
Supplementary for Interpretable and Personalized Apprenticeship Scheduling: Learning Interpretable Scheduling Policies from Heterogeneous User Demonstrations

Anonymous Author(s)

Affiliation

Address

email

1 Additional Experiment Domain Details

2 **Synthetic Scheduling Environment** The synthetic scheduling environment represents one of the
3 hardest scheduling problems. In this environment, two agents must complete a set of 20 tasks which
4 have upper- and lower-bound temporal constraints (i.e., deadline and wait constraints), proximity
5 constraints (i.e., no two agents can be in the same place at the same time), and travel-time constraints.
6 For the purposes of apprenticeship learning, an action is defined as the assignment of an agent to
7 complete a task presently. The decision-maker must decide the optimal sequence of actions according
8 to the decision-maker’s own criteria. For this environment, we construct a set of heterogeneous, mock
9 decision-makers that schedule according to Equation 1.

$$\tau_i^* = \arg \max_{\tau_j \subset \tau_S} (\rho_1 H_{EDF}(\tau_j) + \rho_2 H_{dist}(\tau_j) + H_{ID}(\tau_j, \rho_3)) \quad (1)$$

10 In this equation, our decision-maker selects a task τ_i^* from the set of tasks τ_S . The task-prioritization
11 scheme is based on three criteria: H_{EDF} prioritizes tasks according to deadline (i.e., “earliest-
12 deadline first”), H_{dist} prioritizes the closest task, and H_{ID} prioritizes tasks according to a user-
13 specified highest/lowest index or value based on ρ_3 (i.e., $\rho_3(j) + (1 - \rho_3)(-j)$). The heterogeneity
14 in decision-making comes from the latent weighting vector $\vec{\rho}$. Specifically, $\rho_1 \in \mathbb{R}$ and $\rho_2 \in \mathbb{R}$
15 weight the importance of H_{EDF} and H_{dist} , respectively. $\rho_3 \in \{0, 1\}$ is a mode selector in which
16 the highest/lowest task index is prioritized. By drawing $\vec{\rho}$ from a multivariate random distribution,
17 we can create an infinite number of unique demonstrator types. This adapted environment differs
18 from the synthetic, low-dimensional environment in that there are a rich set of temporal, spatial, and
19 agent-based constraints modeling the job-shop scheduling problem; furthermore, the parameters of
20 the demonstrator’s decision-making process is hidden and comprised of one discrete factor and two
21 continuous factors. In this domain, counterfactuals are generated by consider specific task information
22 such as availability, distance from agent, prerequisites satisfied.

23 **Real-world Data: Taxi Domain** Our environment has three locations: the village, the airport, and
24 the city. The taxi driver has the objective of picking up a passenger from the city or village. There is
25 always a passenger at the city, but the taxi driver may encounter up to 60 minutes of traffic going into
26 the city. There may be a wait time of up to 60 minutes to pick up a passenger at the village; however,
27 there is no traffic on the way to the village, and the wait time is unknown to the taxi driver unless she
28 is at the village. A dataset of 70 human-collected tree policies to solve this task (given with leaf node
29 actions such as “Drive to the City”, “Drive to the Airport”, and “Wait for Passenger”, and decision
30 node criterion depending on the amount of wait time, traffic time, and current location) are used to
31 generate heterogeneous trajectories. We originally collect 98 tree-based policies through an IRB-
32 approved study. However, 28 of these do not produce successful trajectories. The tree-based policies

Table 1: Apprenticeship Performance in Imitating Robot Kinesthetic Ping Pong Demonstrations.

Environment	Our Method	Sammut et al.	Nikolaidis et al.	Tamar et al.	Hsiao et. al.	InfoGAIL Li et. al.	Gombolay et. al.
Ping-Pong	59.60%	18.14%	31.20%	26.17%	17.96%	36.70%	28.60%

33 can be found in this GitHub repository [https://github.com/Personalized-Neural-Trees/Interpretable-](https://github.com/Personalized-Neural-Trees/Interpretable-and-Personalized-Apprenticeship-Scheduling-Learning-Interpretable-Scheduling-Policies)
 34 [and-Personalized-Apprenticeship-Scheduling-Learning-Interpretable-Scheduling-Policies](https://github.com/Personalized-Neural-Trees/Interpretable-and-Personalized-Apprenticeship-Scheduling-Learning-Interpretable-Scheduling-Policies).

