
Leveraging Time Irreversibility with Order-Contrastive Pre-training

Monica Agrawal*
MIT CSAIL

Hunter Lang*
MIT CSAIL

Michael Offin
MSKCC

Lior Gazit
MSKCC

David Sontag
MIT CSAIL

Abstract

Label-scarce, high-dimensional domains such as healthcare present a challenge for modern machine learning techniques. To overcome the difficulties posed by a lack of labeled data, we explore an “order-contrastive” method for self-supervised pre-training on longitudinal data. We sample pairs of time segments, switch the order for half of them, and train a model to predict whether a given pair is in the correct order. Intuitively, the ordering task allows the model to attend to the *least time-reversible* features (for example, features that indicate progression of a chronic disease). The same features are often useful for downstream tasks of interest. To quantify this, we study a simple theoretical setting where we prove a finite-sample guarantee for the downstream error of a representation learned with order-contrastive pre-training. Empirically, in synthetic and longitudinal healthcare settings, we demonstrate the effectiveness of order-contrastive pre-training in the small-data regime over supervised learning and other self-supervised pre-training baselines. Our results indicate that pre-training methods *designed for* particular classes of distributions and downstream tasks can improve the performance of self-supervised learning.

1 Introduction

The advent of electronic health records has led to an explosion in longitudinal health data. This data can power comparative effectiveness studies, provide clinical decision support, enable retrospective research over real-world outcomes, and inform clinical trial de-

sign. However, longitudinal health data is often complex, unstructured, and high-dimensional and thus untapped. Typically, limited *labeled* data is available for downstream tasks of interest, and labels can be prohibitively expensive to obtain: long records are tedious to synthesize, comprehension requires domain expertise, and patient privacy regulations limit data-sharing across institutions (Bleackley and Kim, 2013; Xia and Yetisgen-Yildiz, 2012). Fortunately, given the large amount of unlabeled data, *self-supervision* is a promising avenue.

In self-supervision, models are first pre-trained to optimize an objective over unlabeled data, with the goal of learning representations that capture important semantic structure about the input data modality. For example, in masked language modeling for text, the model is trained to predict the identity of randomly masked tokens. Performing well at this objective should require a representation of sentence syntax and semantics. Once pre-trained, self-supervised representations can be used for downstream supervised tasks. However, despite the success of self-supervision across domains, the development of new self-supervised objectives has been a largely heuristic endeavor. *Why* does pre-training improve performance on downstream tasks? Whether this happens depends on both the self-supervised objective and the downstream task itself. But the assumptions linking the self-supervised objective to the downstream tasks of interest are rarely, if ever, made explicit.

In this work, we design a self-supervised objective with a *particular class* of data distributions and downstream tasks in mind. We aim to make explicit the type of distributions and downstream tasks on which we expect this method to work. We are interested primarily in the types of time-series that arise in longitudinal health data for patients, particularly those with chronic conditions, e.g. cancer, autoimmune disorders, and neurodegenerative diseases. These time-series include long sequences of clinical notes, insurance claims data, biomarker measurements, or combinations thereof. A key differentiating feature of these data is that a given trajectory can *change quickly*.

Proceedings of the 25th International Conference on Artificial Intelligence and Statistics (AISTATS) 2022, Valencia, Spain. PMLR: Volume 151. Copyright 2022 by the author(s).




Figure 1: Depiction of the data-generation and learning process for order-contrastive pre-training. For each trajectory, a pair of consecutive windows is sampled uniformly at random, flipped with probability 0.5, and presented to the model. The model is trained to predict whether the presented pair of windows is in the correct (+1) or incorrect (−1) order.

For example, a patient may develop a new symptom between subsequent healthcare visits, or a certain biomarker value (e.g., blood pressure) may dramatically increase.

Additionally, these changes largely tend to be *irreversible* with respect to time. For example, once the word “metastasis” appears in a clinical note, nearly all subsequent notes tend to comment on the state of that metastasis (so the word “metastasis” appears in those notes as well). To train a good representation for certain downstream problems (e.g., “what is the patient’s current disease state?”), the self-supervised objective should *attend* to these changes, rather than suppress them.

These properties make such data distributions unsuitable for several existing self-supervised objectives. For example, Franceschi et al. (2019) train a model so that the representation of each time segment is more similar to those of its subsegments than the representation of a randomly chosen segment from another trajectory. Similar techniques have been used to learn image representations from video: two subsequent video frames are likely to contain the same objects (Mobahi et al., 2009; Goroshin et al., 2015). These approaches are all similar to the idea of *slow feature analysis* (Wiskott and Sejnowski, 2002) for extracting representations of an input signal that change slowly over time. Representation learning techniques based on the ideas of slow feature analysis are appropriate for some downstream tasks and time-series data types, such as the ones studied in the works above, but not, we argue, for data where the latent variable of interest (such as disease state) can have large changes between subsequent time steps (e.g., between visits to a medical care center). Motivated by the example of chronic diseases in healthcare, we focus on the setting where time-irreversible features are highly useful for downstream classification, and where these features may exhibit large changes between subsequent time steps.

In this work, we introduce a self-supervised objective called *order-contrastive pre-training* (OCP). For each trajectory in the input data, we sample random pairs of time segments, switch the order for half of them,

and train a model to predict whether a given pair is in the correct order (*positives*) or in the incorrect order (*negatives*).

This procedure is shown in Figure 1. OCP is very similar to an existing technique known as *permutation-contrastive learning*, or PCL (Hyvärinen and Morioka, 2017). PCL was also designed to take advantage of temporal dependence between features of the input signal to learn useful representations. The key difference between these two objectives is in the sampling of the *negatives*. Where the negatives in OCP are incorrectly-ordered window pairs, the negatives in PCL are *random* window pairs from the same trajectory, and could be in the correct order. In their simplest forms, the positive samples for the two methods are identical: pairs of consecutive windows in the correct order.

