

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

The alignment problem in curriculum learning

Permalink

<https://escholarship.org/uc/item/45s9p2z5>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 46(0)

Authors

Shafto, Patrick

Sheller, Benjamin

Publication Date

2024

Copyright Information

This work is made available under the terms of a Creative Commons Attribution License, available at <https://creativecommons.org/licenses/by/4.0/>

Peer reviewed

The Alignment Problem in Curriculum Learning

Benjamin Sheller & Patrick Shafto

Department of Mathematics and Computer Science
Rutgers University - Newark

Abstract

In curriculum learning, teaching involves cooperative selection of sequences of data via plans to facilitate efficient and effective learning. One-off cooperative selection of data has been mathematically formalized as entropy-regularized optimal transport and the limiting behavior of myopic sequential interactions has been analyzed, both yielding theoretical and practical guarantees. We recast sequential cooperation with curriculum planning in a reinforcement learning framework and analyze performance mathematically and by simulation. We prove that infinite length plans are equivalent to not planning under certain assumptions on the method of planning, and isolate instances where monotonicity and hence convergence in the limit hold, as well as cases where it does not. We also demonstrate through simulations that argmax data selection is the same across planning horizons and demonstrate problem-dependent sensitivity of learning to the teacher’s planning horizon. Thus, we find that planning ahead yields efficiency at the cost of effectiveness.

Advances in AI and machine learning enable the possibility that artificial systems may autonomously facilitate human goals, including human learning. Design of such systems requires addressing a value alignment problem (Russell, 2019; Christian, 2020), which requires interacting with the system to achieve the desired goals. Toward this end, formalizing models of cooperation among agents that bridge human and machine learning is an important direction for research (Shafto & Goodman, 2008; P. Wang et al., 2020; J. Wang et al., 2020; Goodman & Stuhlmüller, 2013; Ho et al., 2021). We identify a novel value alignment problem in the context of agents that facilitate learning, and mathematically identify and computationally test sufficient conditions for value alignment in curriculum learning.

Learning may be facilitated by the teacher planning ahead, which becomes a problem of reinforcement learning. There exists an extensive literature on curriculum learning (Elman, 1993; Khan et al., 2011; Pentina et al., 2015; Matiisen et al., 2020; Graves et al., 2017); however, this literature focuses on naive learners rather than those that reason cooperatively about the teacher’s selection of data. Theoretical results are limited (Milli & Dragan, 2020) and have not systematically considered the possibility of alignment problems or their solutions. Recent advances in theoretical foundations for cooperative inference admit a more unified formal treatment (P. Wang et al., 2020), which is necessary to understand whether, when, and why alignment problems arise.

We formalize the alignment problem in curriculum learn-

ing via the mathematical condition of consistency. Given a teacher and learner cooperatively communicating, the teacher aims to convey a distribution θ on the finite set of possible hypotheses \mathcal{H} to the learner, over an infinite horizon. That is, if θ_n denotes the learner’s distribution at the n -th round of communication, the alignment problem is to have $\lim_{n \rightarrow \infty} \sum_{h \in \mathcal{H}} |\theta(h) - \theta_n(h)| = 0$. When the teacher is conveying a specific hypothesis h' to the learner, the distribution to be learned is $\theta = \delta_{h'}$, a Dirac distribution.

We investigate the alignment problem in curriculum learning by recasting sequential cooperative Bayesian inference (SCBI) as a Markov decision process (MDP). The first section gives background on one-off cooperative inference and sequential cooperative inference, as well as the interpretation of SCBI as a Markov chain. We then recast SCBI as a Markov decision process, distinct from the trivial realization of a Markov chain as an MDP, contrasting SCBI with no planning ahead vs. using a teaching plan which calculates several steps into the future, and isolating the theoretical basis of misalignment. One main result of this section is that for a class of reward/cost functions, planning infinitely far ahead is equivalent to not planning ahead at all. The other main result is to give a condition for monotonicity in expectation which yields a sufficient requirement for alignment.

Notation. $M \in \mathbb{R}_{>0}^{|\mathcal{D}| \times |\mathcal{H}|}$ is the joint distribution for the teacher and learner between data and hypotheses, with $|\mathcal{D}|$ many rows and $|\mathcal{H}|$ many columns, with $M_{(d,h)}$ the entry of the joint distribution corresponding to datum d and hypothesis h . $M^{\theta,\lambda}$ is the joint distribution M normalized using Sinkhorn scaling to have column sums equal to θ and row sums equal to λ . That is, for every $h \in \mathcal{H}$, $\sum_{d \in \mathcal{D}} M_{(d,h)}^{\theta,\lambda} = \theta(h)$ and for every $d \in \mathcal{D}$, $\sum_{h \in \mathcal{H}} M_{(d,h)}^{\theta,\lambda} = \lambda(d)$. $\pi: \mathcal{P}(\mathcal{H}) \rightarrow \mathcal{P}(\mathcal{D})$ is a teaching strategy used by the teacher for a single round of teaching, while $\pi_R^N: \mathcal{P}(\mathcal{H}) \rightarrow \mathcal{P}(\mathcal{D})$ is the teaching strategy obtained from π by planning N teaching rounds into the future and using the random variable R , representing rewards/costs inherent to the problem. $\delta_{M_{(d,-)}^{\theta,\lambda}/\lambda(d)}$ is the atomic distribution on $\mathcal{P}(\mathcal{H})$ with atom $M_{(d,-)}^{\theta,\lambda}/\lambda(d)$; i.e. $\delta_{M_{(d,-)}^{\theta,\lambda}/\lambda(d)} \in \mathcal{P}(\mathcal{P}(\mathcal{H}))$, the space of distributions on the distributions on hypotheses, where the Markov operator Ψ_π in our formalism is acting. Ψ_π^N denotes $\Psi_\pi: \mathcal{P}(\mathcal{P}(\mathcal{H})) \rightarrow \mathcal{P}(\mathcal{P}(\mathcal{H}))$ composed

with itself N times. Frequently we will shorten the notation $\delta_{M_{(d,-)}^{\theta,\lambda}/\lambda(d)}$ to $\delta_{(d)}$.