35 **Kinesthetic Robot Table Tennis** We collected a real-world data set consisting of 10 human demon-
 36 strators kinesthetically presenting four different table tennis strikes on a Rethink Robotics Sawyer.
 37 The table tennis strike variants consisted of push, topspin, slice, and lob and were conducted using
 38 a forehand motion, giving four different categories of motion. While our approach is primarily for
 39 discrete classification problems, such as decision making, it can naturally be extended to complex
 40 continuous domains, such as low-level robot joint control.

41 To collect data, we first show each demonstrator a sample video of the table tennis strike and allow
 42 them to practice until she feels confident that she can return the ping pong ball over the net. Then, we
 43 reset the robotic arm to a preset initial position and allow the demonstrator to strike a ping pong ball
 44 launched from an automatic ball launcher. Throughout the demonstration, we record the position of
 45 the end-effector.

46 **Survey Scheduling Environment** This domain describes a **ND[ST-SR-TA]** scheduling domain
 47 defined by Korsah [5]. In this domain, synthetic schedulers are given utilities of three tasks where
 48 utility $U \in \{1, 2, 3\}$ and must choose the highest or lowest task index based on a pre-specified latent
 49 decision-making criteria. We generate a set of 100 schedules (each of length 20) from heterogeneous
 50 demonstrators.

51 2 LfD Performance in Kinesthetic Robot Table Tennis

52 Here, we show that Personalized Neural Trees can easily be extended to a variety of domains,
 53 increasing the data-efficiency, accuracy, and utility of learning-from-demonstration with multiple
 54 human demonstrators. We demonstrate this by using a PNT to learn kinesthetic robot table tennis
 55 demonstrations in Table 1. We received 40 demonstrations across 10 demonstrators, representing
 56 four different table tennis strikes. To clean the trajectory of the end effector prior to learning, we
 57 transformed our trajectories into a transformed three-dimensional space using Principal Components
 58 Analysis. Our data was then labeled by selecting the principal component in which the end effector
 59 moved most at each timestep ($|\mathcal{A}| = 6$). As seen in row 4 of Table 1, our approach outperforms all
 60 other benchmarks.

61 3 Sensitivity Analysis of PNTs

62 To analyze the sensitivity of our framework, we use our synthetic scheduling environment and perturb
 63 the amount of data available to the PNT and the amount of noise (correctness) within the data. To
 64 provide a thorough analysis, we validate our approach using k -fold cross-validation. This entails both
 65 choosing a different subset of data to learn from and perturbing different truth-values of state-actions
 66 pairs each fold.

67 As shown in Figure 1, our PNT is reasonably robust to noise for 2, 5, and 15 schedules as there is
 68 not a steep drop in accuracy. We do not see the typical trend where the effect of noise deteriorates
 69 as the amount of data increases. We posit the cause of this deviation as follows: As the number of
 70 demonstrators increases, the embedding space Ω of the PNT tends to represent a richer distribution.
 71 While the heterogeneity among the demonstrators may remain constant (represent the same number
 72 of modes), cases in which the PNT is unable to tease out the demonstrator mode from a single
 73 schedule are more likely (due to the increase in the number of schedules), leading to an embedding
 74 distribution with higher variance. Without noise, the PNT is able to make sense of the embedding
 75 space and learn with high performance; as the amount of noise increases, it is likely more difficult to

