During the training phase, we employ soft masking, allowing all weights $G(x)_i$ to remain non-zero, thus granting the model access to every expert. This approach allows the model to leverage the full set of parameters during training, enhancing the training efficiency. At inference time we switch to using a hard mask, making the weights sparse and allowing for interpretability and efficiency.

4 Experimental Settings

We apply InterpretCC to four domains: education, news/sentiment classification, healthcare, and synthetic data covering *Time Series*, *Text*, and *Tabular* inputs; all for classification tasks. For *Tabular* features, the input is a vector $x \in \mathbb{R}^n$. The mask in the Feature Gating is a sparse vector indicating which tabular feature to use and how important they are (if the weight is non-0) and the groups form a partition over the features. For *Text* features, the input is a sequence of N tokens: $x = [t_1, t_2, \dots, t_N]$. The mask is a sparse vector that indicates which token to use and each group consists of a subset of the tokens. Finally, we consider *Time Series* of n features across T timesteps: $x \in \mathbb{R}^{n \times T}$. In this case, we apply the same mask across all time steps for both Feature Gating and Groups.

EDU (*time series, education domain*). We predict student success in the early weeks of four massive open online courses (MOOCs), using students’ clickstream data (see Table 6 in Appendix A for details about the courses). The raw clickstream input is transformed into weekly time-series features that have proven useful for student success prediction in previous literature (e.g. total video clicks, forum interactions). We select 45 input features used in multiple studies [32, 33, 34, 35]. For early prediction, we only use the first 40% of time steps as input.

Grouping: To derive human-interpretable concepts from these features, we turn to learning science literature. In **routing by paper**, we create 10 distinct feature subsets based on handcrafted initial input features from 10 papers, directing each to a specific expert subnetwork. For **routing by pattern**, we organize features according to five learning dimensions identified by [36, 37]: effort, consistency, regularity, proactivity, control, and assessment-based (see Table 7 for a detailed feature classification). Thirdly, **routing by Large Language Model (LLM)**, uses GPT-4’s capabilities to aid humans in feature grouping [38]. GPT-4 is prompted as an ‘expert learning scientist’ to group the features into self-regulated behavior categories that are easy to understand, which are then used to separate the features for InterpretCC. More details are included in Appendix B.1.

AG News and SST (*text*). For news categorization (**AG News**), we classify news into four categories (‘World’, ‘Sports’, ‘Business’, ‘Sci/Tech’) given a title and description of a real-world article [39]. We use 36,000 training samples and 3,000 test samples evenly distributed across categories. For sentiment prediction (Stanford Sentiment Treebank, **SST**), we use 11,855 sentences labeled by three annotators[40] and predict a binary sentiment from a sentence fragment.

Grouping: The InterpretCC routing model assigns words to subnetworks using the Dewey Decimal Code (DDC) hierarchy of topics for book classification to create 10 subnetworks (see Table 8, Appendix B.2 for more details) [41]. Each word is encoded using SentenceBERT [42] and assigned to a subcategory (i.e. the word ‘school’ is assigned to the subcategory ‘education’ under category 300 for ‘social sciences’) and routed to the appropriate parent network.

Breast Cancer (*tabular, healthcare domain*). The Wisconsin Breast Cancer dataset aims to identify cancerous tissue from an image of a fine needle aspirate (FNA) of a breast mass [43]. It has 30 features (10 from each of 3 cell nuclei) and diagnoses (Malignant: 1, Benign: 0) for 569 patients.

Grouping: For the grouping logic, we group each cell nucleus in a separate subnetwork, enabling features representing the same part of the tissue sample to be considered together.

Synthetic Dataset (*tabular*) We use OpenXAI’s synthetic dataset [44], which includes ground truth labels and explanations, indicating the subset of features influencing each label. This dataset comprises 5000 samples, 20 continuous features, and two classes. It was created using the *SynthGauss* mechanism from 5 cluster neighborhoods (1000 points for each cluster), which ensures three desirable properties for assessing explanations: (1) feature independence, (2) unambiguous and well-separated local neighborhoods, and (3) a unique ground truth explanation for each local neighborhood.

Grouping: We group the feature space by assigning each feature to a cluster neighborhood based solely on the distribution of the training data. The average absolute value of each feature for each cluster is calculated, and the highest feature-cluster value determines the assignment.

| | Dataset | Non-interpretable Base Module | Feature-Based Interpretability | | | Concept-Based Interpretability | | |
|------------------|-----------|-------------------------------|--------------------------------|---------------|----------------------------|--------------------------------|---------------------------|---------------------------|
| | | | NAM | SENN Features | InterpretCC Feature Gating | SENN Concepts | InterpretCC Top K Routing | InterpretCC Group Routing |
| Education | DSP | 82.81 ± 2.61 | 85.20 ± 0.64 | 71.70 ± 0.95 | 90.75 ± 0.01 | 81.50 ± 2.26 | 83.08 ± 1.10 | 84.90 ± 7.59 |
| | Geo | 72.96 ± 1.59 | 65.12 ± 4.07 | 57.90 ± 2.69 | 71.92 ± 0.01 | 70.90 ± 2.45 | 80.44 ± 3.19 | 81.58 ± 0.57 |
| | HWTS | 73.93 ± 3.76 | 73.11 ± 2.13 | 68.63 ± 3.78 | 82.89 ± 0.04 | 75.10 ± 11.67 | 72.59 ± 2.84 | 78.34 ± 0.95 |
| | VA | 74.90 ± 5.28 | 71.39 ± 3.38 | 74.37 ± 1.11 | 77.80 ± 0.01 | 69.99 ± 8.83 | 71.43 ± 1.11 | 72.08 ± 3.71 |
| Health | B. Cancer | 89.70 ± 1.05 | 88.77 ± 7.31 | 80.52 ± 6.21 | 78.19 ± 3.54 | 85.26 ± 1.03 | 84.66 ± 3.02 | 94.85 ± 1.25 |
| Text | AG News | 89.93 ± 3.32 | Not Supported | Not Supported | 85.72 ± 5.31 | Not Supported | 87.25 ± 2.48 | 90.35 ± 1.07 |
| | SST | 91.12 ± 2.03 | Not Supported | Not Supported | 88.21 ± 3.41 | Not Supported | 92.98 ± 0.88 | 91.75 ± 1.86 |
| Synthetic | OpenXAI | 86.67 ± 0.31 | 87.85 ± 1.31 | 83.67 ± 1.86 | 89.51 ± 0.51 | 84.67 ± 4.04 | 90.83 ± 1.93 | 89.47 ± 2.89 |

Table 2: InterpretCC Performance (avg ± std) on EDU (balanced accuracy), Text, Health, and Synthetic (accuracy) datasets compared to a non-interpretable baseline and three interpretable-by-design baselines. All 95% CI overlap with the non-interpretable base module for values in black. Colored values indicate significantly higher (green) or lower (red) performance than the base module. The reported Group Routing results are the best performing InterpretCC variations from Table 3.

5 Experimental Results

Through the following three experiments, we demonstrate that our InterpretCC models do not compromise *performance* compared to black-box models and provide explanations that are *faithful* as well as *human-centered*.

Experimental Setup. We run hyperparameter tuning and three different random seeds for each reported model (reproducibility details in Appendix E). Since EDU MOOC courses have a low passing rate (below 30%), and thus the dataset has a heavy class imbalance, we use balanced accuracy for evaluation. The other datasets are more balanced (AG News, SST, Breast Cancer, synthetic), hence we use accuracy as our evaluation metric. We perform an 80-10-10 train-validation-test data split stratified on the output label, to conserve the class imbalance in each subset. In addition to InterpretCC Feature Gating and Group Routing, we also employ a InterpretCC Top-K expert network solution with k=2 for group routing. This approach is similar to existing mixture-of-expert approaches [26, 15], except that their models make a layer-wise expert choice, which significantly reduces interpretability, while we make a global expert choice.

Base Prediction Module. We choose simple yet performant predictive modules reported in previous literature to isolate the difference in performance due to the interpretable architectures. For the EDU data, previous works uses BiLSTMs on student behavioral data for best predictive performance [6, 35]. Thus, for comparative benchmarking, the most performant BiLSTM setting is used as a baseline model [45]. For the AG News and SST datasets, we use fine-tuned DistilBERT² variations as baselines, also used in related works [46, 47]. For the Breast Cancer dataset, we use a fully connected network as reported in [48].

Interpretable Baselines. In addition to non-interpretable baselines, we compare ICC to two interpretable-by-design methods: SENN [28] and NAM [29], adaptable to any base architecture. SENN generates concept-based explanations and is designed to learn "*interpretable basis concepts*" in parallel with the model optimization. Explanations consist of the concepts most similar to the input. NAM, a General Additive Model [49], uses individual neural networks for each input feature to calculate feature weights. The model’s output is the sum of these weights, and explanations are given by displaying the feature weights.

5.1 Exp. 1: InterpretCC does not compromise on performance

Table 2 shows the models’ metrics (average accuracy, standard deviations, and 95% confidence intervals (CIs)) across three iterations of model training for all eight datasets.

InterpretCC Feature Gating improves performance with respect to the non-interpretable base prediction module for two EDU courses (DSP, HWTS) and the synthetic dataset. It shows comparable performance on all other datasets (indicated by the overlapping CIs) except the Breast Cancer dataset, suggesting that a higher number of the available features is necessary for performant prediction. The

²<https://huggingface.co/distilbert/distilbert-base-uncased>

| Dataset | Baseline | InterpretCC Group Routing | | |
|---------|--------------|---------------------------|---------------------|---------------------|
| | | Paper | Pattern | GPT-4 |
| DSP | 82.81 ± 2.61 | 82.37 ± 6.27 | 82.29 ± 3.72 | 84.90 ± 7.59 |
| Geo | 72.96 ± 1.59 | 69.64 ± 1.23 | 81.58 ± 0.57 | 81.19 ± 1.53 |
| HWTS | 73.93 ± 3.76 | 78.34 ± 0.95 | 72.34 ± 2.77 | 75.12 ± 4.17 |
| VA | 74.90 ± 5.28 | 69.88 ± 2.93 | 72.08 ± 3.71 | 70.98 ± 2.77 |
| Average | 76.65 ± 3.31 | 75.56 ± 2.85 | 77.57 ± 2.69 | 78.05 ± 4.01 |

Table 3: InterpretCC Group Routing Performance: balanced accuracy (average ± std) on variations of routing strategies (paper, pattern, GPT-4) for the EDU datasets in comparison to the non-interpretable baseline.