Background

Curriculum learning involves selecting a sequence of learning problems that lead the learner to a desired knowledge or capability. We formalize these as a sequence of data that lead the learner to a target hypothesis. Throughout, we will assume teachers and learners are probabilistic agents reasoning over discrete and finite spaces of hypotheses $h \in \mathcal{H}$ and data $d \in \mathcal{D}$. Recall, in standard probabilistic inference, learners will update their posterior beliefs $P(h|d)$ in proportion to the product of the prior beliefs, $P(h)$ and the likelihood of the data, $P(d|h)$, as dictated by Bayes rule.

One-off cooperative inference. Cooperative inference between probabilistic agents differs from standard Bayesian inference in the second agent, the teacher, who selects the data, and in that the agents reason about each other's beliefs. Based on the shared joint distribution between data and hypotheses, the teacher reasons about the learner's beliefs, and samples data to pass according to the learner's current distribution on hypotheses, the joint distribution, and the desired hypothesis to be conveyed. The learner then reasons based upon what data they have been passed by the teacher and infers, based on the shared joint distribution, what hypothesis the teacher is attempting to convey. This process may be represented mathematically by the following system of equations:

$$P_L(h|d) = \frac{P_T(d|h)P_{L_0}(h)}{P_L(d)}, \quad P_T(d|h) = \frac{P_L(h|d)P_{T_0}(d)}{P_T(h)} \quad (1)$$

where $P_L(h|d)$ represents the learner's posterior probability for hypothesis h given datum d ; $P_T(d|h)$ is the probability of the teacher selecting datum d to convey hypothesis h ; $P_{L_0}(h)$ represents the learner's prior for hypothesis h ; $P_{T_0}(d)$ is the teacher's prior for selecting data d ; $P_L(d)$ and $P_T(h)$ are normalizing constants. Sinkhorn scaling of matrices (i.e. alternating row-column normalization of the joint distribution) may be used to solve equation (1), and the result is an optimal entropy-regularized plan for transporting beliefs (P. Wang et al., 2019, 2020).

Sequential cooperative inference. In sequential cooperative Bayesian inference (SCBI), a teacher and learner participate in rounds of learning. To convey a particular hypothesis (or belief on the space of possible hypotheses) from the hypothesis-space \mathcal{H} , in each round the teacher passes a datum $d \in \mathcal{D}$ to the learner, and the learner updates their belief distribution accordingly. At the end of each round, the teacher and learner both update their posterior distributions to become their prior distributions in the next round (J. Wang et al., 2020). Each round of learning behaves as in cooperative inference, where the system of equations (1) must be solved. However, at the end of each round, each round differs in having updated the prior, which is one marginal constraint for Sinkhorn scaling in (1).

The process of teaching-learning-updating works as follows: Beginning with the joint distribution M_n of the previous round and distribution $\theta_n \in \mathcal{P}(\mathcal{H})$, which represents the learner's beliefs from the previous round of teaching-learning, where $\mathcal{P}(\mathcal{H})$ is the simplex of probability distributions on \mathcal{H} , the teacher computes the Sinkhorn scaling of M_n with row sums λ and column sums θ_n . Call this M_{n+1} . Here λ is an underlying distribution on \mathcal{D} reflecting inherent biases in selecting particular data points; λ is typically taken to be the uniform distribution. Then the teacher uses the distribution $M_{n+1}\hat{\theta}$ to sample datum d_{n+1} from \mathcal{D} and passes it to the learner, where $\hat{\theta}$ is the desired belief on hypotheses which the teacher wishes to convey, typically a Dirac distribution, corresponding to a corner of the simplex. The learner then calculates M_{n+1} in exactly the same way as the teacher, then multiplies θ_n by the likelihood of selecting d_{n+1} . Normalizing gives a distribution θ_{n+1} . The process then repeats inductively, with n replaced everywhere by $n+1$.

SCBI as a Markov chain. The process of SCBI can be realized as a Markov chain on $\mathcal{P}(\mathcal{H})$ (J. Wang et al., 2020). With $T_d : \mathcal{P}(\mathcal{H}) \rightarrow \mathcal{P}(\mathcal{H})$ the map bringing the learner's prior to posterior when data d is chosen by the teacher; and $\tau : \mathcal{P}(\mathcal{H}) \rightarrow \mathcal{P}(\mathcal{D})$ the map of the teacher's sample distribution based on the learner's prior, and τ_d the d -th component of this map, the Markov transition operator for a fixed hypothesis $h \in \mathcal{H}$, $\Psi(h) : \mathcal{P}(\mathcal{P}(\mathcal{H})) \rightarrow \mathcal{P}(\mathcal{P}(\mathcal{H}))$ is defined as:

$$(\Psi(h)(\mu))(E) := \int_E \sum_{d \in \mathcal{D}} \tau_d(T_d^{-1}(\theta)) d(T_d^*(\mu))(\theta) \quad (2)$$

where T_d^* is the push-forward of T_d on Borel measures¹, μ is a Borel probability measure on the simplex $\mathcal{P}(\mathcal{H})$, and $E \subseteq \mathcal{P}(\mathcal{H})$ is a Borel measurable subset. T_d and τ are computed as above by using the Sinkhorn scaling of the joint distribution between data and hypotheses.