Intuitively, the same time-irreversible features that are useful for the OCP and PCL objectives should also be useful for downstream prediction tasks. To formalize and quantify this, we study a class of data distributions motivated by the preceding discussion. When the representation belongs to a simple hypothesis class (effectively, when the representation is a *feature selector*), we prove a finite-sample bound on the *downstream* error of a representation learned using OCP. Although this setting is much simpler than those that appear in similar work (it involves *linear*, rather than nonlinear, representations of the input data), we show that this model still admits interesting behavior. In particular, we give an example of a data distribution in this setup where OCP and PCL provably learn different representations. Additionally, this model indicates that even when two methods have the same performance with infinite unlabeled data, there is an unlabeled-sample-complexity benefit to using a “clean” distribution of negatives, which matches well with prior work on other contrastive learning algorithms (Chuang et al., 2020).

We supplement this motivating theoretical study with experiments on real-world time-series data. Our results indicate that for the types of data and tasks discussed above, both OCP and PCL representations can enjoy better downstream prediction performance

than those trained using existing self-supervised baselines. Moreover, complementing our theoretical results, we show a real-world scenario where OCP outperforms PCL in the low labeled-data regime despite the seemingly minor difference between the two objectives. Given that OCP and PCL only differ slightly in their negative sampling, these results give further theoretical and empirical evidence for the importance of the negative sampling details in contrastive learning, complementing several recent works (Chuang et al., 2020; Robinson et al., 2021; Liu et al., 2021).

2 Order-pretraining algorithm

We suppose each data point X is a time series, $X = (X^1, \dots, X^\tau)$, where τ is the number of sample points and may vary with X . We also suppose the samples take values in some common set \mathcal{X} . Let a window w be an element of $\{1, \dots, \tau\}$, and let X^w be the corresponding element of X .¹

Given a trajectory X , we use the following generative process to sample a data point (Z, Y) for our contrastive task. First, Y is chosen uniformly at random from $\{-1, 1\}$. Next, random windows W and W' are chosen (in a manner explained below). The segments X^W and $X^{W'}$ corresponding to windows W, W' are combined into a tuple Z . The pair (Z, Y) is then a sample for the contrastive task. A model h , given by a composition of a classifier $c \in \mathcal{C}$ and a representation $g \in \mathcal{G}$, is trained to predict Y from Z :

$$\underset{h=(c,g)}{\text{minimize}} R_{\text{ord}}(h) := \mathbb{E}_{(Z,Y)}[h(Z) \neq Y] \quad (1)$$

Here $h(Z) = h(X^W, X^{W'}) = c(g(X^W), g(X^{W'}))$. That is, h first computes the representation g for each window, then uses a classifier c to predict whether the tuple Z is in the correct order. The representation g can then be re-used on a downstream task. The remaining design choice is to specify the process for sampling windows.

Order-contrastive pre-training. A simple choice for sampling random windows W, W' is to sample a random pair $(W, W + 1)$ in the *correct* order when $Y = 1$, and $(W + 1, W)$ in the *incorrect* order when $Y = -1$. We refer to the optimization problem (1) with this choice of sampling as *order-contrastive pre-training* (OCP). This can easily be generalized to non-consecutive window pairs. The pretraining task (1) is thus to *contrast* windows in the correct order with windows in the incorrect order.

¹For simplicity, we only consider windows of size 1. Our results straightforwardly generalize to windows of arbitrary size ℓ , where $w = (w_1, \dots, w_\ell)$ is a subinterval of $\{1, \dots, \tau\}$ and $X^w = (X^{w_1}, \dots, X^{w_\ell})$.

Permutation-contrastive learning. Another simple choice is to again sample a pair $(W, W + 1)$ in the correct order when $Y = 1$, but sample a *random pair* when $Y = -1$. This is the data generation process for *permutation-contrastive learning* (Hyvärinen and Morioka, 2017). Note that the only differences between OCP and PCL are that in PCL, (i) the negative samples ($Y = -1$) need not be consecutive, and (ii) some are in the correct order. The distributions of positive samples are identical. We refer to the procedure (1) with this sampling as *permutation-contrastive learning* (PCL). This exactly matches the contrastive sample distribution in Hyvärinen and Morioka (2017, equations (10)-(11)). Here, the pretraining task is to contrast consecutive windows in the correct order versus random window pairs.

Comparison. These two sampling methods seem very similar—they only differ slightly in the distribution of negatives (i.e., conditioned on $Y = -1$). However, we show theoretically and empirically in the following sections that they can learn very different representations when used in (1), and they can have different unlabeled sample complexities even if they eventually find the same representation. This gives further evidence of the importance of negative sampling for contrastive learning methods (see, e.g., Chuang et al. (2020)). We give a finite-sample bound for the downstream classification performance of a representation learned using OCP in a simple setup motivated by time series data and predictive tasks in healthcare.

3 Finite-sample guarantee for time-irreversible features

In this section, we study a class of distributions motivated by applications to time-series data in healthcare. We assume for simplicity that each $X^t \in \mathcal{X} = \{0, 1\}^d$. We identify a set of four assumptions for which we can prove a finite sample guarantee for the set of *feature selector* representations \mathcal{G} . Here we use \mathcal{F} to refer to the downstream hypothesis class, and we overload Y to refer to the *downstream* label of interest.

Assumption 1. *There exists a set S of time-irreversible features. Formally, $\forall i \in S, \forall t, \mathbb{P}[X_i^t = 1, X_i^{t+1} = 0] = 0$.*

Assumption 2. *When the features in S are not changing, the other features are time-reversible. More formally, for all t , and all $v, v' \in \{0, 1\}^d$, if $v_S = v'_S$, $\mathbb{P}[X^t = v, X^{t+1} = v'] = \mathbb{P}[X^t = v', X^{t+1} = v]$.*

Assumption 3. *There are no “redundant” features in S . For all U such that $S \not\subset U$, there exists t and*

$v \in \{0, 1\}^d$, $v' \in \{0, 1\}^d$, with $v_{U \cap S} = v'_{U \cap S}$ and:

$$\begin{aligned} \min \mathbb{P}[X_S^t \subsetneq X_S^{t+1}, X_U^t = v_U, X_U^{t+1} = v'_U]; \\ \mathbb{P}[X_S^t \subsetneq X_S^{t+1}, X_U^t = v'_U, X_U^{t+1} = v_U] > 0. \end{aligned}$$