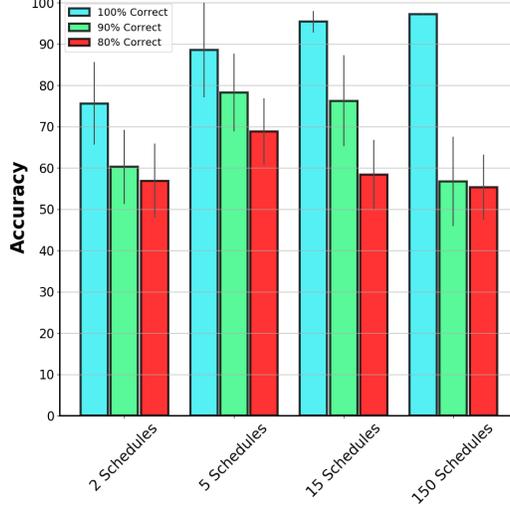


Figure 1: Sensitivity analysis in the synthetic scheduling environment.

76 represent demonstrators compactly within the embedding space. We posit that this increased variance
 77 within the embedding space caused by the combined effect of an increased number of demonstrators
 78 and noise leads to a reduction in performance when noise is held constant and the amount of data
 79 increases.

80 As expected, as the number of schedules increase, the PNTs have higher accuracy. However, from
 81 15 to 150 schedules (a 10x magnitude increase in data), for the case of 100% correct data, there is
 82 only a $\sim 2\%$ increase in accuracy. This result provides support to the claim of data-efficiency in our
 83 apprenticeship scheduling framework.

84 4 Evidence Lower Bound

85 Here, we present the full derivation of the evidence lower bound (ELBO) that is used maximize the
 86 mutual information between ω and trajectories τ .

$$\begin{aligned}
 G(\omega; \tau) &= H(\omega) - H(\omega|\tau) & (2) \\
 &= \mathbb{E}_{\omega \sim P(\omega), a_p^t \sim f_{\theta|\omega}^{PNT}} [\log P(\omega|s_p^t, a_p^t)] + H(\omega) \\
 &= \mathbb{E}_{a \sim f_{\theta|\omega}^{PNT}} [D_{KL}(\log(P(\omega_p|s_p^t, a_p^t)) || \log(q_{\zeta|\theta}^\omega(s_p^t, a_p^t)))] + \mathbb{E}_{\omega \sim P(\omega)} \log(q_{\zeta|\theta}^\omega(s_p^t, a_p^t)) + H(\omega) \\
 &\geq \mathbb{E}_{\omega_p \sim \mathcal{N}(\bar{\mu}_p, \bar{\sigma}_p^2), a \sim f_{\theta|\omega}^{PNT}} [\log(q_{\zeta|\theta}^\omega(\omega_p|s_p^t, a_p^t))] + H(\omega) = L_G(f_{\theta|\omega}^{PNT} || q_{\zeta|\theta}^\omega)
 \end{aligned}$$

87 In our approach, we make use of continuous personalized embeddings which allow for greater
 88 expressivity in the embedding space, Ω . As such, we utilize a mean-squared error (MSE) loss
 89 between a sample from the approximate posterior (modeled as a normal distribution with constant
 90 variance) and the current embedding.

91 We present the approximate normal distribution, $\mathcal{N}_{q_{\zeta|\theta}^\omega}$, in Equation 3, where ω is the mean outputted
 92 by the posterior network, and σ is the standard deviation.

$$\mathcal{N}_{q_{\zeta|\theta}^\omega} = \frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{1}{2} \frac{(x-\omega)^2}{\sigma^2}} \quad (3)$$

93 **Theorem 4.1.** *Minimizing the mean-squared error between a sample from the approximate posterior*
 94 *and the current embedding is equivalent to maximizing the log-likelihood and therefore, the evidence*
 95 *lower bound.*

96 *Proof.* The mean-squared error (MSE) loss is $(x - \omega)^2$, where ω is the sample from the approximate
 97 posterior, and x is the current personalized embedding used to generate the predicted action. This is
 98 equivalent to the exponent numerator in $\mathcal{N}_{q_{\zeta|\theta}^\omega}$. With constant variance, the exponential function is

99 monotonic, and thus, minimizing the exponent will maximize the likelihood of the posterior. Thus,
 100 minimizing the MSE is equivalent to maximizing the likelihood of the posterior. This naturally
 101 extends to the multivariate case. \square