P. Wang et al. (2020) consider the problem of a teacher and learner communicating cooperatively in discrete rounds of teaching/learning. The teacher and learner reason using Bayesian inference at each round, without any reference to what future probability distributions on the available hypotheses might be, and without reference to any costs/rewards the teacher uses in order to determine what distribution to use to sample the data which they are passing to the learner. Although there is a discussion of SCBI as a Markov chain in (P. Wang et al., 2020), there is no extension of the method to a Markov decision process. Here we extend the formalism to include planning ahead by the teacher, as well as rewards/costs the teacher uses in planning in order to bias the learner towards/away from a particular hypothesis.

Curriculum planning as MDP

Curriculum planning involves the teacher planning several teaching moves in advance. We may model this using a Markov decision process (MDP): Let the *state space* of the

¹i.e. $T_d^*(\mu)(E) = \mu(T_d^{-1}(E))$.

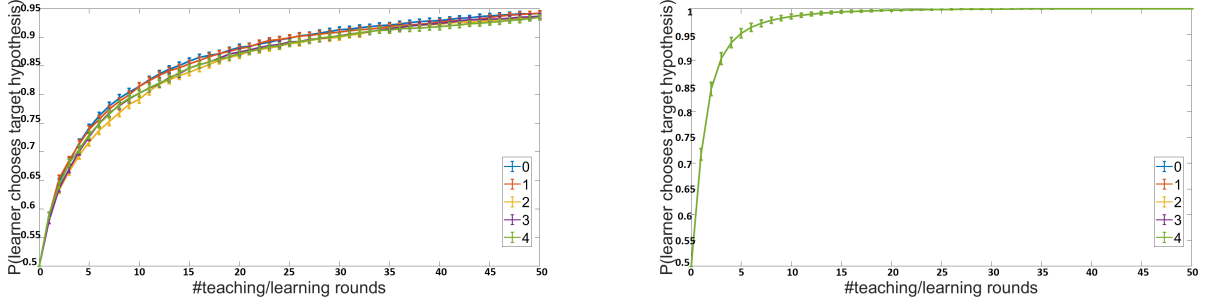


Figure 1: Analysis of convergence to the target hypothesis as a function of longer curricula (0-4) for probabilistic action selection and argmax action selection, given 3 data and 2 hypotheses. We show the probability of the target hypothesis according to the learner versus # teaching/learning rounds, using $X(\theta) = \theta([\text{true hypothesis}])$. **(left)** On average, when the teacher selects actions probabilistically (i.e. Eq. 3), longer curricula yield a marginal difference. **(right)** On average, when the teacher selects actions using argmax, longer curricula is not different from not planning ahead.

process be given by $\mathcal{S} = \mathcal{P}(\mathcal{H})$, the probability distributions on \mathcal{H} . Let the *action space* of the MDP be given by $\mathcal{A} = \mathcal{D}$, the data available for the teacher to pass to the learner (interpretation: an action $d \in \mathcal{D}$ corresponds to passing d to the learner). We fix an underlying distribution $\lambda \in \mathcal{P}(\mathcal{D})$, which represents any inherent bias towards selecting or not selecting some particular data. In SCBI (J. Wang et al., 2020), λ is the uniform distribution. The reward function R may be a combination of positive and negative pieces (to include costs, if some hypotheses are particularly undesirable).

The transition probability $T(\omega|d, \theta)$ of the MDP between probability distributions ω and θ on \mathcal{H} , based upon the teacher selecting datum d , is $\lambda(d)$ if $\omega = T_d(\theta)$ and is zero otherwise. A teaching strategy then consists of a plan $\pi : \mathcal{P}(\mathcal{H}) \rightarrow \mathcal{P}(\mathcal{D})$, effectively, ‘sample datum d using $\pi(\theta)$ when the current distribution on hypotheses is θ .’ In SCBI, the d -th component of the plan π is $\pi_d(\theta) = \tau_d(\theta)$, i.e. the adjustment of the teacher’s distribution according to the learner’s current distribution. The teacher’s strategy can be made deterministic if π_d is an atomic distribution with a single atom for every $d \in \mathcal{D}$. Throughout this paper we will be focusing on plans π which amount to the teacher preparing a curriculum by calculating what might happen in future rounds of learning based upon the current data selection.

Explicitly, the plan $\pi_R^N : \mathcal{P}(\mathcal{H}) \rightarrow \mathcal{P}(\mathcal{D})$ corresponding to planning ahead N moves based on the teaching strategy π and a random variable R is realized by the following procedure:

$$\pi_R^N(\theta)(d) = \text{Norm} \left(M_{(d,h)}^{\theta,\lambda} \mathbb{E}_{\Psi_\pi^N \delta_{(d)}}[R] + C \right) \quad (3)$$