Assumption 4. *The features S are suitable for the downstream classification task (here Y refers to the downstream label):*

$$\operatorname{argmin}_{f \in \mathcal{F}} \mathbb{P}[f(X_S) \neq Y] = \operatorname{argmin}_{f \in \mathcal{F}, g \in \mathcal{G}} \mathbb{P}[(f \circ g)(X) \neq Y]$$

Intuitively, the first two assumptions guarantee that the feature-set S is an optimal choice of representation for the OCP pretraining objective (when \mathcal{G} is the class of feature selector representations), and the third (more technical) assumption guarantees that the optimum is unique (e.g., by preventing the possibility that pretraining leaves out a feature in S that is redundant for the ordering objective, but useful downstream). The last assumption ensures that the features S are suitable for the downstream classification task on the population: the loss achievable by the best $f \in \mathcal{F}$ using X_S as the representation is the same as the loss achievable by using the best (f, g) pair.

When these assumptions are satisfied, we prove a finite-sample bound for a model pretrained using OCP. The bound only depends on the VC-dimension of the *downstream* hypothesis class, $\text{VC}(\mathcal{F})$, rather than on $\text{VC}(\mathcal{F} \times \mathcal{G})$.

Like some results in the nonlinear ICA literature (e.g., Hyvärinen and Morioka (2017)) our results only apply in the regime where there is enough unlabeled data to identify the “correct” representation. It’s then immediate that only \mathcal{F} factors in to the labeled-data dependence. However, our model also allows us to give upper bounds on the *amount* of unlabeled data required to reach that regime. This allows us to more rigorously study other aspects of contrastive pretraining, such the role of *bias* in the negative distribution, which has been shown to affect the performance of other contrastive learning algorithms (Chuang et al., 2020).

We now give a simple example of a class of distributions satisfying these assumptions, grounded in our running application of health time-series data. Despite its simplicity, our findings suggest that this model allows for several interesting phenomena that also occur in practice, which could make it useful for further study of contrastive learning methods on time-series.

3.1 Extraction example

A common task in clinical informatics is to extract for each time t the patient’s structured disease stage, which enables downstream clinical research (Kehl et al., 2019, 2020). Each time point X^t could be an

encoding of the clinical note from a patient’s visit at time t . Let $Y^t \in \{0, 1\}$ be the observed label for time point X^t . The end goal is to train a model f over a representation g to minimize the downstream risk:

$$\underset{(f, g)}{\text{minimize}} R(f, g) := \mathbb{E}_{(X^T, Y^T)} [f(g(X^T)) \neq Y^T].$$

Here we make a prediction for every time point, and the expectation is over the time index T as well as the trajectory (X, Y) .

Model. For each $A \subset [d]$, we denote by X_A^t the random variable corresponding to indices A at time t . Suppose the set of feature indices $[d]$ is partitioned into three types of features:

- A set $S \subset [d]$ of time-irreversible features. We also assume that each $i \in S$ has a nonzero probability of activating on its own, without the other features in S . That is, for each $i \in S$ there exists t with $\mathbb{P}[X_i^t = 0, X_i^{t+1} = 1, X_{S \setminus \{i\}}^t = X_{S \setminus \{i\}}^{t+1}] > 0$. This ensures that assumption 3 is satisfied. Such features include the onset/progression of chronic conditions and markers of aging (Pierson et al., 2019). For example, appearance of the word “metastasis” in a clinical note.
- Noisy versions \hat{S} of S : for each $j \in \hat{S}$, there exists $i \in S$ with $\mathbb{P}[X_j^t = X_i^t] = (1 - \epsilon_i)$, with $\epsilon_i > 0$, for all t . Additionally, X_j^t is conditionally independent of the other variables (for all times) given its parent variable X_i^t . For example, the presence of certain interactions with the health system—such as deciding to attend physical therapy—may be a noisy reflection of the patient’s true disease state, which is captured by X_S^t .
- Background, reversible features B : features such that for all t and all $v, v' \in \{0, 1\}^d$,

$$\begin{aligned} \mathbb{P} \left[(X_B^t, X_{[d] \setminus B}^t) = (v_B, v_{[d] \setminus B}), \right. \\ \left. (X_B^{t+1}, X_{[d] \setminus B}^{t+1}) = (v'_B, v'_{[d] \setminus B}) \right] = \\ \mathbb{P} \left[(X_B^t, X_{[d] \setminus B}^t) = (v'_B, v_{[d] \setminus B}), \right. \\ \left. (X_B^{t+1}, X_{[d] \setminus B}^{t+1}) = (v_B, v'_{[d] \setminus B}) \right] \end{aligned}$$

Consider, for example, common words such as “and”, “chart”, etc., in a clinical note, whose presence or absence gives no order information.

Note that we *do not* make any independence assumptions between the features in this example other than the ones mentioned above.

We prove in Theorem A.1 that this example satisfies Assumptions 1-3. Assumption 4 is true by design when

e.g. $Y^t = f(X_S^t)$ (or a noisy version thereof). We now provide a simple finite-sample bound for OCP when Assumptions 1-4 are satisfied and the representation class \mathcal{G} is the set of *feature selectors*.

3.2 Finite sample bound

Suppose we observe a large set of m *unlabeled* data points $\{X_i\}_{i=1}^m$ drawn independently from the marginal distribution of X , and a much smaller set of n *labeled* data $\{(X_i, Y_i)\}_{i=1}^n$ drawn independently from the joint distribution of X and the downstream label Y (now we use Y to denote the downstream label rather than the pre-training label).

Let the representation hypothesis class $\mathcal{G} = \{U \subset [d] : |U| = |S| = d_0\}$. That is, our representation will *select features* $U \subset [d]$ to be used downstream. We identify each set of indices $g \in \mathcal{G}$ with the mapping $\{0, 1\}^d \rightarrow \{0, 1\}^{d_0}$ given by projection onto those indices. Let \mathcal{F} be the downstream hypothesis class, with $f : \{0, 1\}^{d_0} \rightarrow \{0, 1\}$.² The mapping $f \circ g : \{0, 1\}^d \rightarrow \{0, 1\}$ first selects the input features represented by g , then passes the values of these features through f .