102 5 Interpretability User Study Details

103 Here, we present the details of our novel user study to assess the interpretability of our discretized
 104 PNTs. We design an online questionnaire that asks users to predict a task to schedule given an input
 105 using a decision-tree based method and a neural-network-based method. Each user is randomly
 106 assigned a reasoning level, standard, pointwise, or counterfactual. Standard and counterfactual
 107 reasoning are discussed in the main paper. Pointwise reasoning outputs a probability of taking a
 108 certain action given a feature vector describing that action, x_a^t from state s^t , and a contextual feature
 109 vector capturing features common to all actions \bar{x}^t . We can generate pointwise features through
 110 Equation 4.

$$z^{t,p} := [\omega_p, \bar{x}^t, x_a^t], y_a^t = 1 \quad (4)$$

$$z^{t,p} := [\omega_p, \bar{x}^t, x_{a'}^t], y_{a'}^t = 0 \quad (5)$$

111 The tree and neural network-based models were trained under minimal sizes that were capable of
 112 achieving near-perfect accuracy. Tree models are learned PNTs, which are then discretized. The NN
 113 models are generated by appending personalized embeddings to a NN and following the training
 114 methodology described in Algorithm 1 from the main paper. Then, comparison weights and model
 115 weights for the discrete trees and neural networks, respectively, were rounded to the nearest 0.25.
 116 Rounding yielded $\sim 2\%$ loss in accuracy but allowed for the survey to be conducted within a
 117 reasonable time. For each type of decision-making framework, we provide instructions for how to
 118 utilize the framework to make a prediction. The order in which the user completes the neural network
 119 portion and decision tree portion is randomized. We explore additional hypothesis: counterfactual
 120 tree-based decision-making models are more interpretable (**H4**), quicker to validate (**H5**), and are
 121 more easily utilized (**H6**) than neural-network based models of any reasoning level. We then provide
 122 further comparisons between tree-based methods of different levels of reasoning.

123 We use four metrics throughout our user study: interpretability of the decision-making model,
 124 interpretability of the overall decision-making process, time to compute an output, and correctness.
 125 To verify **H4-H6**, we must compare the counterfactual discretized PNT to a standard neural network,
 126 pointwise neural network, and pairwise neural network. As the first case is shown in the paper
 127 (standard neural network vs. discretized PNT), we provide the results for the other two scenarios
 128 here.

129 6 Survey Results

130 Our IRB-approved anonymous survey was sent out to adult university students. We collected 35
 131 responses, 14 of standard, 11 of pointwise, and 15 of counterfactual. We filter out responses that put
 132 in nonsensical answers (i.e., letters where numbers should be and repeated answers).

133 **H4:** In comparing a NN with pointwise reasoning to a discretized PNT, we test for normality and
 134 homoscedasticity and do not reject the null hypothesis in either case, using Shapiro-Wilk ($p > 0.9$
 135 and $p > 0.3$) and Levene’s Test ($p > 0.2$ and $p > 0.3$). We perform a paired t-test and find that
 136 counterfactual tree-based models were rated statistically significantly higher than pointwise neural
 137 networks on users’ Likert scale ratings for model interpretability and overall process interpretability
 138 ($p < 0.05$ and $p < 0.01$). In comparing a NN with pairwise reasoning to a discretized PNT, we
 139 test for normality and homoscedasticity and do not reject the null hypothesis in either case, using
 140 Shapiro-Wilk ($p > 0.1$ and $p > 0.1$) and Levene’s Test ($p > 0.4$ and $p > 0.4$). We perform a paired
 141 t-test and find that counterfactual tree-based models were rated statistically significantly higher than
 142 pointwise neural networks on users’ Likert scale ratings for model interpretability and overall process
 143 interpretability ($p < 0.01$ and $p < .05$). These results support **H4**.