Here $\text{Norm}(\cdot)$ is normalization to ensure that $\sum_d \pi_R^N(\theta)(d) = 1$, $M^{\theta,\lambda}$ is the Sinkhorn scaling of M so that for every $d \in \mathcal{D}$ and $h \in \mathcal{H}$, $\sum_{d' \in \mathcal{D}} M_{(d',h)}^{\theta,\lambda} = \theta(h)$ and $\sum_{h' \in \mathcal{H}} M_{(d,h')}^{\theta,\lambda} = \lambda(d)$. The h in Equation (3) is the target hypothesis the teacher desires to convey, $R : \mathcal{P}(\mathcal{H}) \rightarrow \mathbb{R}$ is a random variable representing a reward or cost, and $\Psi_\pi : \mathcal{P}(\mathcal{P}(\mathcal{H})) \rightarrow \mathcal{P}(\mathcal{P}(\mathcal{H}))$ is the Markov operator induced by the teaching plan π , i.e. for a Borel measure $\mu \in \mathcal{P}(\mathcal{P}(\mathcal{H}))$

and $E \subseteq \mathcal{P}(\mathcal{H})$ a measurable subset,

$$\Psi_\pi(\mu)(E) = \int_E \sum_{d \in \mathcal{D}} \pi_d(T_d^{-1}(\theta)) d(T_d^* \mu)(\theta), \quad (4)$$

and C is a constant so that what is being normalized to a distribution is non-negative (typically chosen to be $\varepsilon + \min \{ M_{(d,h)}^{\theta,\lambda} \mathbb{E}_{\Psi_\pi^N \delta_{(d)}}[R] | d \in \mathcal{D}, h \in \mathcal{H} \}$, for some small ε). Frequently we will drop the subscript R , to simplify notation.

Note that the behavior of the teaching/learning interaction will vary, depending upon the normalization constant C . On one extreme, as $C \rightarrow 0^+$, $\pi_R^N(\theta)(d) \rightarrow \text{Norm} \left(M_{(d,h)}^{\theta,\lambda} \mathbb{E}_{\Psi_\pi^N \delta_{(d)}}[R] \right)$. Because of the potential non-positivity, but the overall normalization, this corresponds to a signed probability measure. At the other extreme, as $C \rightarrow \infty$, $\pi_R^N(\theta)(d) \rightarrow \frac{1}{|\mathcal{D}|}$, independently of θ ; i.e. the teacher’s distribution is the uniform distribution on $|\mathcal{D}|$, regardless of the learner’s beliefs, hence the teacher’s choice of data is random. In particular, if C is much greater than the expectation term for some particular choice of data d , then $\pi_R^N(\theta) \approx \frac{1}{|\mathcal{D}|}$. In order to make the distribution positive, we must have:

$$C > \left| \min_{d \text{ s.t. } M_{(d,h)}^{\theta,\lambda} \mathbb{E}_{\Psi_\pi^N \delta_{(d)}}[R] \leq 0} M_{(d,h)}^{\theta,\lambda} \mathbb{E}_{\Psi_\pi^N \delta_{(d)}}[R] \right| \quad (5)$$

On the other hand, in order to preserve the teacher’s knowledge and avoid random data selection, we would also like:

$$C \leq \min_{d | M_{(d,h)}^{\theta,\lambda} \mathbb{E}_{\Psi_\pi^N \delta_{(d)}}[R] > 0} M_{(d,h)}^{\theta,\lambda} \mathbb{E}_{\Psi_\pi^N \delta_{(d)}}[R] \quad (6)$$

However, for random variables R which place too high a cost on some outcomes, versus too small a reward, it may be impossible to simultaneously meet these two conditions. That is, ensuring a positive probability distribution may create a more uniform distribution for the data which already have a positive probability of selection, prior to the addition of the normalization constant C .

In Equation (3), if $R(\omega) = \omega(h)$, this corresponds to the teacher calculating N steps ahead in order to choose the data at the current step which will increase the likelihood the learner will be able to infer the target hypothesis in the future. Furthermore, we can replace the expectation of the random variable $\omega(h)$ with an expectation that accounts for rewards/costs inherent to the problem. For example, if π is the SCBI teaching strategy, then $\pi(\theta)(d) = \frac{M_{(d,h)}^{\theta,\lambda}}{\theta(h)}$, where h is the target hypothesis (P. Wang et al., 2020).

The expectation term of (3) may be simplified as follows, letting $\tilde{\theta} = M_{(d,-)}^{\theta,\lambda}$, and assuming that the original teaching plan follows the scheme of SCBI:

$$\mathbb{E}_{\Psi^N \delta_{\tilde{\theta}}}[R(\theta)] = \sum_{d_1, \dots, d_N \in \mathcal{D}} \tau_{d_N}(\tilde{\theta}) \tau_{d_{N-1}}(T_{d_N}(\tilde{\theta})) \dots \tau_{d_1}(T_{d_2} \circ T_{d_3} \circ \dots \circ T_{d_N}(\tilde{\theta})) R(T_{d_1} \circ \dots \circ T_{d_N}(\tilde{\theta})). \quad (7)$$

This formula is useful for simulations, as it allows us to replace the integral arising from the expectation $\mathbb{E}_{\Psi^N \delta_{(d)}}[R]$ by a finite (though growing exponentially in the number of steps planning ahead) sum. The teacher, when planning a curriculum consisting of finitely many data points, can only guide the learner to a finite number of distributions on hypotheses.