Our end goal is to design a learning algorithm \mathcal{A} with a *downstream excess risk bound*. If (\hat{f}, \hat{g}) is the hypothesis output by \mathcal{A} , we want an upper bound on the excess risk: $R(\hat{f}, \hat{g}) - \min_{f \in \mathcal{F}, g \in \mathcal{G}} R(f, g)$. If we let \mathcal{A}_{ds} be empirical risk minimization (ERM) over $\mathcal{F} \times \mathcal{G}$ on the small labeled sample (i.e., the method directly optimizing the downstream objective over $\mathcal{F} \times \mathcal{G}$ without pre-training),³ a standard result (e.g., Shalev-Shwartz and Ben-David (2014, Theorem 6.8)) implies:

$$R(\hat{f}, \hat{g}) - \min_{f \in \mathcal{F}, g \in \mathcal{G}} R(f, g) \leq O\left(\sqrt{\frac{\text{VC}(\mathcal{F} \times \mathcal{G})}{n}}\right) \quad (2)$$

with high probability over the sampling of the data. On the other hand, let \mathcal{A}_{pt} be the algorithm that first uses unlabeled data to *pre-train* a representation \hat{g} by minimizing (1), then minimizes the downstream risk over $\mathcal{F} \times \{\hat{g}\}$ (i.e., a 2-phase ERM learner). The following theorem states that under Assumptions 1-4, we

²Assume \mathcal{F} is closed under permutations of the input dimensions. This ensures that we only need to identify the features belonging to S , and don't need to put them in a particular order to do well at downstream prediction. The set of linear hypotheses has this property.

³Assume for simplicity that for each trajectory (X, Y) , a single time T is chosen uniformly at random and $\{(X^T, Y^T)\}$ are passed to the learner, so the learner sees i.i.d. samples. A more detailed treatment would handle the dependence between multiple time points to get bounds that decrease as $\tilde{O}(1/\sqrt{n\tau})$ when possible (e.g., Mohri and Rostamizadeh, 2010).

can give a more parsimonious upper bound on the excess risk. In what follows, we use $R_{ord}(g)$ to refer to $\inf_{c \in \mathcal{C}} R_{ord}(c, g)$, and likewise for \hat{R}_{ord} . We give details on the choice of \mathcal{C} in Appendix A. Since \mathcal{G} is the class of feature selectors and we assumed \mathcal{F} is closed under permutations of the input dimensions, we sometimes replace $g \in \mathcal{G}$ below with sets $U \subset [d]$.

Theorem 1. *Suppose Assumptions 1-4 are satisfied, and let $\epsilon_0 = \min_{U: S \subseteq U} R_{ord}(U) - R_{ord}(S)$ be the difference in OCP error between S and the next-best representation. Suppose we have an unlabeled dataset of m i.i.d. pretraining points $\{(Z_i, Y_i)\}_{i=1}^m$, with:*

$$m > \frac{2\left(\log\binom{d}{|S|} + \log\frac{4}{\delta}\right)}{\epsilon_0^2},$$

and a labeled dataset of n downstream training points $\{(X_i^t, Y_i^t)\}_{i=1}^n$. Let \mathcal{G} be all sets of size $|S|$ features chosen from the full set of d features. Let \hat{g} be the minimizer of the empirical OCP pretraining objective:

$$\hat{g} = \operatorname{argmin}_{g \in \mathcal{G}} \hat{R}_{ord}(g)$$

Let \hat{f} be the minimizer of the empirical downstream objective when using the fixed representation \hat{g} :

$$\hat{f} = \operatorname{argmin}_{f \in \mathcal{F}} \hat{R}(f, \hat{g})$$

Then for any $\delta > 0$, with probability at least $1 - \delta$, (\hat{f}, \hat{g}) has excess risk:

$$R(\hat{f}, \hat{g}) - \inf_{(f, g) \in \mathcal{F} \times \mathcal{G}} R(f, g) \leq O\left(\sqrt{\frac{\text{VC}(\mathcal{F}) + \log\frac{1}{\delta}}{n}}\right) \quad (3)$$

Proof (sketch). Assumptions 1-2 are used to show that S is one of the optima for the population OCP objective, i.e., that

$$S \in \operatorname{argmin}_{g \in \mathcal{G}} R_{ord}(g).$$

Then, Assumption 3 is used to show that S is actually the *unique* optimum of size $|S|$ (and hence $\epsilon_0 > 0$). The condition on m (the amount of pretraining data), together with a standard learning bound for finite classes, is enough to guarantee that OCP identifies S with high probability over the sampling of the pre-training data. That is, the choice of m guarantees that with high probability,

$$S = \operatorname{argmin}_{g \in \mathcal{G}} \hat{R}_{ord}(g).$$

Assumption 4 guarantees that choosing S in the pre-training step does not incur additional error on the downstream task compared to the optimal (f, g) pair

(since it states that S is the optimal g for the *downstream* task). The result follows from a standard uniform convergence bound (e.g., Shalev-Shwartz and Ben-David (2014, Theorem 6.8)) applied to \mathcal{F} . The full proof is given in Appendix A. \square

The pretrain + finetune bound (3) has a better dependence on the labeled dataset size n than the downstream ERM learner bound (2). The large unlabeled dataset allows for the learning of a good representation \hat{g} without using any labeled data. Even in this simple feature selection setting, this bound may be much tighter than the direct-downstream bound when \mathcal{F} is a fairly complex hypothesis class and $d_0 \ll d$. Even for \mathcal{F} linear, $\mathcal{F} \times \mathcal{G}$ is roughly $\Theta(d_0 \log(d/d_0))$ (see e.g. (Abramovich and Grinshtein, 2018)), so (3) can even save over (2) in this case. In fact, we show in Section 4 that OCP can improve the performance of sparse linear models in a real-world low-labeled data setting (compared to direct downstream prediction without pre-training).