144 **H5:** In comparing a NN with pointwise reasoning to a discretized PNT, we perform a Wilcoxon
 145 signed-rank test on the per-model time to determine an output and find that tree-based models were
 146 not statistically significantly quicker to validate than neural networks ($p = 0.37$). In comparing a

147 NN with pairwise reasoning to a discretized PNT, we perform a Wilcoxon signed-rank test on the
148 per-model time to determine an output and find that tree-based models were statistically significantly
149 quicker to validate than neural networks ($p = 0.001$). This result provides partial support **H5**.

150 **H6:** In comparing a NN with pointwise reasoning to a discretized PNT, we perform a Wilcoxon
151 signed-rank test on the per-model time to determine an output and find that tree-based models were
152 statistically significantly achieved higher overall correctness scores compared to NN based models
153 ($p < 0.05$), supporting **H6**. In comparing a NN with pairwise reasoning to a discretized PNT, we
154 test for normality and homoscedasticity and do not reject the null hypothesis in either case, using
155 Shapiro-Wilk ($p > 0.05$) and Levene’s Test ($p > 0.2$). We perform a paired t-test and find that
156 users using tree-based models statistically significantly achieved higher overall correctness scores
157 compared to NN based models ($p < 0.001$), supporting **H6**.

158 7 Hyperparameters and Architecture Details

159 We compare our personalized apprenticeship scheduling approach to several baselines [2, 4, 6, 7, 10,
160 11]. Throughout this section, we will discuss the architecture, implementation details, and learning
161 rates for all baselines and our algorithm in each domain. The runtime mentioned is in respect to a
162 desktop with a NVIDIA RTX 2080Ti GPU and an Intel i7 processor.

163 7.1 Synthetic Low-Dimensional Environment

164 Each apprenticeship learning algorithm below is given 50 schedules to learn from and tests on a set
165 of 50 unseen demonstrations.

- 166 • For the method of Sammut et al. [10], we utilize an multi-layer perceptron (MLP) with 3
167 linear layers connected by ReLU activation functions. After the last linear layer, we utilize a
168 log softmax function to compute the log probability of which task to schedule. Each linear
169 layer has 10 hidden units. We utilize the Adam optimizer with a learning rate of $1e^{-3}$. The
170 runtime for training and verifying this model is under 30 minutes.
- 171 • For the method of Nikolaidis et al. [7], we utilize k-means clustering to separate the data into
172 two clusters. Two neural networks (one for each cluster) are trained to imitate demonstrator
173 data within the cluster. Each network utilizes the same architecture and learning rate used
174 in the baseline of Sammut et al. [10]. The runtime for training and verifying this model is
175 under 30 minutes.
- 176 • For the method of Li et al. [6], we utilize an simulator-free version of infoGAIL. The policy,
177 discriminator, and approximate posterior are modeled by MLPs with 2 linear layers (32
178 hidden units) connected by a ReLU activation function, and an output activation function of
179 a softmax, sigmoid, and softmax respectively. We initialize the number of discrete modes
180 to 2. We utilize learning rates of $1e^{-4}$, $1e^{-3}$, $1e^{-4}$ respectively. For the hyperparameters
181 of infoGAIL, we initialize λ_1 to 1, γ to 0.95, and λ_2 to 0. The runtime for training and
182 verifying this model is under 30 minutes.
- 183 • For the method of Tamar et al. [11], we utilize a neural network with 3 linear layers (10, 2,
184 2 hidden units, respectively) connected by ReLU activation functions. We use $N=5$ samples
185 as our hyperparameter to estimate the intention probability distribution $\mathcal{P}(z)$. We utilize
186 a learning rate of $1e^{-3}$ alongside Stochastic Gradient Descent (SGD). The runtime for
187 training and verifying this model is under 30 minutes.
- 188 • For the method of Hsiao et al. [4], we utilize a bidirectional LSTM with attention followed
189 by a linear layer as specified in their paper. For the decoder, we utilize three linear layers
190 connected by ReLU activation functions. We utilize a learning rate of $1e^{-3}$ alongside
191 Stochastic Gradient Descent (SGD). The runtime for training and verifying this model is
192 under 30 minutes.
- 193 • For the method of Gombolay et al. [2], we utilize a standard decision tree (counterfactuals
194 are not possible when $|A| \leq 2$) of depth 10. The runtime for training and verifying this
195 model is under 30 minutes.
- 196 • For our Personalized Neural Trees, we utilize a max depth of 6 (32 leaves) and embedding
197 dimension of 2 ($d = 2$). We set learning rates of θ to $1e^{-3}$, ω to $1e^{-2}$, and ζ to $1e^{-3}$.