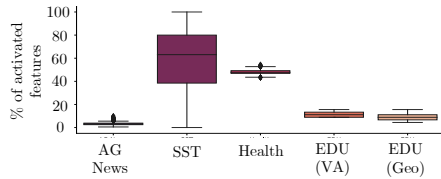


Figure 2: InterpretCC Feature Gating Sparsity: % of features activated per data point across five representative datasets.

interpretable-by-design baselines SENN and NAM never outperform the base model, and always under-perform or perform comparably to InterpretCC.³

InterpretCC Group Routing outperforms (for Geo course and Breast Cancer dataset) or performs comparably (95% CI overlap) to the fine-tuned, non-interpretable base module. It also consistently performs at least comparatively to SENN which, besides never outperforming the base model, has a relatively high variance. We further observe that the selected grouping method impacts performance (Table 3). We achieve a 10% increase in performance compared to the base model when grouping using patterns or GPT-4 for the Geo course. On average, over the four EDU courses, the automated LLM grouping and the pattern-based human-defined grouping perform comparably, showing that using automated grouping methods does not mean compromising on performance.

InterpretCC performs comparable to black-box models and outperforms interpretable baselines.

5.2 Exp. 2: InterpretCC provides faithful and user-friendly explanations

Table 4 showcases the faithfulness of InterpretCC models in comparison to SENN and NAM on ground truth explanations from the synthetic dataset. All 95% CI overlap in ground truth accuracy (how close the explanation is to ground truth feature importance, similarly to [50]) and Pairwise Rank Agreement (PRA) [44], except for SENN Concepts performing statistically significantly worse than InterpretCC Feature Gating in accuracy and InterpretCC Group Routing on PRA. This indicates that ICC’s identification of important features, and of their order of importance, is close to the ground truth and better than all baselines.

Besides being faithful, our models enable concise and hence user-friendly explanations by sparse feature (group) activations [6, 17]. Figure 2 shows the percentage of activated features per data point for ICC Feature Gating. For EDU, only ~ 10% of the 45 features are activated, with low standard deviation. For **Breast Cancer**, more features seem to be needed to make the prediction: 39.7% of the features are activated per data point. Contrarily to other datasets, for text-based data, the number of available features (words) varies per data point. **AG News** consists of news articles (average number of words per sample: 35), and only a small percentage of words is activated for each article. **SST** contains much shorter sentences (7 words on average), with high variance, leading to a higher percentage of features selected per data point (59.8%) as well as a high standard deviation. The achieved sparsity, concisely indicating the most important features in the original data set (especially for EDU and AG News), is a major advantage of our feature gating approach. In comparison, post-hoc explainers often select a wider range of important features (e.g., in previous work [6] on EDU, LIME and SHAP [1, 2] indicated broad importance over all input features).

ICC Group Routing activates different subnetworks with different weights for each data point. Figure 3 illustrates the number of activations and the average weight for each subnetwork for the text data sets (see Appendix G for detailed analysis on the Breast Cancer and EDU datasets). For **AG News** (Figure 3 top), the average activation weight is similar across all subnetworks (min 0.10, max 0.21). However, some subnetworks are activated much more frequently (400 - Language: 18, 335 times). This indicates that most data points will be routed through the same subset of subnetworks, while the remaining subnetworks are important for specific data points only. **SST** (Figure 3 bottom) shows

³We do not report results for text features with NAM and SENN, since having a network for every word reach word as a concept would be infeasible.

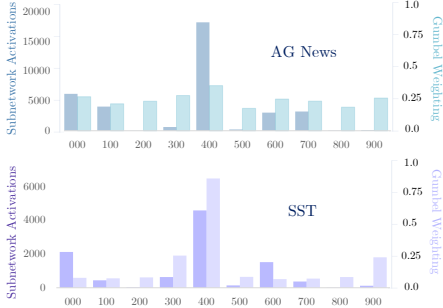


Figure 3: **AG News** and **SST**: # of subnetwork activations (left) vs. average activation weights (right), grouped by subnetworks based on the Dewey Decimal Code.



Figure 4: **AG News**: Qualitative examples of two articles, where highlighted text represents a subnet activation. The pie chart represents the weight of each subnetwork in the resulting prediction.

similar subnetwork activation patterns. However, in contrast to **AG News**, the distribution of average weights is not uniform: only three networks are activated with weights larger than 0.15.

Figure 4 illustrates two entries of **AG News** and corresponding ICC explanations. In the top example (Perseid meteor shower), the words ‘stars’, ‘meteor’, and ‘SPACE’ are routed to the *Pure Science* (500) subnetwork with a 50% activation weight, resulting in the correct prediction of ‘Sci/Tech’ category. Likewise, for the bottom article, both the *Technology* and *Arts* subnetworks are highly weighted, resulting in the correct prediction of the ‘Business’ category. Interestingly, subnetwork *Language* (400) is also activated. We suspect the high weights showcased for 400 in Figure 3 are representative of words the DDC does not have a close relation to in SentenceBERT embedding space. Examples of explanations for the EDU domain for all models can be found in Appendix D, Figures 6 and 7. We also compare ICC behavior across the three grouping strategies for the DSP course (Appendix G.2.2) and analyze the network sparsity at different prediction horizons (G.2.1). Finally, we show the variation of the selection of feature groups across all MOOCs, indicating the adaptability of ICC in selecting relevant subnetworks whatever the grouping method (G.2.3).

| Model | PRA [44] | Ground Truth Accuracy |
|----------------|---------------------|-----------------------|
| InterpretCC FG | 69.71 ± 3.89 | 94.84 ± 4.11 |
| InterpretCC GR | 74.55 ± 1.21 | 89.51 ± 0.56 |
| SENN Features | 71.99 ± 3.25 | 85.83 ± 2.56 |
| SENN Concepts | 45.30 ± 7.49 | 65.19 ± 9.23 |
| NAM | 70.16 ± 5.63 | 87.39 ± 2.45 |

Table 4: **Synthetic faithfulness evaluation** across five models on ground truth explanations using PRA and accuracy.

| | NAM | SENN | ICC GR | ICC FG | Weight |
|-----------------|-------------|--------------------|--------------------|--------------------|--------|
| Usefulness | 3.25 ± 0.98 | 3.3 ± 1.11 | 3.53 ± 1.11 | 3.88 ± 0.94 | 0.28 |
| Trustworthiness | 3.28 ± 0.93 | 3.64 ± 0.92 | 3.36 ± 1.06 | 3.78 ± 0.9 | 0.23 |
| Actionability | 3.08 ± 0.96 | 3.25 ± 1.06 | 3.37 ± 1.04 | 3.77 ± 0.95 | 0.21 |
| Completeness | 3.18 ± 1.02 | 3.76 ± 1.09 | 3.1 ± 1.19 | 3.67 ± 1.07 | 0.16 |
| Conciseness | 3.13 ± 1.06 | 2.82 ± 1.31 | 3.72 ± 1.06 | 3.68 ± 1.05 | 0.12 |
| Global | 3.2 ± 0.81 | 3.38 ± 0.85 | 3.41 ± 0.88 | 3.78 ± 0.77 | |

Table 5: **Model score for each criterion** (average ± std), and criteria weight according to their ranking. All scores range from 1 (lowest) to 5 (highest).

InterpretCC provides sparse and hence user-friendly explanations, while not compromising on explanation faithfulness.

5.3 Exp. 3: InterpretCC explanations are preferred by humans

Setting. To validate the user-centeredness of our approach, we conducted a user study comparing InterpretCC’s explanations with the ones from other interpretable-by-design methods, SENN and NAM. We focused on the education domain and time series input type, using the DSP course in the EDU dataset. We trained InterpretCC Feature Gating and Group Routing (with pattern-based feature grouping strategy, since it is heavily grounded on expert knowledge) as well as SENN and NAM and randomly selected four test samples (*i.e.*, four students) for prediction.

We recruited 56 teachers using Prolific⁴, (see Appendix D for detailed information about the participants’ demographics and backgrounds (Figure 5) and the content of the study). We showed them each model’s prediction of the student’s success or failure along with its explanation. The

⁴www.prolific.com/

explanations were given as a short text and a graph showing the features and concepts used by the model. Note that the choices we made for the presentation of the explanations might have an influence on the participants’ perception of the explanations. Examples of how the explanations were presented and a discussion on this limitation can be found in Appendix D, Figures 6, 7, 9 and 8. We asked participants to compare these explanations according to five criteria (aligned with [51]), and to rank the criteria in terms of importance.

- [1] **Usefulness:** This explanation is useful to understand the prediction.
- [2] **Trustworthiness:** This explanation lets me judge if I should trust the model.
- [3] **Actionability:** This explanation helps me know how to give feedback to the student.
- [4] **Completeness:** This explanation has sufficient detail to understand why the prediction was made.
- [5] **Conciseness:** Every detail of this explanation is necessary.

Results. Using Friedman’s Chi-Square test, we verify that the ranking of the criteria is consistent among the participants ($p < 0.05$). We report normalized criteria importance in Table 5 (last column). We observe that Usefulness, Trustworthiness and Actionability are consistently ranked the highest by the participants. We compute the *Global satisfaction* score as a weighted average of the criteria as a global satisfaction measure for each model (Table 5, last row).

We observe that ICC FG ranks the highest in the top 3 most important criteria as well as in global satisfaction. Furthermore, ICC GR ranks second in Usefulness and Actionability and first in Conciseness. Overall, InterpretCC models are favored over interpretable baselines in 4 out of 5 criteria and in terms of global satisfaction. An ANOVA performed separately for each criterion as well as the global satisfaction measure (Table 10, Appendix D.3), indicates that there is indeed a significant difference between the models for each criterion. Post-hoc pairwise Tukey’s Honest Significant Difference (HSD) tests confirm that ICC FG significantly outperforms NAM and SENN on all criteria but completeness (see Figure 11, Appendix D.3 for details).