In this section we gave a simple example of a class of distributions, together with an *assumption* linking the distribution to the downstream task (Assumption 4—the time-irreversible features are the most useful ones for downstream classification) for which we can *prove* that OCP gives a more parsimonious bound on the labeled sample complexity. While the example in Section 3.1 seems straightforward, we show now that it still admits interesting behavior. In particular, there are distributions that satisfy Assumptions 1-3, but where PCL and OCP learn different representations.

PCL versus OCP: different infinite-data optima. There are examples of the model from Section 3.1 where PCL and OCP learn provably different representations even with infinite unlabeled samples, despite the minor difference in their sampling schemes.⁴ Intuitively, the existence of a periodic feature (such as a procedure always performed at a particular time of day) is strongly predictive of whether two samples are consecutive, but need not be predictive of whether a pair of consecutive samples are in the correct order. Concretely, consider a feature X_i such that $X_i^{t+1} = 1 - X_i^t$, and $X_i^1 \sim \text{Ber}(0.5)$. Inclusion of this feature doesn’t violate Assumptions 1-3—indeed, X_i would qualify as a “background” feature under our model—so Theorem 1 guarantees that OCP finds the correct representation. However, in PCL, every non-consecutive sample is a negative. But only

⁴The example we use includes *nonstationary* features, which violates the assumptions under which PCL is proven in Hyvärinen and Morioka (2017) to find the “right” representation, so this does not contradict those results.

non-consecutive samples can have $X_i^t = X_i^{t'}$, so X_i^t is helpful for the PCL objective. We treat this example more formally in Appendix A, but our synthetic results in Section 4 also show that a background periodic feature can affect the PCL representation.

“Debiased” negatives. PCL has some negatives that are actually in the correct order. Prior work on contrastive learning has called this “bias” in the negative distribution (Chuang et al., 2020). What’s the role of this “bias?” Does it affect the learned representations? Does it affect the amount of *unlabeled* data required to find a good representation? For distributions satisfying assumptions 1-3 and when \mathcal{G} is the class of feature-selectors, we answer these questions in the negative and positive, respectively.

In particular, consider the analogue of OCP that instead of *always* choosing $(W+1, W)$ when $Y = -1$ (Y as used in OCP, not the downstream label), instead just chooses a random pair (W, W') with $|W - W'| = 1$ (i.e., a random consecutive pair). We refer to this as OCP-biased, since some of the negatives are actually in the correct order. However, the following theorem shows the estimator obtained by minimizing this objective is *not* biased in a statistical sense:

Theorem 2 (informal). *When assumptions 1-3 are satisfied, S is also the unique optimal representation for OCP-biased.*

However, it *does* affect the bound on unlabeled sample complexity required to obtain a good representation:

Proposition 1 (informal). *The upper bound on the sample complexity required for OCP-biased to identify S is worse than the upper bound for OCP.*

While this proposition only compares upper bounds, our synthetic experiments (see Figure 2) indicate that OCP is more sample-efficient than OCP-biased. We prove these results in Appendix A.

4 Experiments

4.1 Synthetic data

We demonstrate the importance of negative sampling over two synthetic datasets from the model in Section 3. Each distribution contains $|S| = 4$ alongside a number of noisy features. We generate pre-training datasets of different sizes (50 to 16,000) and sample pairs from each dataset according to OCP, PCL, and OCP-biased. We then conduct a logistic regression with L0 penalty over the sample pairs and analyze how many variables in S were correctly recovered. The top panel of Figure 2 shows a distribution where PCL

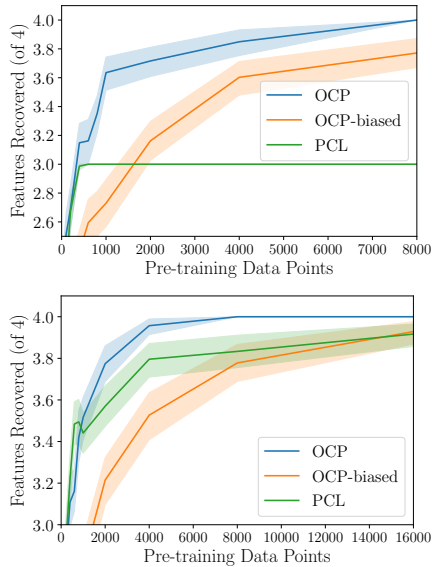


Figure 2: Synthetic experiments admitting different behavior across pre-training setups. (Top): PCL is unable to ever recover all four true features in S ; (Bottom): PCL recovers the true representation, but requires higher sample complexity than OCP.

does not recover S in the infinite data limit—this distribution includes a periodic background feature in B that is selected by PCL. The bottom panel of Figure 2 shows a distribution where PCL is able to recover all of S , but requires a larger sample complexity than OCP. In both cases, OCP and OCP-biased find the same representation, but the former has better dependence on unlabeled data. We provide the details and explanations for these experiments in Appendix B.

4.2 Real-world data

We show OCP yields significant improvements in the low-label regime on extraction from clinical notes.

Progression dataset. We utilize a dataset of fully de-identified clinical notes from Memorial Sloan Kettering Cancer Center. This research was reviewed by the MIT Committee on the Use of Humans as Experimental Subjects and determined to be IRB-exempt. The dataset contains data for 82,839 patients with cancer, with a median of 12 radiology notes each. Each radiology note focuses on one body area (e.g., chest CT scan). In addition, we have a subset of 135 patients with progressive lung cancer with 1095 labeled radiology notes. Each note was labeled post-hoc by a dedicated thoracic oncologist as ‘indicating progression’ (19%), ‘not indicating progression’ (79.5%), or ‘ambiguous’ (1.5%).