198 We find empirically that setting the learning rate of ω slightly higher allows for better LfD
199 accuracy. For our approximate posterior, $q_{\zeta|\theta}^{\omega}$, we set the value of σ_p to zero. The runtime
200 for training and verifying this model is under 30 minutes.

201 7.2 Synthetic Scheduling Environment

202 Each apprenticeship learning algorithm below is given 150 schedules to learn from and tests on a set
203 of 100 unseen demonstrators.

- 204 • For the method of Sammut et al. [10], we utilize an multi-layer perceptron (MLP) with six
205 linear layers connected by ReLU activation functions. After the last linear layer, we utilize a
206 log softmax function to compute the log probability of which task to schedule. Each linear
207 layers have 128, 128, 32, 32, 32, and 20 hidden units, respectively. We utilize the Adam
208 optimizer with a learning rate of $1e^{-4}$. The runtime for training and verifying this model is
209 approximately 30 minutes.
- 210 • For the method of Nikolaidis et al. [7], we utilize k-means clustering to separate the data
211 into three clusters. Three neural networks (one for each cluster) are trained to imitate
212 demonstrator data within the cluster. Each network utilizes the same architecture and
213 learning rate used in the baseline of Sammut et al. [10]. The runtime for training and
214 verifying this model is approximately 30 minutes.
- 215 • For the method of Li et al. [6], we again utilize a simulator-free version of infoGAIL. The
216 policy follows the same network structure used in the Sammut et al. [10] baseline. The
217 discriminator and approximate posterior are modeled by MLPs with six linear layers (128,
218 128, 128, 32, 32, 32 hidden units, respectively) connected by a ReLU activation function,
219 and an output activation function of a sigmoid, and softmax respectively. We initialize the
220 number of discrete modes to 3. We utilize learning rates of $1e^{-4}$, $1e^{-3}$, $1e^{-4}$ respectively.
221 For the hyperparameters of infoGAIL, we initialize λ_1 to 1, γ to 0.95, and λ_2 to 0. The
222 runtime for training and verifying this model is approximately 24-48 hours.
- 223 • For the method of Tamar et al. [11], we utilize a neural network with 5 linear layers (128,
224 32, 32, 32, 32, 20, 2, 2 hidden units, respectively) connected by ReLU activation functions.
225 We use $N=5$ samples as our hyperparameter to estimate the intention probability distribution
226 $\mathcal{P}(z)$. We utilize a learning rate of $1e^{-3}$ alongside Stochastic Gradient Descent (SGD). The
227 runtime for training and verifying this model is approximately 3 hours.
- 228 • For the method of Hsiao et al. [4], we utilize a bidirectional LSTM with attention followed
229 by a linear layer as specified in their paper. For the decoder, we utilize six linear layers
230 connected by ReLU activation functions. We utilize a learning rate of $1e^{-3}$ alongside
231 Stochastic Gradient Descent (SGD). The runtime for training and verifying this model is
232 approximately 3 hours.
- 233 • For the method of Gombolay et al. [2], we utilize a pairwise decision tree of depth 10. The
234 counterfactuals are set to one-hot encodings of each action, as done in the original paper.
235 The runtime for generating and verifying this model is approximately 5 minutes.
- 236 • For our Personalized Neural Trees, we utilize a max depth of six (32 leaves) and embedding
237 dimension of 3 ($d = 3$). We set learning rates of θ to $1e^{-2}$, ω to $1e^{-2}$, and ζ to $1e^{-2}$. We
238 find empirically that pretraining the policy network first and then adding in the posterior
239 at a later epoch results in both good performance and mutual information maximization.
240 This is opposed to training both models at once from scratch. For our approximate posterior,
241 $q_{\zeta|\theta}^{\omega}$, we set the value of σ_p to zero. The runtime for training and verifying this model is
242 approximately 24 hours.