Participants prefer InterpretCC explanations in terms of usefulness, trustworthiness, actionability, and conciseness over interpretable baselines.

6 Discussion and Conclusion

We propose InterpretCC, a family of interpretable-by-design models that puts interpretability and human understanding at the forefront of the design. Through our experiments with feature gating models and interpretable group routing (mixture-of-expert) models, we demonstrate that our optimally interpretable modular architectures do not compromise on performance. In a real-world setting, we show that InterpretCC models are preferred over other intrinsically interpretable models in 4 of 5 explainability criteria.

Limitations. InterpretCC Feature Gating is best when individual features are important for actionable decisions based on the explanation. However, it can lose sparseness if all features are equally important: consider the extreme case where the prediction is a sum over the full feature space and all features are independent. While this is rare in real-world datasets, it is a critical design consideration. InterpretCC Group Routing requires more human effort than Feature Gating and is, therefore, more adapted to scenarios with large feature spaces, where broader concepts are more relevant than individual features. For both InterpretCC models, we find that tuning the learning rate and feature threshold parameter was especially key to achieving strong results.

User-defined feature groups aim at deriving explanations useful to the user, but might compromise performance if the user specifies a grouping that carries minimal signal. This is by design: regardless of the grouping, InterpretCC optimizes for explanation actionability and understandability over performance. We believe that an accurate prediction at the cost of explanation usefulness is not worth it in an applied setting. Moreover, we have shown how the feature groups can be defined automatically using LLMs to reduce human effort, but still provide human-centric explanations.

We acknowledge that the presentation of explanations in our user study (Section 5.3) has influence over our preliminary results on user perception. We had to significantly post-process the explanations of SENN and NAM to provide them in a format understandable for a non-technical audience. An extensive study over many different prediction tasks and participants’ domains of expertise is necessary for generalizable conclusions.

Future Work. For the text datasets, we fine-tune DistilBERT models to solve the task given groups of features. For more complex tasks, for example requiring very long context size, multi-step

reasoning ability, or strong prior domain-specific knowledge, DistilBERT can be swapped with large decoder models (LLMs), either through fine-tuning or in-context learning. Parameter-efficient fine-tuning such as LoRA would allow fine-tuning even large LMs with limited computational cost, while in-context learning would use the same model instance for each feature and feature group. Overall, we encourage the machine learning community to design models for interpretability at many different granularities and user-specified requirements. InterpretCC provides one such family of models as a tradeoff between human specification, explanation certainty, and performance.

References

- [1] Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. "Why Should I Trust You?": Explaining the predictions of any classifier. *KDD*, 2016.
- [2] Scott M. Lundberg and Su-In Lee. A unified approach to interpreting model predictions. In *NeurIPS*, 2017.
- [3] Amit Dhurandhar, Pin-Yu Chen, Ronny Luss, Chun-Chen Tu, Paishun Ting, Karthikeyan Shanmugam, and Payel Das. Explanations based on the missing: Towards contrastive explanations with pertinent negatives. In *NeurIPS*, 2018.
- [4] Michael A Hedderich, Jonas Fischer, Dietrich Klakow, and Jilles Vreeken. Label-descriptive patterns and their application to characterizing classification errors. In *ICML*, 2022.
- [5] Satyapriya Krishna, Tessa Han, Alex Gu, Javin Pombra, Shahin Jabbari, Steven Wu, and Himabindu Lakkaraju. The disagreement problem in explainable machine learning: A practitioner’s perspective. *arXiv preprint arXiv:2202.01602*, 2022.
- [6] Vinitra Swamy, Bahar Radmehr, Natasa Krco, Mirko Marras, and Tanja Käser. Evaluating the explainers: Black-box explainable machine learning for student success prediction in MOOCs. In *International Conference on Educational Data Mining*, 2022.
- [7] Ramprasaath R. Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra. Grad-CAM: Visual explanations from deep networks via gradient-based localization. *International Journal of Computer Vision*, 2019.
- [8] Been Kim, Martin Wattenberg, Justin Gilmer, Carrie Cai, James Wexler, Fernanda Viegas, and Rory Sayres. Interpretability beyond feature attribution: Quantitative testing with concept activation vectors (tcav). *ICML*, 2018.
- [9] Amirata Ghorbani, James Wexler, James Zou, and Been Kim. Towards automatic concept-based explanations. *NeurIPS*, 2019.
- [10] Moritz Böhle, Mario Fritz, and Bernt Schiele. B-cos networks: Alignment is all we need for interpretability. In *CVPR*, 2022.
- [11] Meike Nauta, Jörg Schlötterer, Maurice van Keulen, and Christin Seifert. Pip-net: Patch-based intuitive prototypes for interpretable image classification. In *CVPR*, 2023.
- [12] Chaofan Chen, Oscar Li, Daniel Tao, Alina Barnett, Cynthia Rudin, and Jonathan K Su. This looks like that: deep learning for interpretable image recognition. *NeurIPS*, 2019.
- [13] Yoshihide Sawada and Keigo Nakamura. Concept bottleneck model with additional unsupervised concepts. *IEEE Access*, 2022.
- [14] Aya Abdelsalam Ismail, Sercan O Arik, Jinsung Yoon, Ankur Taly, Soheil Feizi, and Tomas Pfister. Interpretable mixture of experts. *Transactions on Machine Learning Research*, 2023.
- [15] Margaret Li, Suchin Gururangan, Tim Dettmers, Mike Lewis, Tim Althoff, Noah A Smith, and Luke Zettlemoyer. Branch-train-merge: Embarrassingly parallel training of expert language models. *Interpolate Workshop at NeurIPS*, 2022.
- [16] Weiqiu You, Helen Qu, Marco Gatti, Bhuvnesh Jain, and Eric Wong. Sum-of-parts models: Faithful attributions for groups of features. *arXiv preprint arXiv:2310.16316*, 2023.

- [17] Tim Miller. Explanation in artificial intelligence: Insights from the social sciences. *Artificial Intelligence*, 2019.
- [18] Vinitra Swamy, Jibril Frej, and Tanja Käser. The future of human-centric eXplainable Artificial Intelligence (XAI) is not post-hoc explanations. *arXiv preprint arXiv:2307.00364*, 2023.
- [19] Jon Donnelly, Alina Jade Barnett, and Chaofan Chen. Deformable protopnet: An interpretable image classifier using deformable prototypes. In *CVPR*, 2022.
- [20] Fel Thomas, Thibaut Boissin, Victor Boutin, Agustin Martin Picard, Paul Novello, Julien Colin, Drew Linsley, Tom Rousseau, Remi Cadene, Lore Goetschalckx, et al. Unlocking feature visualization for deep network with magnitude constrained optimization. In *NeurIPS*, 2023.
- [21] Yoshua Bengio, Nicholas Léonard, and Aaron Courville. Estimating or propagating gradients through stochastic neurons for conditional computation. *arXiv preprint arXiv:1308.3432*, 2013.
- [22] Andrew Davis and Itamar Arel. Low-rank approximations for conditional feedforward computation in deep neural networks. *arXiv preprint arXiv:1312.4461*, 2013.
- [23] Emmanuel Bengio, Pierre-Luc Bacon, Joelle Pineau, and Doina Precup. Conditional computation in neural networks for faster models. In *ICLR Workshop Track*, 2016.
- [24] Mike Lewis, Shruti Bhosale, Tim Dettmers, Naman Goyal, and Luke Zettlemoyer. Base layers: Simplifying training of large, sparse models. In *International Conference on Machine Learning*, pages 6265–6274. PMLR, 2021.
- [25] William Fedus, Barret Zoph, and Noam Shazeer. Switch transformers: Scaling to trillion parameter models with simple and efficient sparsity. *The Journal of Machine Learning Research*, 2022.
- [26] Albert Q Jiang, Alexandre Sablayrolles, Antoine Roux, Arthur Mensch, Blanche Savary, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Emma Bou Hanna, Florian Bressand, et al. Mixtral of experts. *arXiv preprint arXiv:2401.04088*, 2024.
- [27] Basil Mustafa, Carlos Riquelme, Joan Puigcerver, Rodolphe Jenatton, and Neil Houlsby. Multimodal contrastive learning with limoe: the language-image mixture of experts. *NeurIPS*, 2022.
- [28] David Alvarez Melis and Tommi Jaakkola. Towards robust interpretability with self-explaining neural networks. *NeurIPS*, 2018.
- [29] Rishabh Agarwal, Levi Melnick, Nicholas Frosst, Xuezhou Zhang, Ben Lengerich, Rich Caruana, and Geoffrey E Hinton. Neural additive models: Interpretable machine learning with neural nets. *NeurIPS*, 2021.
- [30] William Fedus, Jeff Dean, and Barret Zoph. A review of sparse expert models in deep learning. *arXiv preprint arXiv:2209.01667*, 2022.
- [31] Eric Jang, Shixiang Gu, and Ben Poole. Categorical reparameterization with gumbel-softmax. In *ICLR*, 2017.
- [32] Sébastien Lallé and Cristina Conati. A data-driven student model to provide adaptive support during video watching across moocs. In *Artificial Intelligence in Education*, 2020.
- [33] Mina Shirvani Boroujeni, Kshitij Sharma, Łukasz Kidziński, Lorenzo Lucignano, and Pierre Dillenbourg. How to quantify student’s regularity? In Katrien Verbert, Mike Sharples, and Tomaž Klobučar, editors, *Adaptive and Adaptable Learning*, 2016.
- [34] Fu Chen and Ying Cui. Utilizing student time series behaviour in learning management systems for early prediction of course performance. *Journal of Learning Analytics*, 2020.
- [35] Mirko Marras, Tuan Tu, Julien Vignoud, and Tanja Käser. Can feature predictive power generalize? benchmarking early predictors of student success across flipped and online courses. *International Conference on Educational Data Mining*, 2021.