Experimental setup. We investigate extraction of these binary progression labels from the *Impression*

section of the note. The labeled data was split via 5-fold cross-validation: each fold contained sets of sizes 64% (train), 16% (validation), and 20% (test); for a given fold, no patient examples were ever split between sets. On each fold, we used the test set to benchmark models trained using different amounts of the labeled training data: from just 5 training patients ($\frac{1}{16}$) to all of the training patients. We excluded patients with downstream labels from pretraining. For contrastive pre-training schemes, a pretraining window pair was sampled once per each unique body area (e.g. chest, brain) that was scanned at least twice, capped at five locations per patient. This resulted in $\approx 158,000$ samples for pretraining.

Pre-training for feature selection. We first validate our modeling assumptions from Section 3 using a linear model. The goal of this section is to roughly validate our assumptions and the setup of our theoretical model. We compare downstream progression extraction performance of (i) a vanilla logistic regression model and (ii) a logistic regression model only using the features selected by OCP. We test on all five folds for five training dataset fractions.

For each experiment, our dataset is featurized using the unigrams and bigrams that occur in at least 5% of the labeled training data set. They are vectorized using the *term frequency-inverse document frequency* weighting scheme, via scikit-learn (Pedregosa et al., 2011, BSD 3-clause license). We conduct feature selection as an optional intermediate step preceding progression extraction. For OCP, we train a logistic regression model with L1 penalty over the 158,000 pre-training pairs of consecutive radiology notes. The regularization constant was set such that there were 50 ± 5 features with nonzero weights. In addition to OCP-derived features, we select the 50 most common features, and 5 random subsets of 50 features to serve as a comparison.

We train scikit-learn logistic regression models for downstream progression extraction over each feature set; further details are in Appendix C. Results can be seen in Figure 3a. Even with a simple bag-of-words representation, feature selection with OCP outperforms *directly* training a tuned logistic regression model on the available labeled data (“All features”), especially for small dataset sizes. A paired *t*-test finds that the model with OCP-selected features is significantly better than the direct-downstream model on a sixteenth of the data ($p < 0.05$). Note that selecting the most common features or a random set of features does not compare, showing that OCP does not improve performance by simply reducing the feature dimension in a redundant space.

(a) Mean note-level AUC of regularized logistic regression over different dataset sizes. Averaged over the 5 folds, performance was optimal for each dataset size when restricted to the features with nonzero coefficients recovered by OCP.

| <i>Available features</i> | <i>Fraction of training data</i> | | | | |
|---------------------------|----------------------------------|------------|------------|------------|-------------|
| | 1 | 1/2 | 1/4 | 1/8 | 1/16 |
| OCP subset | 0.864 | 0.860 | 0.847 | 0.808 | 0.786 |
| All features | 0.856 | 0.851 | 0.818 | 0.723 | 0.726 |
| Most common | 0.767 | 0.767 | 0.728 | 0.687 | 0.658 |
| Random subset | 0.740 | 0.747 | 0.727 | 0.639 | 0.634 |

(b) Example features that OCP selected (top) or excluded (bottom) for downstream prediction.

| | |
|-----------------------|---|
| Selected Terms | mass, increased, decreased, stable, change, new, suspicious |
| Excluded Terms | discussed, imaging, left, mri, also, follow |

Figure 3: Linear representation space experiment to validate assumptions of our model apply to real-world data. Quantitatively, we find downstream wins from restricting the model feature space to those found useful for the order-contrastive task. Qualitatively, the features important for the order pre-training are the same we would expect to be useful for the downstream extraction task.

We manually examined the OCP-selected features and their coefficients (Figure 3b). The features included (e.g. *increased*, *decreased*) strongly indicate disease progression, while those discarded (e.g. *discussed*) largely seem to be noise. Of the nonzero coefficients, 76% have a positive weight; this indicates that the pre-training model focuses mostly on features that have been *turned on* to conduct the ordering task, fitting with our motivating theoretical setting.

Nonlinear representations. We now study the use of OCP for pre-training nonlinear representations. We compare performance of a BERT model pre-trained using OCP to several other self-supervision methods. We investigate the BERT base model and the BERT base model after it is pre-trained using: (i) *FT LM*: fine-tuned masked language modeling over an equivalent number of impressions, (ii) *Pt-Contrastive*: a patient-level contrastive objective (identical positive sampling to OCP and PCL, but each negative is a random note of the same note type from a *different patient*, similar to Diamant et al. (2022)), (iii) *PCL*: contrastive pre-training with PCL sampling (each negative is a random pair of notes of the same type from the same patient), (iv) *OCP*: contrastive pre-training with OCP sampling (each negative is a pair of notes of the same type in the incorrect order). All pre-training is conducted over three seeds, and all three contrastive objectives were trained with the same number of pairs (158,000). Implementation for language modeling and contrastive pre-training came from Wolf et al. (2020, Apache-2.0 License) with full details in Appendix C. After model pre-training/fine-tuning, the self-supervised representation layers were frozen, and a single L2-regularized linear layer was added on top. The goal of freezing was to isolate the effect of pre-training to understand representation quality, due to the instability of training BERT on small downstream tasks (Zhang et al., 2021).

Results can be seen in Table 1. The top row shows

that OCP has only a modest drop in performance even when trained on the data from just 5 patients ($\frac{1}{16}$). Since correlations exist in AUC across the 5 folds, 3 seeds, and 5 dataset sizes, standard statistical comparison testing is inappropriate. Instead, we present the mean increase in AUC from OCP, as well as the percentage of the time OCP outperformed the comparisons. Unsurprisingly, BERT alone (trained on non-clinical text) unsurprisingly does not perform well out-of-the-box; fine-tuning with language modeling improves performance, but still suffers in the low data regime. Among the contrastive objectives, the cross-patient objective is the weakest, which may follow since its pre-training task was the easiest (82% accuracy on validation). It could rely on features that differed between patients, instead of being forced to focus on the temporal features that differed *within* a patient’s timeline. PCL performs equivalently to OCP at large data sizes, but at the smaller data set sizes, it loses to OCP a large majority of the time.