Optimal policy

We may compute the optimal policy as follows:

$$\mathbb{E} \left[\sum_{t=0}^{\infty} \gamma^t R_{a_t}(\theta_t, \theta_{t+1}) \right] = \sum_{t=0}^{\infty} \gamma^t R(\theta_t) P_{a_t}(\theta_t, \theta_{t+1}) = \sum_{t=0}^{\infty} \gamma^t R(\theta_t) \lambda(d(t)) \delta_{\theta_{t+1}, T_{d(t)}(\theta_t)} \pi_{d(t)}(\theta_t) \quad (8)$$

Assuming that λ is uniform, i.e. the teacher has no underlying bias with regards to the data, and also using the fact that $\delta_{\theta_{t+1}, T_{d(t)}(\theta_t)}$ is only nonzero when $T_{d(t)}(\theta_t) = \theta_{t+1}$, this becomes:

$$= \frac{1}{|\mathcal{D}|} \sum_{t=0}^{\infty} \gamma^t R(T_{d(t-1)} \circ T_{d(t-2)} \circ \dots \circ T_{d(0)}(\theta_0)) \cdot \pi_{d(t)}(T_{d(t-1)} \circ T_{d(t-2)} \circ \dots \circ T_{d(0)}(\theta_0)) \cdot \pi_{d(t-1)}(T_{d(t-2)} \circ \dots \circ T_{d(0)}(\theta_0)) \dots \pi_{d(0)}(\theta_0)$$

The optimal policy which maximizes the expectation term in Equation 8 is therefore the argmax over all possible functions $d: \mathbb{N} \rightarrow \mathcal{D}$, where the action at step t is $d(t-1)$, and argmax over functions $\pi: \mathcal{P}(\mathcal{H}) \rightarrow \mathcal{P}(\mathcal{D})$. Note that Equation 8 is the same as Equation 7 with a factor of $\frac{1}{|\mathcal{D}|}$ in front, $N \rightarrow \infty$, reverse indexing, and a discount factor $\gamma \in [0, 1]$. This implies that taking an argmax over the distribution in Equation 3 gives a finite horizon approximation of the optimal policy.

Theorem 1 *The optimal policy is given by the following:*

$$\lim_{\gamma \rightarrow 1} \operatorname{argmax}_{\pi} \operatorname{argmax}_{d: \mathbb{N} \rightarrow \mathcal{D}} \mathbb{E} \left[\sum_{t=0}^{\infty} \gamma^t R(\theta_t) \right] = \operatorname{argmax}_{\pi} \lim_{N \rightarrow \infty} \operatorname{argmax}_{d: \{0, \dots, N-1\} \rightarrow \mathcal{D}} \pi^N(\theta_0) \quad (9)$$

Proof: Note that normalization does not affect which entries of $\pi^N(\theta)$ are larger than others, so we have:

$$\begin{aligned} & \lim_{N \rightarrow \infty} \operatorname{argmax}_{d: \{0, \dots, N-1\} \rightarrow \mathcal{D}} \pi_{d(N-1)}^N(\theta_0) \\ &= \lim_{N \rightarrow \infty} \operatorname{argmax}_{d: \{0, \dots, N-1\} \rightarrow \mathcal{D}} \left(M_{(d(N-1), h)}^{\theta, \lambda} \mathbb{E}_{\Psi_{\pi}^N \delta_{(d)}}[R] + C \right) \\ &= \lim_{N \rightarrow \infty} \operatorname{argmax}_{d: \{0, \dots, N-1\} \rightarrow \mathcal{D}} \left(M_{(d(N-1), h)}^{\theta, \lambda} \mathbb{E}_{\Psi_{\pi}^N \delta_{(d(N-1))}}[R] \right) \\ &= \lim_{N \rightarrow \infty} \operatorname{argmax}_{d: \{0, \dots, N-1\} \rightarrow \mathcal{D}} \left(T_{d(N-1)}(\theta)(h) \mathbb{E}_{\Psi_{\pi}^N \delta_{(d(N-1))}}[R] \right) \end{aligned}$$

Because $T_{d(N-1)}(\theta)(h)$ is multiplied by every term of the expectation sum, we obtain that the above is equal to:

$$= \lim_{N \rightarrow \infty} \operatorname{argmax}_{d: \{0, \dots, N-1\} \rightarrow \mathcal{D}} \left(\mathbb{E}_{\Psi_{\pi}^N \delta_{(d(N-1))}}[R] \right) = \lim_{\gamma \rightarrow 1} \mathbb{E} \left[\sum_{t=0}^{\infty} \gamma^t R(\theta_t) \right]$$

Taking the argmax over policies $\pi: \mathcal{P}(\mathcal{H}) \rightarrow \mathcal{P}(\mathcal{D})$ then yields the result. \square

The policy $\pi_{d(t)}(\theta_0)$ which satisfies the argmax above at time step t should be the Dirac distribution $\delta_{d(t)}$, where $d(t)$ is the optimal choice of datum at time t . We may therefore obtain an estimate of the optimal policy by taking the argmax of π^N as N becomes progressively larger.

SCBI vs SCBI with planning. When preparing a curriculum, if the teacher does not plan ahead, they may base their teaching strategy for each round on Bayesian inference. This amounts to an initial teaching strategy of SCBI. However, if the same Bayesian teacher plans a curriculum which includes infinitely many data points, and uses the random variable $R(\theta) = \theta(h_{\text{target}})$ in their planning in order to effectively sample data points leading the learner to a stronger belief in the true hypothesis, the usual process of SCBI is recovered.