- [36] Mohammad Asadi, Vinitra Swamy, Jibril Frej, Julien Vignoud, Mirko Marras, and Tanja Käser. Ripple: Concept-based interpretation for raw time series models in education. In *AAAI*, 2023.
- [37] Paola Mejia, Mirko Marras, Christian Giang, and Tanja Käser. Identifying and comparing multi-dimensional student profiles across flipped classrooms. In *Artificial Intelligence in Education*, 2022.
- [38] Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. GPT-4 technical report. *arXiv preprint arXiv:2303.08774*, 2023.
- [39] Xiang Zhang, Junbo Zhao, and Yann LeCun. Character-level convolutional networks for text classification. *NeurIPS*, 2015.
- [40] Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D Manning, Andrew Y Ng, and Christopher Potts. Recursive deep models for semantic compositionality over a sentiment treebank. In *NeurIPS*, 2013.
- [41] Mohinder Partap Satija. *The theory and practice of the Dewey decimal classification system*. Elsevier, 2013.
- [42] Nils Reimers and Iryna Gurevych. Sentence-bert: Sentence embeddings using siamese bert-networks. *arXiv preprint arXiv:1908.10084*, 2019.
- [43] William Wolberg, Mangasarian Olvi, Nick Street, and W. Street. Breast Cancer Wisconsin (Diagnostic). UCI Machine Learning Repository, 1995.
- [44] Chirag Agarwal, Satyapriya Krishna, Eshika Saxena, Martin Pawelczyk, Nari Johnson, Isha Puri, Marinka Zitnik, and Himabindu Lakkaraju. Openxai: Towards a transparent evaluation of model explanations. *NeurIPS*, 2022.
- [45] Vinitra Swamy, Mirko Marras, and Tanja Käser. Meta transfer learning for early success prediction in moocs. In *ACM Conference on Learning@Scale*, 2022.
- [46] Zhilin Yang, Zihang Dai, Yiming Yang, Jaime Carbonell, Russ R Salakhutdinov, and Quoc V Le. Xlnet: Generalized autoregressive pretraining for language understanding. *Advances in neural information processing systems*, 32, 2019.
- [47] HF Canonical Model Maintainers. distilbert-base-uncased-finetuned-sst-2-english (revision bfdd146), 2022.
- [48] Abien Fred M Agarap. On breast cancer detection: an application of machine learning algorithms on the wisconsin diagnostic dataset. In *Proceedings of the 2nd international conference on machine learning and soft computing*, 2018.
- [49] Trevor J Hastie. Generalized additive models. In *Statistical models in S*. Routledge, 2017.
- [50] Yang Liu, Sujay Khandagale, Colin White, and Willie Neiswanger. Synthetic benchmarks for scientific research in explainable machine learning. *NeurIPS*, 2021.
- [51] Jibril Frej, Neel Shah, Marta Knezevic, Tanya Nazaretsky, and Tanja Käser. Finding paths for explainable mooc recommendation: A learner perspective. In *Learning Analytics and Knowledge Conference*, 2024.
- [52] Mona L Scott. Dewey decimal classification. *Libraries Unlimited*, 1998.
- [53] Thomas Verelst and Tinne Tuytelaars. Dynamic convolutions: Exploiting spatial sparsity for faster inference. In *CVPR*, 2020.
- [54] Mohammad Asadi, Vinitra Swamy, Jibril Frej, Julien Vignoud, Mirko Marras, and Tanja Käser. Ripple: Concept-based interpretation for raw time series models in education. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2023.
- [55] Vinitra Swamy, Sijia Du, Mirko Marras, and Tanja Kaser. Trusting the explainers: teacher validation of explainable artificial intelligence for course design. In *LAK23: 13th International Learning Analytics and Knowledge Conference*, pages 345–356, 2023.

Impact Statement

This paper introduces work aimed at advancing the field of human-centered machine learning, addressing the critical balance between interpretability and model performance in neural networks. We demonstrate that these aspects do not necessarily need to be a trade-off, particularly in fields where AI’s can influence decisions that can have effects on people’s lives, such as finance, healthcare, and education. Focusing on the applications presented in this paper, in healthcare, the ability to understand and trust AI-driven recommendations can greatly improve care delivery. In education, the application of interpretable models reduces the risks of unintended biases, help users to identify actionable items, and increases trust.

A key impacts is regulatory compliance, where guidelines and regulations around AI usage are increasing. Models like InterpretCC can help industries and organizations meet these legal standards by increasing the transparency of their AI-based decision-making tools.

ICC enhances trust and human-centered actionability in predictive tools by clearly listing the features impacting predictions. This empowers decision-makers to take informed actions. Different contexts and applications lead to different interpretability needs; different levels of granularity, and different sets or groups of features. ICC Group Routing allows users to define their own set of feature, adapted to their needs. Moreover, by identifying a finite set of key features, ICC may reduce spurious predictions and allow users to better understand the causes of unexpected model behavior. Overall, an interpretable model like ICC allows users to assess the model’s robustness more easily and faithfully.

Our tool also offers the potential to identify and mitigate biased behaviors within models, particularly if discriminatory predictions arise from certain feature uses. As different contexts may require varying levels of interpretability, the Group Routing feature of ICC allows users to define and adjust feature sets according to their specific needs, ensuring flexibility and relevance in diverse applications.

Comparable in societal and technological impact to other interpretable-by-design methods, such as those discussed by Agarwal et al. in their 2021 study, ICC represents a significant step forward in making AI more accessible and understandable for users and practitioners alike.

However, we acknowledge that despite the versatility of ICC across various data types, domains, and datasets demonstrated in this work, we cannot guarantee that it will always perform on par with or better than baseline models. It is important to note that the model weights we release are not safe for direct use in production tasks, such as detecting breast cancer, without further validation and adaptation. We encourage other researchers and practitioners to adopt explainable AI methods.

A Additional Details on Datasets

Here, we provide additional statistics regarding each dataset used in our study. In particular, we highlight their availability details and terms of use.

EDU. We predict student success during the early weeks of four massive open online courses (MOOCs), using students’ clickstream data (see Table 6 for details about the courses). Contrary to the other 4 datasets, this dataset is kept private for student privacy reasons, as required by HREC 058-2020/10.09.2020 and HREC 096-2020/09.04.2022.

| Title | Identifier | Topic | Level | Language | No. Weeks | No. Students | Passing Rate [%] |
|---------------------------------------|------------|-------|-------|----------|-----------|--------------|------------------|
| Digital Signal Processing | DSP | CS | MSc | English | 10 | 4,012 | 23.1 |
| Éléments de Géomatique | Geo | Math | BSc | French | 11 | 452 | 45.1 |
| Household Water Treatment and Storage | HWTS | NS | BSc | French | 5 | 2,438 | 47.2 |
| Villes Africaines | VA | SS | BSc | En/Fr | 12 | 5,643 | 9.9 |

Table 6: Course Details and Statistics.

Topic abbreviations: Math: Mathematics; NS: Natural Science; CS: Computer Science; SS: Social Science; Arch: Architecture; Bus: Economics and Business.

AG News is a news classification dataset, where given a title and description of a real-world article, it has to be classified into one of the four categories: ‘World’, ‘Sports’, ‘Business’, ‘Sci/Tech’ [39]. It is freely available at the following location: http://groups.di.unipi.it/~gulli/AG_corpus_of_news_articles.html, but only for non-commercial use.

SST. The Stanford Sentiment Treebank dataset aims at predicting the sentiment from a sentence fragment sourced from a movie review. The dataset is freely available here: <https://huggingface.co/datasets/sst>. This popular benchmark is an extension of the Movie Review Database (MRD) [40]. It includes two sets of labels: one for binary sentiment classification and one for multiclass. We use binary classification to demonstrate a different setting than the multiclass classification of AG News.

Breast Cancer. The Wisconsin Breast Cancer dataset attempts to identify the presence of cancerous tissue from an image of a fine needle aspirate (FNA) of a breast mass [43]. This dataset is freely available here: <https://archive.ics.uci.edu/dataset/17/breast+cancer+wisconsin+diagnostic>, and is licensed under a Creative Commons Attribution 4.0 International (CC BY 4.0) license.

Synthetic Dataset. We generate 5000 points of OpenXAI’s synthetic dataset [44] through the class `generate_gaussians` with 5 cluster centers, which includes both ground truth labels and ground truth explanations, and is available here: <https://github.com/AI4LIFE-GROUP/OpenXAI> (under MIT license).

B InterpretCC Group Routing Schema

In the following, we discuss the exact grouping schematic employed on the 4 EDU datasets (including three schemes) and the 2 text datasets.

| Dimensions | Corresponding measures | Student patterns | Code | Field of Study | | |
|--------------------|-----------------------------|------------------|------------|---|------------|---------------------------|
| Effort | Total time online | Higher intensity | 000 | Computer Science, Information and General Works | | |
| | Total video clicks | Lower intensity | | | | |
| Consistency | Mean session duration | Uniform | | | 100 | Philosophy and Psychology |
| | Relative time online | First half | | | 200 | Religion |
| | Relative video clicks | Second half | | | 300 | Social Sciences |
| Regularity | Periodicity of week day | Higher peaks | | | 400 | Language |
| | Periodicity of week hour | Lower peaks | | | 500 | Pure Science |
| | Periodicity of day hour | | | | 600 | Technology |
| Proactivity | Content anticipation | Anticipated | | | 700 | Arts and recreation |
| | Delay in lecture view | Delayed | 800 | Literature | | |
| Control | Fraction time spent (video) | Higher intensity | 900 | History and geography | | |
| | Pause action frequency | Lower intensity | | | | |
| | Average change rate | | | | | |
| Assessment | Competency strength | Higher intensity | | | | |
| | Student shape | Lower intensity | | | | |

Table 7: **EDU Routing by Pattern** uses learning dimensions from [37] to create interpretable feature groupings.

Table 8: **Text Routing** by the Dewey Decimal Classification system [52]. Each code represents a sub-network in the text variations of the InterpretCC framework.

B.1 Feature Grouping Schema for EDU Datasets

The data is composed of student clickstream data from four different MOOCs. The raw clickstream input is transformed into 45 weekly time-series features that have proven useful for student success prediction in previous literature (e.g. total video clicks, forum interactions).

Routing by pattern. We organize features according to four learning dimensions identified by [36] and detailed in Table 7—effort, consistency, regularity, proactivity, control, and assessment—based on [37], with a focus on these dimensions in an extended experiment.