243 7.3 Taxi Domain

244 Each apprenticeship learning algorithm below is given 25 successful trajectories from each user and
245 tested on a set of 25 unseen trajectories from each demonstrator.

- 246 • For the method of Sammut et al. [10], we utilize the same architecture and learning rate as
247 that of the synthetic scheduling environment. The runtime for training and verifying this
248 model is approximately 30 minutes.

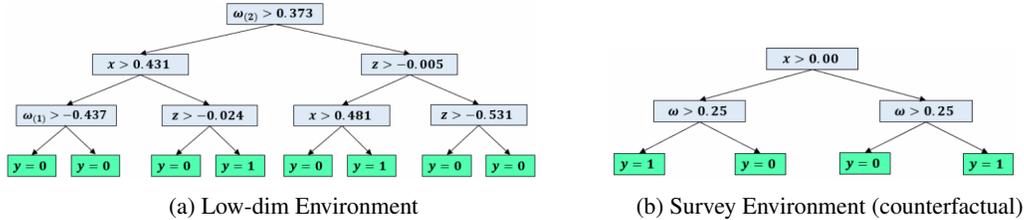


Figure 2: This figure depicts the learned PNT model after translation to an interpretable form.

- 249 • For the method of Nikolaidis et al. [7], we utilize k-means clustering to separate the data
250 into three clusters. Three neural networks (one for each cluster) are trained to imitate
251 demonstrator data within the cluster. Each network utilizes the same architecture and
252 learning rate used in the baseline of Sammut et al. [10]. The runtime for training and
253 verifying this model is approximately 30 minutes.
- 254 • For the method of Li et al. [6], we utilize the same architecture and learning rate as that of
255 the synthetic scheduling environment. The runtime for training and verifying this model is
256 approximately 24-48 hours.
- 257 • For the method of Tamar et al. [11], we utilize the same architecture and learning rate as
258 that of the synthetic scheduling environment. The runtime for training and verifying this
259 model is approximately 3 hours.
- 260 • For the method of Hsiao et al. [4], we utilize the same architecture and learning rate as that
261 of the synthetic scheduling environment. The runtime for training and verifying this model
262 is approximately 3 hours.
- 263 • For the method of Gombolay et al. [2], we utilize a pairwise decision tree of depth 13. The
264 counterfactuals are set to one-hot encodings of each action, as done in the original paper.
265 The runtime for generating and verifying this model is approximately 5 minutes.
- 266 • For our Personalized Neural Trees, we utilize a max depth of 8 (128 leaves) and embedding
267 dimension of 3 ($d = 3$). As counterfactual task information is not readily available, we
268 utilize one-hot encodings for each action. We set learning rates of θ to $1e^{-2}$, ω to $1e^{-1}$,
269 and ζ to $1e^{-2}$. We find empirically that pretraining the policy network first and then adding
270 in the posterior at a later epoch results in both good performance and mutual information
271 maximization. For our approximate posterior, $q_{\zeta|\theta}^{\omega}$, we set the value of σ_p to zero. The
272 runtime for training and verifying this model is approximately 12 hours.

273 8 Interpretable Models

274 As machine learning is being increasingly deployed into the real world, interpretability is required
275 for these systems to gain human trust [1, 3, 8]. Interpretability refers to attempts that help the user
276 understand why a machine learning model behaves the way it does. A clear visualization of a policy
277 is one way to help a human form an accurate representation of its capabilities [9]. Furthermore, an
278 interpretable model of resource allocation or planning tasks would be highly useful for a variety
279 of reasons, from decision explanations to training purposes. In Figure 2, we display interpretable
280 models generated through discretization for the low-dimensional environment and survey scheduling
281 environment.