Theorem 2 *Suppose that $\pi: \mathcal{D} \rightarrow \mathcal{H}$ is a teaching plan and R is a reward/cost such that for every $d \in \mathcal{D}$, $\lim_{N \rightarrow \infty} \mathbb{E}_{\Psi_{\pi}^N \delta_{(d)}}[R]$ exists, and is independent of d . Then $\pi_R^{\infty} := \lim_{N \rightarrow \infty} \pi_R^N$ exists and is equal to the SCBI teaching strategy.*

Proof: Let $\theta \in \mathcal{P}(\mathcal{H})$, and $d \in \mathcal{D}$, and suppose that the teacher is teaching hypothesis $h \in \mathcal{H}$. To ease the readability, we will drop the subscript R on the teaching plan. Then:

$$\lim_{N \rightarrow \infty} \pi^N(\theta)(d) = \lim_{N \rightarrow \infty} \frac{M_{(d,h)}^{\theta,\lambda} \mathbb{E}_{\Psi_{\pi}^N \delta_{(d)}}[R]}{\sum_{d'} M_{(d',h)}^{\theta,\lambda} \mathbb{E}_{\Psi_{\pi}^N \delta_{(d')}}[R]} = \frac{M_{(d,h)}^{\theta,\lambda}}{\sum_{d'} M_{(d',h)}^{\theta,\lambda}} = \frac{M_{(d,h)}^{\theta,\lambda}}{\theta(h)} \quad (10)$$

This is the formula used in the standard SCBI teaching strategy at each round of learning. \square




Figure 2: Analysis of convergence as before using reward $R(\theta) = 10 \cdot d_{L^1}(\theta, \delta_{h_{\text{bad}}}) - d_{L^1}(\theta, \delta_{h_{\text{true}}})$. On average, when the teacher selects actions probabilistically (**upper left**) or deterministically (argmax) (**upper right**) curricula yield worse outcomes. The average reward is shown in (**lower**). Planning ahead leads to misalignment between learners and teachers, when the teacher attempts to avoid an undesirable hypothesis. Note that further planning ahead reduces misalignment, consistent with the fact that planning sufficiently far ahead is equivalent to not planning at all.

Corollary 3 *With π the standard SCBI teaching strategy and $R(\theta) = \theta(h)$, where h is the target hypothesis, planning infinitely far ahead is identical to not planning ahead at all.*

Proof: By the proof of ((J. Wang et al., 2020), Theorem 3.5), $\lim_{N \rightarrow \infty} \mathbb{E}_{\Psi_{\pi}^N}[\theta(h)] = 1$, for any Borel measure $\mu \in \mathcal{P}(\mathcal{P}(\mathcal{H}))$. The claim then follows from Theorem 2 above. \square

The corollary says that the limit of planning infinitely far ahead using SCBI is identical to SCBI with no planning ahead! Intuitively, this makes sense: in the infinite limit, the teacher is considering all possible infinite strings of data to pass to the learner; however, most of the strings will be indistinguishable as the learner approaches a particular belief on the true hypothesis, and so only the short-term behaviour of the strings is important. Furthermore, because the proof in P. Wang et al. (2020) also shows that the convergence is monotone increasing; this implies that not planning ahead is the optimal teaching strategy when the reward at each round is equal to the probability that the learner would select the target hypothesis out of all hypotheses.

Figure 1 compares planning ahead up to four steps for two hypotheses and three data points, using $R(\theta) = \theta(h)$, with h the hypothesis to be taught, assuming actions are selected probabilistically. The vertical axis is the learner’s probability of choosing h as the correct hypothesis if they were to randomly guess, while the horizontal axis represents the number of rounds of teaching/learning. The plot was created by randomly generating 1500 initial joint distributions, performing 50 rounds of teaching-learning with each 30 times, then aver-

aging over the learner’s performance.

An interesting feature of planning ahead with $R(\theta) = \theta(h_{\text{target}})$ is that if the teacher uses the deterministic procedure of argmax instead of sampling from the distribution to choose data to pass to the learner, the choice is the same for every number of steps planning ahead; i.e., in this case, local maximization is global maximization (see Figure 1).

Other random variables are available to use in the expectation: for example, if there is one particular hypothesis which the teacher wishes to avoid, while biasing the learner toward the true hypothesis, the random variable could be $R(\theta) = 10 \cdot d_{L^1}(\theta, \delta_{h_{\text{bad}}}) - d_{L^1}(\theta, \delta_{h_{\text{true}}})$, where $d_{L^1}(\cdot, \cdot)$ is the L^1 distance between θ and δ_h represented as points of a simplex. In this case, there is non-monotonic behaviour, as the teacher overcompensates for trying to move the learner away from the ‘bad’ hypothesis, which subsequently may lead the learner closer to a neutral hypothesis than to the true hypothesis. See Figure (2), in particular the trajectories corresponding to planning ahead one step and two steps, where the probability that the learner selects the true hypothesis decreases.

Guaranteed alignment via monotonicity. One of the key results of (P. Wang et al., 2020) is the consistency of SCBI. Consistency here refers to the convergence in expectation of the learner’s belief distribution to a Dirac distribution on the target hypothesis over the infinite limit of teaching rounds. In particular, P. Wang et al. (2020) shows that if monotonicity holds, i.e. $\mathbb{E}_{\Psi_{\pi}^N}[\theta(h)] - \mathbb{E}_{\mu}[\theta(h)] > 0$, where h is the target hypothesis, then consistency and hence alignment follows. By writing out the equation above for monotonicity with π re-

placed by π_R^N , for some choice of rewards/costs R , we obtain a condition for monotonicity and hence for alignment. For $\mathbb{E}_{\Psi_{\pi_R\mu}}[\theta(h)]$ and $\mathbb{E}_\mu[\theta(h)]$, we may obtain an explicit condition for when $\mathbb{E}_{\Psi_{\pi_R\mu}}[\theta(h)] \geq \mathbb{E}_\mu[\theta(h)]$. Hence:

Theorem 4 *Monotonicity holds if and only if for any $\mu \in \mathcal{P}(\mathcal{H})$:*

$$\int_{\Delta} \sum_d \pi_d(\theta) \cdot (T_d(\theta)(h) - \theta(h)) d\mu(\theta) > 0 \quad (11)$$

Proof: Throughout this proof, we rewrite $\pi(\theta)(d)$ as $\pi_d(\theta)$ for ease of readability. Expanding $\mathbb{E}_{\Psi_{\pi\mu}}[\theta(h)] - \mathbb{E}_\mu[\theta(h)]$ using the definition of Ψ_π yields:

$$\begin{aligned} & \int_{\Delta} \sum_d \pi_d(T_d^{-1}(\theta)) \theta(h) d(T_d^*\mu)(\theta) - \int_{\Delta} \theta(h) d\mu(\theta) \\ &= \sum_d \int_{T_d^{-1}(\Delta)} \pi_d(\theta) T_d(\theta(h)) d\mu(\theta) - \int_{\Delta} \theta(h) d\mu(\theta) \\ &= \int_{\Delta} \sum_d \pi_d(\theta) \cdot (T_d(\theta)(h) - \theta(h)) d\mu(\theta) \end{aligned}$$

To get from the penultimate line to the final line, we use the fact that $T_d : \Delta \rightarrow \Delta$ is a homeomorphism of the simplex. \square

From this equation, we can see that if the $\delta_{M_{(d,-)}^{\theta,\lambda}}$ expectation of R is overly negative for a hypothesis for which $T_d(\theta)(h) > \theta(h)$, while $T_{d'}(\theta)(h) < \theta(h)$ for other hypotheses, monotonicity will be broken. This implies that if the teacher places a heavy cost on a hypothesis lying close in belief-space to the target hypothesis, the curriculum of the teacher may over-emphasize moving away from the heavy cost hypothesis, at the expense of potentially converging to a more neutral hypothesis. Here two hypotheses h_1 and h_2 ‘lying close in belief-space’ means that with respect to the initial joint distribution M , the L^1 distance on the simplex $\mathcal{P}(\mathcal{D})$ between $M_{(-,h_1)}$ and $M_{(-,h_2)}$ is small. This implies that whatever data the teacher passes to the learner will affect both rows similarly.

Related work

The mathematical theory of cooperative communication as optimal transport has been explored (P. Wang et al., 2020; Yang et al., 2018; P. Wang et al., 2019). Sequential cooperation as utilized in this paper, including proofs of stability and rates of convergence, may be found in (J. Wang et al., 2020).

Cooperative inference and reinforcement learning appear together in several contexts. For example, cooperative inverse reinforcement learning is centered around cooperative inference (Hadfield-Menell et al., 2016; Fisac et al., 2020), and cooperative inference and reinforcement learning often appear in the coordination of multiple agents, e.g. (Pesce & Montana, 2020). Milli & Dragan (2020) consider misalignment in terms of levels of theory of mind recursion. Our work differs in considering the problem of developing curricula to foster learning and offers novel theoretical and simulation-based insights into the possibility of misaligned curricula.

In curriculum learning (Elman, 1993), a learner is presented with carefully chosen examples or tasks, often increasing in complexity and carefully curated by the teacher. Such strategies are used in human cognition when teaching-learning complex tasks (Khan et al., 2011), and curriculum learning is a useful method in machine learning for gradually revealing complex concepts (Bengio et al., 2009). Curriculum learning is used, for example, in neural networks in order to maximize learning efficiency (Graves et al., 2017), and multitask learning (Pentina et al., 2015). Curriculum learning has been included into the framework of a partially observable Markov decision process in Matiisen et al. (2020); Narvekar & Stone (2018). In Matiisen et al. (2020), the student is not directly observable by the teacher and the actions of the teacher correspond to teaching the learner on a certain task for a specific number of iterations. In Narvekar & Stone (2018), curriculum design is viewed as transfer learning, and the problem of learning a meta-policy is addressed. Our approach is distinctive in focusing on the alignment problem and offering strong mathematical proofs.

It is attractive to consider curricula that actively avoid common misconceptions. We are not aware of prior work in curriculum learning that does so, but it is common in reinforcement learning to consider worlds with negative reward, which is equivalent. In education, dealing with misconceptions most typically occurs through direct negation; for example, by including incorrect examples for students to correct (Heemsoth & Heinze, 2014) and through refutation texts (Tippett, 2010). Evidence suggests that such efforts are only effective when students know enough about the domain and are most effective at avoiding the misconception rather than inducing the correct belief.

Conclusion

We investigate the possibility of alignment problems in curriculum learning by building recent theoretical advances in modeling cooperation. Recasting sequential cooperative Bayesian inference (SCBI) as a Markov decision process (MDP) enables the inclusion of curriculum learning to a teacher and learner cooperatively learning in multiple rounds of interactions. Through theoretical and simulation-based analysis, we show that curriculum planning introduces brittleness that leads to misalignment when the curriculum introduces additional costs, for example when avoiding a misconception, or when the teacher has imperfect knowledge of the learner. SCBI without curriculum planning offers competitive rates of convergence and desirable theoretical guarantees across our theoretical and simulation analyses. We also show that under simple assumptions on the reward/cost, e.g. taking the reward to be the probability that the learner will select the correct hypothesis, myopia appears to be optimal: planning ahead can actually decrease the rate of convergence, while converging to no planning as the number of steps ahead tends to infinity.