Routing by GPT-4. To leverage LLMs for automated feature groupings, we use the following prompt alongside the feature descriptions from the appendix table of Swamy et. al [45].

Prompt: You are a talented learning scientist, specializing in self-regulated learning behavior. You want to group the following features together into smaller groups of logical concepts based on educational theory with headings. What are your feature groupings?

Response: Certainly, based on educational theory and self-regulated learning behavior, I would group the provided features into the following logical categories:

- Temporal Engagement:
 - DelayLecture
 - RegPeakTimeDayHour
 - RegPeriodicityDayHour
- Study Patterns:
 - NumberOfSessions
 - AvgTimeSessions
 - TotalTimeSessions
 - StdTimeSessions
 - StdTimeBetweenSessions
- Video Engagement:
 - AvgWatchedWeeklyProp
 - StdWatchedWeeklyProp
 - AvgReplayedWeeklyProp
 - StdReplayedWeeklyProp
 - AvgInterruptedWeeklyProp
 - StdInterruptedWeeklyProp
- Competency and Learning Progress:
 - CompetencyStrength
 - CompetencyAlignment
 - CompetencyAnticipation
 - ContentAlignment
 - ContentAnticipation
- Interaction Patterns:
 - TotalClicks
 - TotalClicksProblem
 - TotalClicksVideo
 - TotalClicksWeekday
 - TotalClicksWeekend
 - TotalTimeProblem
 - TotalTimeVideo
 - TotalClicksVideoLoad
 - TotalClicksVideo
 - FrequencyEventVideo
 - FrequencyEventLoad
 - FrequencyEventPlay
 - FrequencyEventPause
 - FrequencyEventStop
 - FrequencyEventSeekBackward
 - FrequencyEventSeekForward
 - FrequencyEventSpeedChange
 - AvgSeekLength
 - StdSeekLength
 - AvgPauseDuration
 - StdPauseDuration
 - AvgTimeSpeedingUp
 - StdTimeSpeedingUp
- Assessment Behavior:
 - StudentSpeed
 - StudentShape

Each grouping represents a different aspect of self-regulated learning behavior, focusing on how students engage with learning resources, interact with content, demonstrate competency, and approach assessments. This categorization aligns with principles of self-regulated learning and can help in analyzing and understanding students’ behaviors and strategies within an educational context.

Routing by Paper. In routing by paper, we create 4 distinct feature subsets based on handcrafted initial input features from 4 learning science research papers highlighted as predictive for MOOCs in [45]. We filter each feature set to a expert subnetwork for the research paper the set originated from [32, 35, 34, 33].

B.2 Feature Grouping Schema for Text Datasets

For news categorization (AG News) and sentiment prediction (SST) feature grouping, we assign words to subnetworks. For this, we use the Dewey Decimal Code (DDC) for librarians and its hierarchy of topics for book classification to create 10 subnetworks, as showcased by topic in Table 8 [41]. Each word is assigned to a subcategory (i.e. the word ‘school’ is assigned to the subcategory ‘education’ under category 300 for ‘social sciences’) and routed to the appropriate parent network. The decision to use the DDC was to use subnetworks that were standardized, pervasive in daily life and clearly human-understandable. To conduct this assignment, we utilize SentenceBERT to encode the subtopics for each DDC heading (i.e. all of 010, 020, 030, etc. for the category 000) [42]. The choice of SentenceBERT is motivated towards capturing the broader context of multi-word category headings in a lightweight model. During training and inference, we again use SentenceBERT to

encode each word in the input instance, then assign each word to the most similar DDC category in embedding space with cosine similarity.

C Sparsity Criterion Experiments

For the feature gating architecture to further improve interpretability, we would like the network to learn sparse feature activations. That is, for a given input x we would like to reduce the number of features that affect the model prediction. To achieve this we apply regularization to the generated feature mask.

One natural choice to enforce sparse feature activations is to apply L_1 -norm regularization to the feature mask, which penalizes a high number of nonzero elements. Another choice is to use annealed regularization as presented by Verelst and Tuytelaars [53], which might aid the model to first work through a prediction optimization phase that is not confounded by any additional error terms before moving towards a sparsity-enforcing phase.

We experiment with annealing L1 and L1 regularization across four courses, and find that traditional L1 regularization is more stable (at least in the time-series setting). The Baseline BiLSTM results are not reported as confidence intervals here as they are directly sourced from a recent benchmarking paper by Swamy et al., with confirmed similarity by Asadi et al. [45, 54]. We reproduce this benchmark above with similar values in Table 9.

| EDU Dataset | Baseline | InterpretCC Feature Gating | |
|-------------|----------|----------------------------|----------------|
| | | Annealing | L1 |
| 40% EP | BiLSTM | | |
| DSP | 82 | 87.76 +/- 3.12 | 90.75 +/- 0.01 |
| Geo | 76.2 | 81.13 +/- 5.39 | 71.92 +/- 0.01 |
| HWTS | 72 | 77.58 +/- 0.01 | 82.89 +/- 0.04 |
| VA | 73.8 | 84.81 +/- 0.01 | 77.80 +/- 0.01 |

Table 9: Annealing L1 regularization in comparison with L1 regularization across EDU datasets.

D User Study

In this section, we discuss the details of the user study presented in Section 5.3, discussing the design, content, and additional analyses of the results (including an ANOVA and Tukey HSD tests). We designed the study over four rounds of pilots, with 8 individuals from diverse backgrounds, continuously updating the study design following their feedback. The survey has been approved by the the Human Research Ethics Committee (HREC) under application number HREC 065-2022/27.09.2022.

D.1 Design of the Study

We recruit 56 participants using Prolific⁵, selecting the ones who identified their current profession as a teacher and who have at least a bachelor’s degree. Our target participants have expertise in teaching, as they would be well-suited to understand both the educational context of the study and the consequences of black-box models for student outcomes. During the study, we ask the participants whether they have ever taken or prepared material for an online course (MOOC), their level of education, and what level they are teaching at (from primary school to graduate school). Detailed demographics distribution can be found in Figure 5. The sample of participants is gender-balanced, and about half of them have taken or participated in creating a MOOC. The median completion time is 22 minutes, and the average reward per hour is £14.55.

At the beginning of the study, the participants are shown the following welcome message and disclaimer:

⁵www.prolific.com/

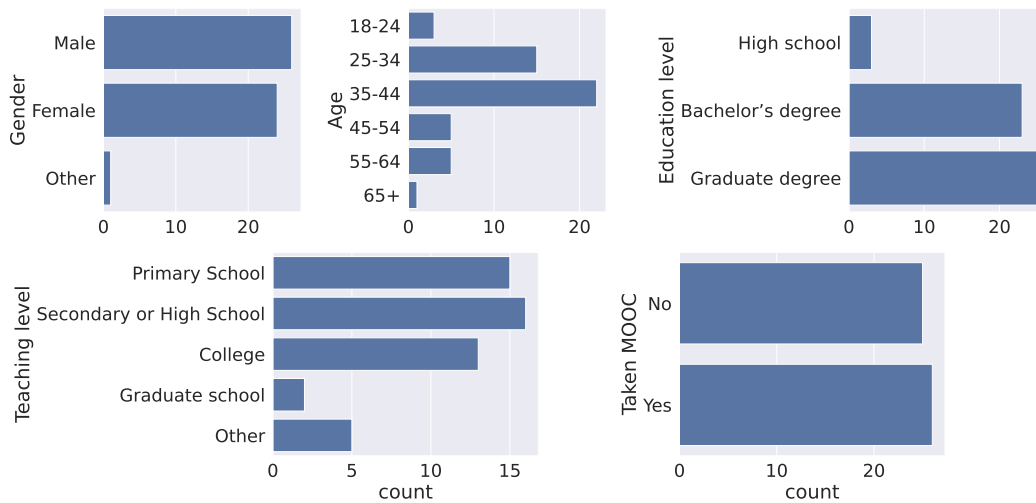


Figure 5: Demographics of teachers that participated in the user study.

Dear participant,

Thank you for participating in our study on model explanations. We are very grateful for your participation and your invaluable insight. Please read this Explanatory Statement in full before moving forward. If you would like further information regarding any aspect of this project, please contact us via the email address provided below.

We are a group of researchers from the ML4ED Laboratory at EPFL dedicated to improving education through technology. **The goal of this study is to evaluate different explanations to help a student improve their performance in an online course.**

- This survey has been approved by the EPFL Human Research Ethics Committee (HREC) under application number HREC 065-2022/27.09.2022. HREC reviews research proposals involving human participants to ensure that they are ethically acceptable.
- All the personal information will be kept confidential and anonymized. Only demographic information is being recorded and will only be reported as aggregate in a way that prevents identification of any individual participant. You can freely withdraw at any time and any collected data you provided so far will be destroyed.
- All data will be collected and stored safely and reported in an anonymous form, in accordance with the Swiss Federal law on data protection (“Loi fédérale sur la protection des données” – RS 235.1).
- Only anonymized or aggregated data may be used in follow-up research (subject to ethics approval), and made available to other researchers for further analysis and for verification of the conclusions reached by the research team.
- Only the principal investigator and the aforementioned researchers have access to the original data under strict confidentiality. Results from the project may be published in conference papers and/or journal articles. In any case, no personal data will be published (only aggregated, anonymous and/or anonymized data will be published).
- Personal data of participants will be stored for 5 years from the date of collection. During this time, participants have the right to access their data and request information about the processing of their personal data. In order to exercise this right, you need to contact the Principal Investigator.

By participating in this survey, you agree that your data can be used for scientific purposes.

In the following study, you will be asked to compare explanations for approximately 35 minutes. Please ensure that you have enough time to finish the study correctly. Unfinished or only partially answered studies will not be considered as taken part.

We ask you to approach the questions and exercises with seriousness and to complete them to the best of your ability. We will subsequently check questionnaires for seriousness and will have to discard questionnaires that do not meet this requirement.

Thank you for your help. If you encounter any problem with the survey, or if you want to give extra feedback, or receive additional information, feel free to contact us (vinitra.swamy@epfl.ch).