282 9 Future Work

283 During the deployment of a discretized PNT, we required pre-inferred embeddings to understand
284 decision-maker behavior. As this involves a sample of the decision-maker’s data and the use of
285 backpropagation with a pre-discretized PNT to infer demonstrator style, we feel this can be improved
286 by producing the demonstrator embedding through the means of our approximate posterior $q_{\zeta|\theta}^{\omega}$,
287 modeled as a $PNT \setminus \omega$. This can be discretized following the framework of Section 4.3 of our paper,
288 producing an interpretable model that predicts a mean and covariance of an embedding given a single
289 state-action pair. This discretized posterior then takes in a state-action pair and produce the latent

290 embedding that generated this action. In this way, the interpretable discretized PNT has a method to
291 naturally infer the demonstrator's embedding.

292 **References**

- 293 [1] Lisa Anne Hendricks, Ronghang Hu, Trevor Darrell, and Zeynep Akata. Grounding visual explanations.
294 In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 264–279, 2018.
- 295 [2] Matthew Gombolay, Reed Jensen, Jessica Stigile, Sung-Hyun Son, and Julie Shah. Decision-making
296 authority, team efficiency and human worker satisfaction in mixed human-robot teams. In *Proceedings*
297 *of the International Joint Conference on Artificial Intelligence (IJCAI)*, New York City, NY, U.S.A., July
298 9-15 2016.
- 299 [3] Lisa Anne Hendricks, Ronghang Hu, Trevor Darrell, and Zeynep Akata. Generating counterfactual
300 explanations with natural language. *arXiv preprint arXiv:1806.09809*, 2018.
- 301 [4] Fang-I Hsiao, Jui-Hsuan Kuo, and Min Sun. Learning a multi-modal policy via imitating demonstrations
302 with mixed behaviors. *ArXiv*, abs/1903.10304, 2019.
- 303 [5] G. Ayorkor Korsah. *Exploring bounded optimal coordination for heterogeneous teams with cross-schedule*
304 *dependencies*. PhD thesis, Carnegie Mellon University, Pittsburgh, PA, January 2011.
- 305 [6] Yunzhu Li, Jiaming Song, and Stefano Ermon. Infogail: Interpretable imitation learning from visual
306 demonstrations. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and
307 R. Garnett, editors, *Advances in Neural Information Processing Systems 30*, pages 3812–3822. Curran
308 Associates, Inc., 2017.
- 309 [7] Stefanos Nikolaidis, Ramya Ramakrishnan, Keren Gu, and Julie Shah. Efficient model learning from
310 joint-action demonstrations for human-robot collaborative tasks. In *Proceedings of the Tenth Annual*
311 *ACM/IEEE International Conference on Human-Robot Interaction, HRI '15*, pages 189–196, New York,
312 NY, USA, 2015. ACM.
- 313 [8] Chris Olah, Arvind Satyanarayan, Ian Johnson, Shan Carter, Ludwig Schubert, Katherine Ye, and Alexander
314 Mordvintsev. The building blocks of interpretability. *Distill*, 3(3):e10, 2018.
- 315 [9] Steffi Paepcke and Leila Takayama. Judging a bot by its cover: An experiment on expectation setting
316 for personal robots. In *Proceedings of the 5th ACM/IEEE International Conference on Human-robot*
317 *Interaction, HRI '10*, pages 45–52, Piscataway, NJ, USA, 2010. IEEE Press. ISBN 978-1-4244-4893-7.
- 318 [10] Claude Sammut, Scott Hurst, Dana Kedzier, and Donald Michie. *Learning to Fly*, page 171–189. MIT
319 Press, Cambridge, MA, USA, 2002.
- 320 [11] Aviv Tamar, Khashayar Rohanimanesh, Yinlam Chow, Chris Vigorito, Ben Goodrich, Michael Kahane, and
321 Derik Pridmore. Imitation learning from visual data with multiple intentions. In *International Conference*
322 *on Learning Representations*, 2018.