5 Related work

Order pre-training. Others have found order-based self-supervision useful for more complex time-series data, but without theoretical study. For example, learning the order of frames within a video yields representations useful for downstream activity classification (Fernando et al., 2015; Misra et al., 2016; Lee et al., 2017; Wei et al., 2018). Most similar to our work, Hyvärinen and Morioka (2017) introduced *permutation-contrastive learning* (PCL) and proved *nonlinear* identifiability for representations learned using PCL in an ICA setting. That is, they gave distributional conditions where PCL provably recovers the “correct” nonlinear representation of the input given infinite unlabeled data. Our theoretical and empirical results indicate that there can be nontrivial differences between OCP and PCL’s downstream performance. Deeper understanding of what data distribu-

Table 1: Performance of deep methods on cancer progression extraction. The first row contains the mean AUC of OCP \pm its std dev. The following rows contain the mean AUC advantage of OCP over each comparison method, and the percentage of time OCP outperforms that method, across the 3 seeds and 5 folds.

| AUC diff. (OCP Win %) | Fraction of training data | | | | |
|-----------------------------|---------------------------|----------------|----------------|----------------|----------------|
| | 1 | 1/2 | 1/4 | 1/8 | 1/16 |
| OCP AUC | 0.87 \pm .03 | 0.86 \pm .04 | 0.84 \pm .04 | 0.82 \pm .03 | 0.81 \pm .03 |
| OCP – BERT | 0.08 (93%) | 0.12 (100%) | 0.12 (100%) | 0.18 (100%) | 0.22 (100%) |
| OCP – FT LM | 0.03 (80%) | 0.04 (82%) | 0.04 (82%) | 0.08 (93%) | 0.10 (89%) |
| OCP – Pt-Contrastive | 0.03 (86%) | 0.03 (77%) | 0.05 (91%) | 0.09 (91%) | 0.12 (97%) |
| OCP – PCL | 0.00 (53%) | 0.00 (46%) | 0.03 (64%) | 0.03 (76%) | 0.06 (87%) |

tions and downstream tasks are “right” for PCL versus OCP (and for other contrastive sampling methods) is an interesting direction for future theoretical study.

Pre-training for medical time-series. Several other pre-training objectives have been explored on clinical time-series data. A contrastive learning setup similar to our *patient-contrastive* baseline has shown promising results on electrocardiograph signals (Diamant et al., 2022; Kiyasseh et al., 2021). Banville et al. (2021) studied a contrastive objective for electroencephalography signals, in which windows of a signal are judged to be similar if they occur within a certain time gap, and dissimilar if they are far away in time. Intuitively, this objective is well-suited to data where the true representation “changes slowly” with time, as with Franceschi et al. (2019) (discussed in Section 1). Other objectives include auto-encoding (Fox et al., 2019) and masked prediction over text and tabular data, to mixed results (Steinberg et al., 2021; Huang et al., 2019; Yoon et al., 2020; McDermott et al., 2021). Multi-task pre-training supplies improvements, but unlike our work, it relies on additional labeled data from closely-related downstream tasks (McDermott et al., 2021).

Self-supervision theory. Like our work, Saunshi et al. (2019); Liu et al. (2021) and Tosh et al. (2021a,b) give downstream finite-sample error bounds for representations learned using particular contrastive learning objectives. Our motivating theoretical setting and proof techniques are simpler than the ones considered in these works, but we show that our setup in Section 3 is (i) complex enough to allow for some of the same nontrivial behavior observed by contrastive methods in practice (Sections 3.2, 4.1) and (ii) it has some practical applications (Section 4.2).

6 Limitations and Conclusion

We have shown both theoretically and empirically that order-contrastive pre-training is an effective self-supervised method for certain types of time-series data and downstream tasks. On real-world longitudinal

health data, we find that representations from OCP significantly outperform others, including the similar PCL, in the small data regime. Concretely, being able to structure variables from longitudinal, label-scarce, data in health records could enable us to evaluate large scale retrospective datasets and potentially inform future clinical trials and patient care.

However, OCP is not always suitable. For example, cases of temporal leakage (e.g., the date in a note) can lead to weak OCP (and PCL) representations downstream, since they provide a shortcut during pre-training. While dates are straightforward to censor, more complex global nonstationarities irrelevant to downstream tasks would present a challenge to these methods. We additionally wish to emphasize that not all clinical tasks have time-irreversible expressions in the data (e.g., acute/temporary conditions, such as pregnancy), so the motivating model assumptions (particularly Assumption 4) should be considered before applying OCP.

Our theoretical setup and results in Section 3 also serve to highlight that contrastive pre-training methods can be very sensitive to the precise sampling details, and provide a simple model for studying these details that is still complex enough to capture some empirical phenomena. This suggests that obtaining broader theoretical guidelines for selecting a contrastive distribution is an interesting direction for future work.

Acknowledgements

We would like to thank Rebecca Boiarsky for helpful comments on the manuscript, Richard Do for advising as a subject matter expert on radiology reports, and the Institute of Advanced Study for their hospitality in hosting MA. DS and HL were partially supported by AitF award CCF-1723344, and MA by a Takeda Fellowship. The authors acknowledge funding for Memorial Sloan Kettering Cancer Center received through the NIH/NCI institutional P30 CA008748 grant.