D.2 Content of the study

First, we explain the setting of the study to the participants with the following introductory message:

You are a teaching assistant helping with a Massive Open Online Course (MOOC). This course is taught at the Master’s level with quizzes and video lectures taking place over 10 weeks. Since it’s a difficult course with a low passing rate (23.1%), the teaching team wants to help students who are not doing well to perform better in the course by giving them personalized assistance, and encourage students who are already performing well to continue.

To do this, we have models to predict student success or failure using various weekly behavior features, such as number of video clicks or how accurately questions are answered on the weekly quizzes. If potential failure is predicted early (in our case, in the first 4 weeks of the course), we can use the explanation of the prediction to give additional support (i.e. specific tutoring or assignment reminders) to help the student pass the course.

We train four interpretable machine learning models. Each model predicts a student’s performance at the end of the course, in the form of “pass” or “fail”, but also gives us which factors contribute to student success or failure. We want to compare these explanations according to several criteria:

- Usefulness: This explanation is useful to understand the prediction.
- Trustworthiness: This explanation lets me judge if I should trust the model.
- Actionability: This explanation helps me know how to give feedback to the student.
- Completeness: This explanation has sufficient detail to understand why the prediction was made.
- Conciseness: Every detail of this explanation is necessary.

We randomly sample 4 students from the test set. Among the 4 selected students, 3 failed the class and one passed. For each student, we predict their success or failure with each model and generate an explanation. We provide them to the participants along with the models’ prediction of the student’s success or failure. The ground truth (student’s performance) and the models’ performance are not provided to the participants so that we do not bias their assessment.

The content of the explanation obtained by each method differs greatly. We simplify the explanations and render them in textual and graph format to make them as easy to understand as possible to a human. For InterpretCC explanations, we provide the full list of single features / feature groups used by the models. For Feature Gating, we show the evolution of the student’s behavior across weeks for each of the features used (see Figure 6). For Group Routing, we compute a generic score for each concept used by the model, by averaging the normalized behavior features that compose the concept. We show the evolution of that concept measure across weeks for the student. We also provide the definition of the concept and of the features that compose it (see Figure 7). For SENN, we select the top 5 groups of students, that we call *concepts*. We showcase all 180 feature-weeks for each concept along with their importance in that concept, highlighting the salient ones. We also provide the importance of each concept for the model’s prediction (see Figure 9). Finally, NAM

assigns feature importance to all 180 feature-weeks used to make the prediction. We select the 5 feature-weeks found to have highest positive impact, 5 feature weeks with the lowest impact, and the 5 feature-weeks with highest negative impact. We show their importance in a barplot (see Figure 8).

Note that the choice we made for the presentation of the explanations might have an influence on the participants' perception of the explanations. In an ideal setting, we would provide a very detailed description of how each model uses the features and how the explanation is derived, so that the participants can fairly assess the explanation's quality. However, in a realistic setting, the user facing the explanation might not have the time or prior knowledge necessary to understand these elements. To tackle that trade-off between ensuring thoroughness and accessibility of the content of each model's explanation, we opted for a balanced approach in presenting each explanation with a simple graph and an explanatory text that we kept as short as possible.

For each new sample (student taking the course), we provide the list of 4 model explanations in random order. We ask participants to compare these explanations according to five criteria using likert scales. A screenshot of the answer section is shown in Figure 10. We include a practice question to train the participants in how to answer the study and filter inattentive experts. We excluded from the analysis 5 participants who failed to answer correctly to half of the practice questions.

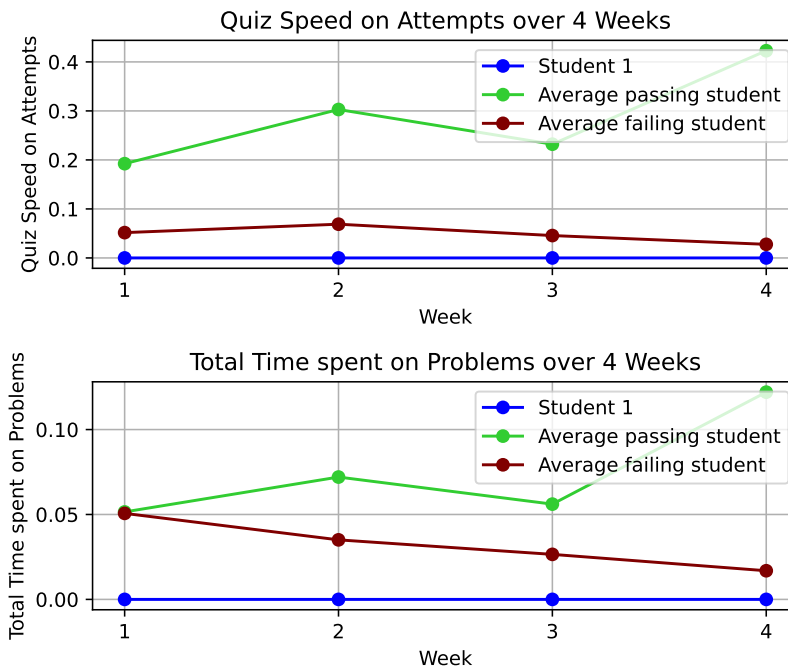


Figure 6: Explanation visualizations given to the participants for the InterpretCC Feature Gating method, for one student.

Along with the InterpretCC FG graphs, we provide the following explanatory text:

This student is predicted to fail the course. **The model found the following 2 features to be the most predictive for this student, and only used these features to make the prediction:**

- Quiz Speed on Attempts: The average time passed between two consecutive attempts for the same quiz.
- Total Time spent on Problems: The total (cumulative) time that a student has spent on problem events.

The plot shows the evolution of the student's behavior for each feature across the 4 weeks.

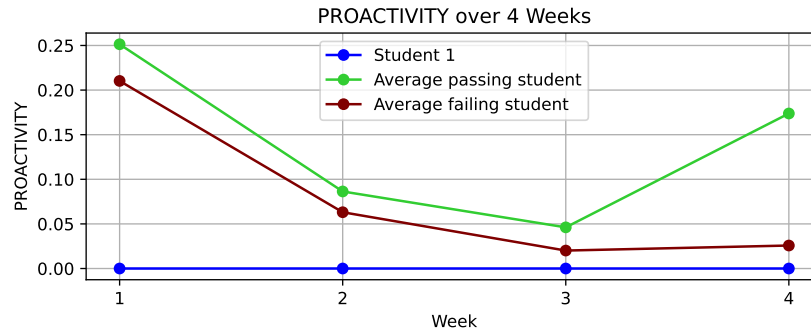


Figure 7: Graph given to the participants for the InterpretCC Group Routing with pattern-based matching method, for one student.

Along with the InterpretCC GR graph, we provide the following explanatory text:

This student is predicted to fail the course. **For this student, the model decided to only use 1 group(s) of features to make the prediction as this was the one(s) it found most important. The groups of features were designed by expert literature in learning sciences:**

- Concept: PROACTIVITY

Definition: Proactivity measures the extent to which students are on time or ahead of the schedule, as engagement in pre-class activities has shown to be associated with exam performance.

- It includes the following features: The number of videos covered by the student from those that are in subsequent weeks and The average delay in viewing video lectures after they are released to students.

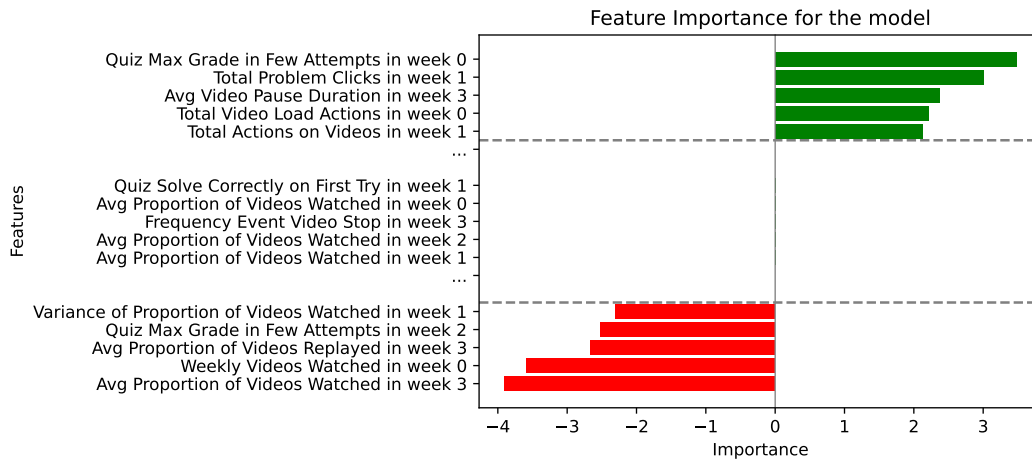


Figure 8: Graph given to the participants for the NAM method, for one student.

Along with the NAM graph, we provide the following explanatory text:

This student is predicted to fail the course. **The model used all 180 feature-weeks (45 features from 4 weeks) to make the prediction. It has assigned a level of importance for each feature-week, showing how much it impacts the predictions, independently of the student's behavior.** Out of the 180 feature-weeks, the plot shows the 5 feature-weeks found to have the highest positive impact, 5 feature weeks with the lowest impact, and the 5 feature-weeks with the highest negative

impact. For example, Quiz Max Grade in Few Attempts in week 0 has an importance score of 3.48.