References

- Bengio, Y., Laouradour, J., Colloberty, R., & Weston, J. (2009). Curriculum learning. In L. Bottou & M. Littman (Eds.), *Proceedings of the 26th annual international conference on machine learning (icml 2009)* (p. 41–48). New York, NY: Association for Computing Machinery.
- Christian, B. (2020). *The alignment problem: Machine learning and human values*. WW Norton & Company.
- Elman, J.-L. (1993). Learning and development in neural networks: the importance of starting small. *Cognition*, 48, 71–99.
- Fisac, J.-F., Gates, M.-A., Hamrick, J.-B., Liu, C., Hadfield-Menell, D., Palaniappan, M., ... Dragan, A.-D. (2020). Pragmatic-pedagogic value alignment. In N. Amato, G. Hager, S. Thomas, & M. Torres-Torriti (Eds.), *Robotics research. springer proceedings in advanced robotics, vol 10* (pp. 49–57). Switzerland: Springer, Cham.
- Goodman, N.-D., & Stuhlmüller, A. (2013). Knowledge and implicature: Modeling language understanding as social cognition. *Topics in cognitive science*, 5(1), 173–184.
- Graves, A., Bellemare, M.-G., Menick, J., Munos, R., & Kavukcuoglu, K. (2017). Automated curriculum learning for neural networks. In D. Precup & Y. W. Teh (Eds.), *Proceedings of the 34th international conference on machine learning (icml 2017)* (p. 1311–1320). online: JMLR.org.
- Hadfield-Menell, D., Russell, S., Abbeel, P., & Dragan, A.-D. (2016). Cooperative inverse reinforcement learning. In D. Lee, M. Sugiyama, U. Luxburg, I. Guyon, & R. Garnett (Eds.), *Advances in neural information processing systems (neurips 2016)* (pp. 3909–3917). Red Hook, NY: Curran Associates Inc.
- Heemsoth, T., & Heinze, A. (2014). The impact of incorrect examples on learning fractions: A field experiment with 6th grade students. *Instructional Science*, 42(4), 639–657.
- Ho, M. K., Cushman, F., Littman, M. L., & Austerweil, J. L. (2021). Communication in action: Planning and interpreting communicative demonstrations. *Journal of Experimental Psychology: General*, 150(11), 2246.
- Khan, F., Zhu, J., & Mutlu, B. (2011). How do humans teach: on curriculum learning and teaching dimension. In J. Shawe-Taylor, R. Zemel, P. Bartlett, F. Pereira, & K. Weinberger (Eds.), *Advances in neural information processing systems 24 (nips 2011)* (pp. 1449–1457). Red Hook, NY: Curran Associates Inc.
- Matiisen, T., Oliver, A., Cohen, T., & Schulman, J. (2020). Teacher-student curriculum learning. *IEEE Transactions on Neural Networks and Learning Systems*, 31(9), 3732–3740.
- Milli, S., & Dragan, A.-D. (2020). Literal or pedagogic human? analyzing human model misspecification in objective learning. In *Uncertainty in artificial intelligence* (pp. 925–934).
- Narvekar, S., & Stone, P. (2018). Learning curriculum policies for reinforcement learning. *arXiv preprint arXiv:1812.00285*.
- Pentina, A., Sharmanska, V., & Lampert, C.-H. (2015). Curriculum learning of multiple tasks. In L. O’Conner (Ed.), *2015 IEEE conference on computer vision and pattern recognition (cvpr 2015)* (p. 5492–5500). Piscataway, NJ: Conference Publishing Services.
- Pesce, E., & Montana, G. (2020). Improving coordination in small-scale multi-agent deep reinforcement learning through memory-driven communication. *Machine Learning*, 109, 1727–1747.
- Russell, S. (2019). *Human compatible: Artificial intelligence and the problem of control*. Penguin.
- Shafte, P., & Goodman, N. (2008). Teaching games: Statistical sampling assumptions for learning in pedagogical situations. In B. C. Love, K. McRae, & V. M. Sloutsky (Eds.), *Proceedings of the 30th annual conference of the cognitive science society* (pp. 1632–1637). Austin, TX: Cognitive Science Society.
- Tippett, C.-D. (2010). Refutation text in science education: A review of two decades of research. *International journal of science and mathematics education*, 8(6), 951–970.
- Wang, J., Wang, P., & Shafte, P. (2020). Sequential cooperative bayesian inference. In H. D. III & A. Singh (Eds.), *Proceedings of machine learning research (pmlr 2020)* (pp. 10039–10049). online: PMLR.
- Wang, P., Paranamana, P., & Shafte, P. (2019). Generalizing the theory of cooperative inference. In K. Chaudhuri & M. Sugiyama (Eds.), *Proceedings of machine learning research (pmlr 2019)* (p. 1841–1850). online: PMLR.
- Wang, P., Wang, J., Paranamana, P., & Shafte, P. (2020). A mathematical theory of cooperative inference. In H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, & H. Lin (Eds.), *Advances in neural information processing systems pre-proceedings (neurips 2020)*. online: PMLR.
- Yang, S. C.-H., Yu, Y., Givchi, A., Wang, P., Vong, W. K., & Shafte, P. (2018). Optimal cooperative inference. In A. Storkey & F. Perez-Cruz (Eds.), *Proceedings of the 21st international conference on artificial intelligence and statistics (aistats 2018)* (p. 376–385). online: PMLR.