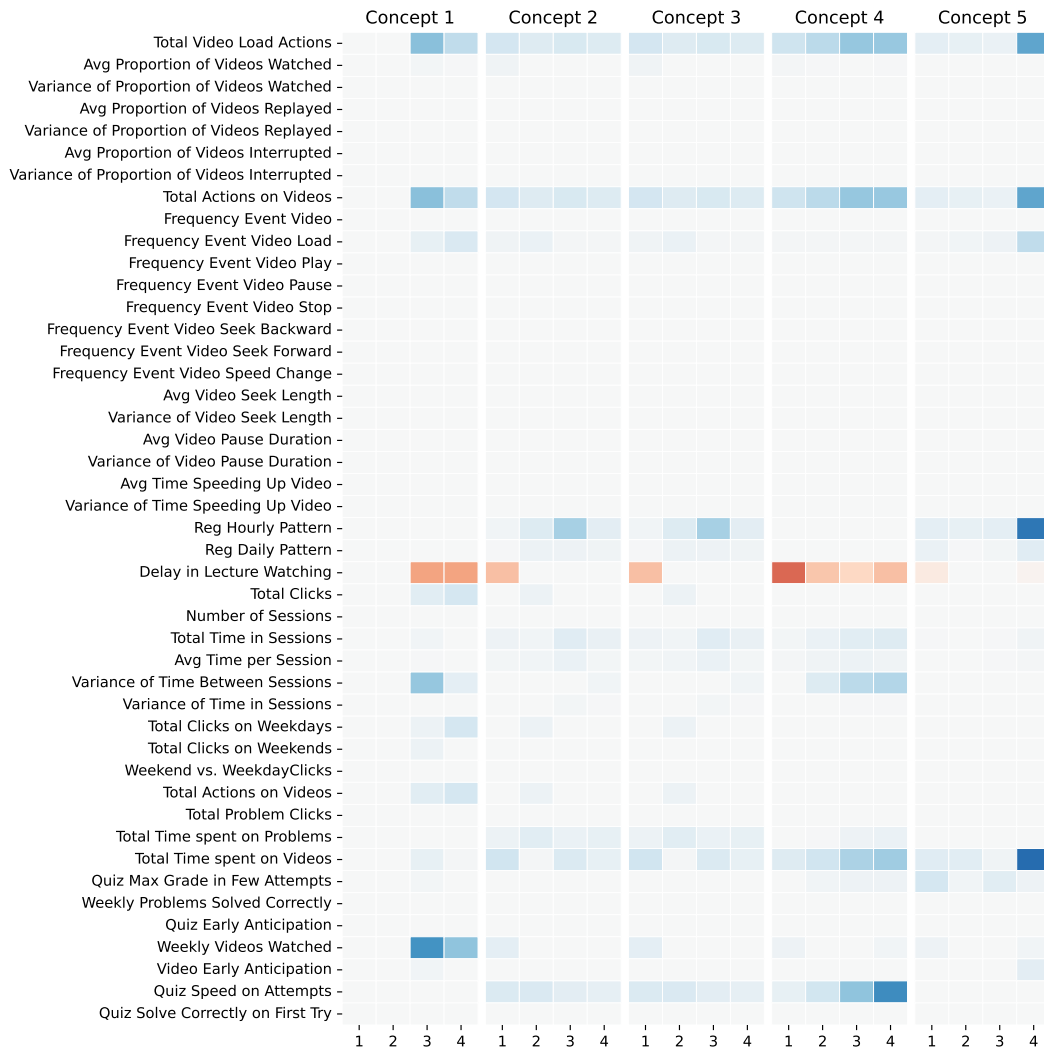


Figure 9: Graph given to the participants for the SENN method, for one student.

Along with the SENN graph, we provide the following explanatory text:

This student is predicted to fail the course. **The model used all 180 feature-weeks (45 features from 4 weeks) to make the prediction. It groups them into 5 concepts automatically and assigned a score to each concept.** Each concept can be interpreted as a group of features that are important for the prediction. The plot shows the importance of each feature-week for each concept, from high positive importance (blue) to high negative importance (red).

Here is the list ordered by absolute value of importance:

- Concept 3 (-0.21)
- Concept 2 (0.08)
- Concept 1 (-0.01)
- Concept 5 (-0.01)
- Concept 4 (0.01)

| | EXP 1 | | | | | EXP 2 | | | | | EXP 3 | | | | | EXP 4 | | | | |
|-----------------|-----------------------|-----------------------|-----------------------|-----------------------|-----------------------|-----------------------|-----------------------|-----------------------|-----------------------|-----------------------|-----------------------|-----------------------|-----------------------|-----------------------|-----------------------|-----------------------|-----------------------|-----------------------|-----------------------|-----------------------|
| | 1 | 2 | 3 | 4 | 5 | 1 | 2 | 3 | 4 | 5 | 1 | 2 | 3 | 4 | 5 | 1 | 2 | 3 | 4 | 5 |
| Usefulness | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |
| Trustworthiness | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |
| Actionability | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |
| Completeness | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |
| Conciseness | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |

Figure 10: Format of the Likert scale question asked for each criterion and explanation.

D.3 Significance Testing for User Study Results

We perform an ANOVA to determine the effect of the model and the sample (the student) on the score given by the participants, for each criterion and on average. Table 10 shows the p-values testing the significance of the effect of the model and sample on the participants' scores for each criterion. It can be interpreted the following way. In the first row, if the p-value is lower than the significance level (0.05), then there is a statistically significant difference in scores across the different models for that criterion. In the second row, if the p-value is below 0.05, there is a statistically significant difference in scores across the different students. Finally, a p-value lower than 0.05 in the final row shows that the effect of the model on the scores depends on the student. According to the table, the model has always a significant impact of the value given to each criterion. Then, we apply Tukey's Honest Significant Difference (HSD) Test to determine, for each pair of explanation, if their scores are significantly different (Figure 11).

| | Usefulness | Trustworthiness | Actionability | Completeness | Conciseness | Global |
|---------------|------------|-----------------|---------------|--------------|-------------|--------|
| Model | 0.000 | 0.001 | 0.000 | 0.000 | 0.000 | 0.000 |
| Student | 0.004 | 0.018 | 0.144 | 0.003 | 0.075 | 0.075 |
| Model:Student | 0.143 | 0.058 | 0.063 | 0.000 | 0.171 | 0.171 |

Table 10: ANOVA results for the user study.

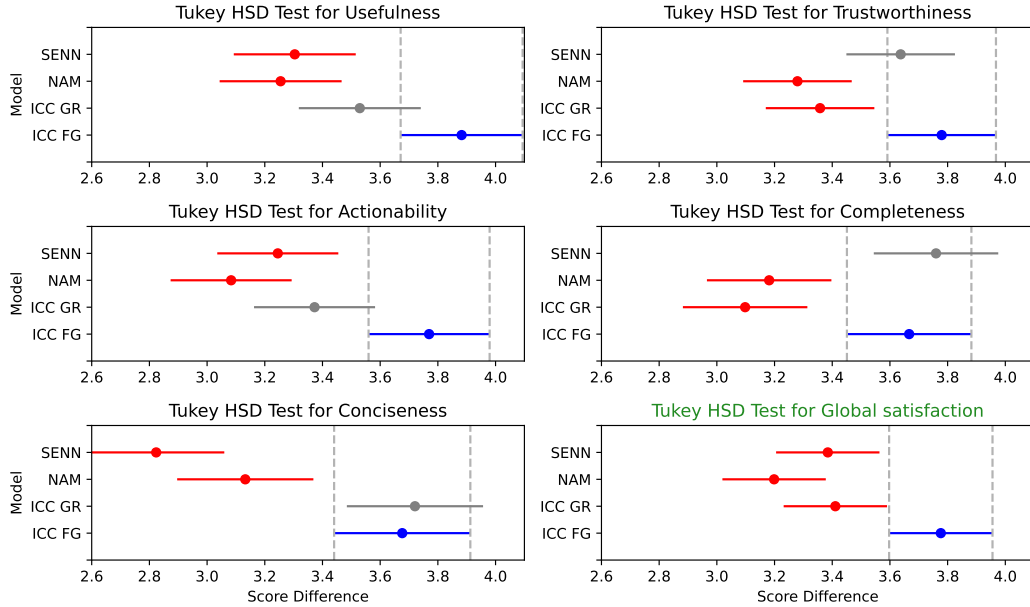


Figure 11: Significance test using Tukey’s Honest Significant Difference (HSD) Test. It indicates which pairs of models have significantly different means. We highlight the top model on average (*Global satisfaction*, in the last plot), ICC FG (InterpretCC Feature Gating) in blue, and the models that are significantly worse according to each criterion in red. Example of interpretation: for the criterion *Usefulness*, ICC FG is scored significantly better than SENN and NAM, but is not significantly better than ICC GR.

E Reproducibility and Compute Details

All EDU, Health, and Synthetic experiments were run on a single NVIDIA A100 GPU with 32 GB Memory, each model taking approximately 20 minutes (or often less) to train. Text experiments for InterpretCC took approximately an hour to train. SENN Concepts took the longest time of all models, with each model running within 3-4 hours.

All interpretable models reported in 2 and 3 has been hyperparameter tuned over the following parameters with early-stopping (where applicable):

- learning rate: 1e-3, 1e-4, 2e-5, 1e-5
- layer size: 16, 32, 64
- number of concepts (only for SENN): 5, 6, 7
- batch size: 8, 16, 32, 64)
- gumbel softmax threshold (only for InterpretCC): 0.1, 0.3, 0.5, 0.7

Each Feature-Based model was run for 100 epochs with early stopping, and each Concept-Based model was run for 20 epochs with early stopping. Other details related to preprocessing and thresholds are included directly in our repository.

F Gumbel Softmax trick and its application to InterpretCC

To make the feature gating and routing architectures compatible with backpropagation, we need to make the masks differentiable. These discrete decisions can be trained end-to-end using the Gumbel Softmax trick [31]. This method adapts soft decisions into hard ones while enabling backpropagation, i.e. provides a simple way to draw samples from a categorical distribution.

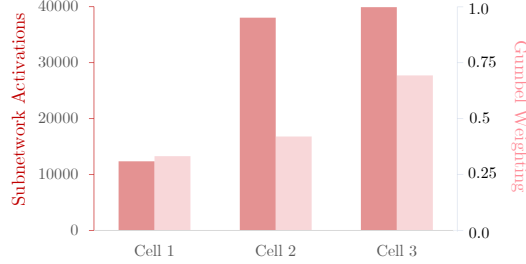


Figure 12: **Breast Cancer**: Number of activations (left) vs. average activation weight (right) per cell.

Given a categorical distribution with class probabilities $\pi = [\pi_1 \pi_2 \dots \pi_N]$, one can draw discrete samples z as follows:

$$z = \text{ONEHOT} \left(\arg \max_i [g_i + \log \pi_i] \right)$$

where $g_1 \dots g_N$ are i.i.d. samples drawn from the Gumbel(0, 1) distribution. Then, the softmax function is used as a differentiable approximation to $\arg \max$ to generate a N -dimensional sample vector y such that

$$y_i = \frac{\exp((\log(\pi_i) + g_i)/\tau)}{\sum_{j=1}^N \exp((\log(\pi_j) + g_j)/\tau)} \quad \text{for } i = 1, \dots, N$$

where τ is a softmax temperature parameter that is fixed at $\tau = 1$ for experiments in this project.

Notice that for the gating mechanism, an independent sample is drawn for each ‘gate’ instead of for each datapoint in routing. For example in feature gating, for each feature i , a soft-decision $a_i \in (-\infty, +\infty)$ is outputted by the discriminator layers. The probability π_1 that the feature should be activated as well as the complement probability π_2 (feature is not activated) can then be computed by using the sigmoid function:

$$\pi_1 = \sigma(a_i) \quad \pi_2 = 1 - \pi_1 = 1 - \sigma(a_i)$$

The corresponding (1-dimensional) sample y for each i can thus be reduced to

$$y = \sigma \left(\frac{a_i + g_1 - g_2}{\tau} \right)$$

In other words, the discriminator layers from Fig. 1 actually feed into an adapted Gumbel Sigmoid where σ_i is the corresponding y sample as described above.

For routing, the discriminator layers actually output the route logits to a Gumbel Softmax, which constructs the categorical sample vector (of dimension equal to the number of routes and i -th entry y_i defined as above).

Finally, we can use a straight-through estimator during training. In other words, binary (or hard/quantized) samples are then used for the forward pass while gradients are obtained from the soft samples for backpropagation. This means that, given soft decisions σ_i , architectures that use a mask $M = [m_1 \dots m_N]$ with $m_i = \mathbb{1}_{\{\sigma_i \geq 0.5\}}$ differ in value during the forward and backward pass:

$$m_i = \begin{cases} \mathbb{1}_{\{\sigma_i \geq 0.5\}} & \text{(forward pass),} \\ \sigma_i & \text{(backward pass)} \end{cases}$$

G Additional Group Routing Experiments

We conduct additional analyses for the high impact, real-world applications in Health and EDU.

G.1 Grouping Analysis for Breast Cancer Dataset

For the **Breast Cancer** data set, the subnetworks grouping features from Cell 1 and Cell 2 are activated much more frequently than the third subnetwork (see Fig. 12). Furthermore, Cell 3 also gets activated with higher weights than the other two cells (Cell 1: 0.25, Cell 2: 0.40, Cell 3: 0.70). Smoothness and texture of the tissue images were the most important features across cells.

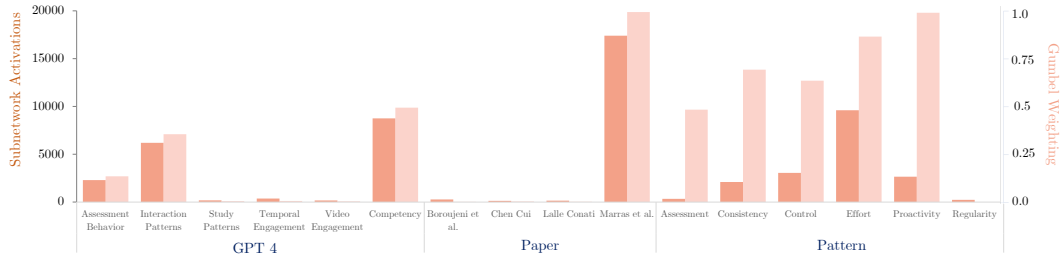


Figure 13: **EDU**: Number of subnetwork activations (left) and Gumbel Softmax activation weights (right) across different groupings (GPT-4, Paper, Pattern) for course DSP 1 of the EDU domain.

G.2 Grouping Analysis for EDU Datasets

We conduct sparsity and group routing paradigm analyses on a representative course of the EDU datasets (DSP). We additionally experiment with grouping by paper and pattern over all courses, including two additional MOOC courses with low predictive performance (Structures and Ventures) featured in [45].

G.2.1 Sparsity Analysis on DSP

We conduct a sparsity analysis on the course used for the user study, DSP, and featured in several research works in the learning science community [6, 45, 33]. In line with [45], we evaluate the sparsity at both the 40% and 60% prediction setting, representing early prediction at 4 and 6 weeks of the course respectively.

| Setting | Model | Performance | | | Activated Features | |
|---------|--------------------|-------------|-------|-------|--------------------|----------|
| | | ACC | BAC | AUC | μ | σ |
| 40% | Baseline | 0.835 | 0.653 | 0.858 | 97 | 0 |
| | L1 (1e-5) | 0.790 | 0.711 | 0.817 | 8.54 | 0.58 |
| | L1 (1e-4) | 0.763 | 0.748 | 0.784 | 2.20 | 1.64 |
| | Annealed MSE | 0.768 | 0.770 | 0.823 | 13.26 | 3.72 |
| | Truncated AMSE | 0.805 | 0.743 | 0.823 | 5.95 | 1.03 |
| 60% | Baseline | 0.944 | 0.925 | 0.982 | 97 | 0 |
| | L1 (1e-5) | 0.914 | 0.935 | 0.963 | 37.60 | 3.97 |
| | L1 (1e-4) | 0.914 | 0.917 | 0.957 | 31.38 | 4.45 |
| | Annealed L1 (1e-5) | 0.910 | 0.927 | 0.957 | 38.36 | 3.61 |
| | Annealed MSE | 0.892 | 0.927 | 0.952 | 21.45 | 3.23 |
| | Truncated AMSE | 0.787 | 0.788 | 0.839 | - | - |

Table 11: **InterpretCC Feature Gating** comparison of performance metrics between different sparsification methods next to their average and standard deviation of number of activated features for both 40% and 60% early success prediction settings; baseline benchmarks also provided for contrast.

An annealed mean-squared regularization proved most effective, although it activated more features on average than L_1 -norm regularization which more effectively reduced the feature space while achieving desirable balanced accuracy performance in this setting. By truncating the initial feature space to only the activated features and using the same architecture, performance is almost maintained even though the average number of activated features per datapoint is more than halved (from around 13 to 6). However, this method was not as effective for the 60% setting. Truncating the feature space largely reduces predictive capability (e.g. almost a 0.15 drop in balanced accuracy). L_1 regularization in this case proved best. Using an annealed regularization did not significantly improve or change model performance as well.

G.2.2 Multiple Grouping Paradigms on DSP

To illustrate the influence of different feature groupings, we conduct a deep dive for course DSP 1 of the EDU domain. Figure 13 illustrates the number of subnetwork activations and corresponding weights for three different groupings.

For the first two groupings (GPT-4, Paper), the subnetwork activations (number of times the route was activated) closely mirror the Gumbel Sigmoid adaptive weighting, indicating that a few networks are activated with high weights for prediction. In the *group by GPT-4* setting, we see behaviors of competency, interaction patterns, and assessment frequently activated for student pass-fail predictions. Although ‘interaction patterns’ is the largest category (most number of features chosen by GPT-4), it still comes second to competency (focusing on student achievement). In the *group by paper* setting, we see a clear preference for Marras et al. with over 17, 500 students predicted using this network (dark orange) and high weight given to the predictions from the network (light orange).

In contrast, in the third grouping (Pattern), we see a differentiation between the number of activations (dark orange) and the weight of the activations (light orange). Notably, the patterns of ‘Effort’, ‘Proactivity’, ‘Consistency’ and ‘Control’ all have higher than 50% weight when they are activated, which means they contribute a lot to the overall prediction when chosen.

G.2.3 Grouping by paper and pattern over all courses

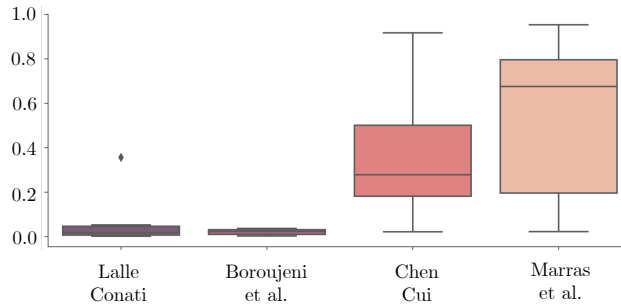


Figure 14: EDU analysis of *group by paper* routing averaged over all six courses for each paper grouping. The y-axis represents the proportion of points for which the subnetwork is activated.

In Figure 14, we see InterpretCC routing by research paper (grouping the features based on the paper they were proposed in). The Marras et al. and Chen Cui feature sets have clearly been identified as important over the majority of courses, echoing findings in other learning science literature using BiLSTM and random forest architectures [35, 34, 55]. The large standard deviations in the box-plots indicate that for at least some courses (in this case Structures and Venture), Chen Cui and Marras were not found significantly important. Notably, the same courses that have low accuracies on routing in 2 are those that have low scores on the two most popular feature sets, showing a consensus among performant InterpretCC models and a validation of the identification of importance.

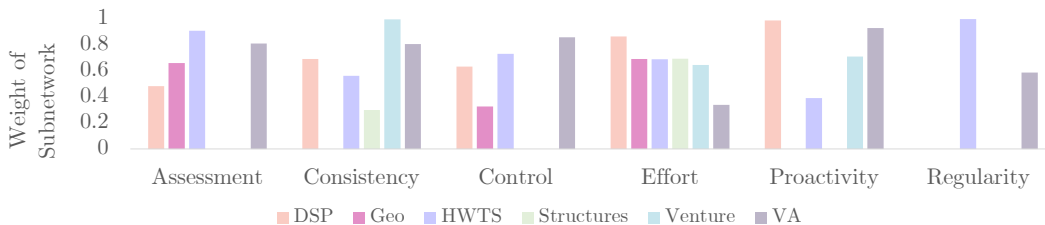


Figure 15: EDU Weighting distribution across subnetworks using *group by pattern* for 6 courses.

In Fig. 15, we see a widely varying distribution of patterns selected across courses, showcasing the ability of InterpretCC to adaptively select subnetwork weights depending on the dataset.