References

- Abramovich, F. and Grinshtein, V. (2018). High-dimensional classification by sparse logistic regression. *IEEE Transactions on Information Theory*, 65(5):3068–3079.
- Banville, H., Chehab, O., Hyvärinen, A., Engemann, D.-A., and Gramfort, A. (2021). Uncovering the structure of clinical eeg signals with self-supervised learning. *Journal of Neural Engineering*, 18(4):046020.
- Bleackley, J. and Kim, S. Y. R. (2013). The merit and agony of retrospective chart reviews: A medical student’s perspective. *British Columbia Medical Journal*, 55:374–375.
- Chuang, C.-Y., Robinson, J., Lin, Y.-C., Torralba, A., and Jegelka, S. (2020). Debaised contrastive learning. In *NeurIPS*.
- Diamant, N., Reinertsen, E., Song, S., Aguirre, A. D., Stultz, C. M., and Batra, P. (2022). Patient contrastive learning: A performant, expressive, and practical approach to electrocardiogram modeling. *PLOS Computational Biology*, 18(2):e1009862.
- Fernando, B., Gavves, E., Oramas, J. M., Ghodrati, A., and Tuytelaars, T. (2015). Modeling video evolution for action recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5378–5387.
- Fox, I., Rubin-Falcone, H., and Wiens, J. (2019). Learning through limited self-supervision: Improving time-series classification without additional data via auxiliary tasks.
- Franceschi, J.-Y., Dieuleveut, A., and Jaggi, M. (2019). Unsupervised scalable representation learning for multivariate time series. In *Advances in Neural Information Processing Systems*, volume 32.
- Goroshin, R., Bruna, J., Tompson, J., Eigen, D., and LeCun, Y. (2015). Unsupervised learning of spatiotemporally coherent metrics. In *Proceedings of the IEEE international conference on computer vision*, pages 4086–4093.
- Huang, K., Altosaar, J., and Ranganath, R. (2019). Clinicalbert: Modeling clinical notes and predicting hospital readmission. *arXiv preprint arXiv:1904.05342*.
- Hyvärinen, A. and Morioka, H. (2017). Nonlinear ica of temporally dependent stationary sources. In *Artificial Intelligence and Statistics*, pages 460–469. PMLR.
- Kehl, K. L., Elmarakeby, H., Nishino, M., Van Allen, E. M., Lepisto, E. M., Hassett, M. J., Johnson, B. E., and Schrag, D. (2019). Assessment of deep natural language processing in ascertaining oncologic outcomes from radiology reports. *JAMA oncology*, 5(10):1421–1429.
- Kehl, K. L., Xu, W., Lepisto, E., Elmarakeby, H., Hassett, M. J., Van Allen, E. M., Johnson, B. E., and Schrag, D. (2020). Natural language processing to ascertain cancer outcomes from medical oncologist notes. *JCO Clinical Cancer Informatics*, 4:680–690.
- Kiyasseh, D., Zhu, T., and Clifton, D. A. (2021). Clocs: contrastive learning of cardiac signals across space, time, and patients. In *International Conference on Machine Learning*, pages 5606–5615. PMLR.
- Lee, H.-Y., Huang, J.-B., Singh, M., and Yang, M.-H. (2017). Unsupervised representation learning by sorting sequences. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 667–676.
- Liu, B., Ravikumar, P., and Risteski, A. (2021). Contrastive learning of strong-mixing continuous-time stochastic processes. In *International Conference on Artificial Intelligence and Statistics*, pages 3151–3159. PMLR.
- McDermott, M., Nestor, B., Kim, E., Zhang, W., Goldenberg, A., Szolovits, P., and Ghassemi, M. (2021). A comprehensive ehr timeseries pre-training benchmark. In *Proceedings of the Conference on Health, Inference, and Learning, CHIL ’21*, page 257–278, New York, NY, USA. Association for Computing Machinery.
- Misra, I., Zitnick, C. L., and Hebert, M. (2016). Shuffle and learn: unsupervised learning using temporal order verification. In *European Conference on Computer Vision*, pages 527–544. Springer.
- Mobahi, H., Collobert, R., and Weston, J. (2009). Deep learning from temporal coherence in video. In *Proceedings of the 26th Annual International Conference on Machine Learning*, pages 737–744.
- Mohri, M. and Rostamizadeh, A. (2010). Stability bounds for stationary φ -mixing and β -mixing processes. *Journal of Machine Learning Research*, 11(2).
- Mohri, M., Rostamizadeh, A., and Talwalkar, A. (2018). *Foundations of machine learning*. MIT press.
- Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., and Duchesnay, E. (2011). Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830.

- Pierson, E., Koh, P. W., Hashimoto, T., Koller, D., Leskovec, J., Eriksson, N., and Liang, P. (2019). Inferring multidimensional rates of aging from cross-sectional data. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 97–107. PMLR.
- Robinson, J., Chuang, C.-Y., Sra, S., and Jegelka, S. (2021). Contrastive learning with hard negative samples. *ICLR*.
- Saunshi, N., Plevrakis, O., Arora, S., Khodak, M., and Khandeparkar, H. (2019). A theoretical analysis of contrastive unsupervised representation learning. In *International Conference on Machine Learning*, pages 5628–5637. PMLR.
- Shalev-Shwartz, S. and Ben-David, S. (2014). *Understanding machine learning: From theory to algorithms*. Cambridge university press.
- Steinberg, E., Jung, K., Fries, J. A., Corbin, C. K., Pfohl, S. R., and Shah, N. H. (2021). Language models are an effective representation learning technique for electronic health record data. *Journal of Biomedical Informatics*, 113:103637.
- Tosh, C., Krishnamurthy, A., and Hsu, D. (2021a). Contrastive estimation reveals topic posterior information to linear models. *Journal of Machine Learning Research*, 22(281):1–31.
- Tosh, C., Krishnamurthy, A., and Hsu, D. (2021b). Contrastive learning, multi-view redundancy, and linear models. In *Algorithmic Learning Theory*, pages 1179–1206. PMLR.
- Wei, D., Lim, J. J., Zisserman, A., and Freeman, W. T. (2018). Learning and using the arrow of time. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8052–8060.
- Wiskott, L. and Sejnowski, T. J. (2002). Slow feature analysis: Unsupervised learning of invariances. *Neural computation*, 14(4):715–770.
- Wolf, T., Debut, L., Sanh, V., Chaumond, J., Delangue, C., Moi, A., Cistac, P., Rault, T., Louf, R., Funtowicz, M., Davison, J., Shleifer, S., von Platen, P., Ma, C., Jernite, Y., Plu, J., Xu, C., Scao, T. L., Gugger, S., Drame, M., Lhoest, Q., and Rush, A. M. (2020). Transformers: State-of-the-art natural language processing. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.
- Xia, F. and Yetisgen-Yildiz, M. (2012). Clinical corpus annotation: challenges and strategies. In *Proceedings of the Third Workshop on Building and Evaluating Resources for Biomedical Text Mining (BioTxtM’2012) in conjunction with the International Conference on Language Resources and Evaluation (LREC), Istanbul, Turkey*.
- Yoon, J., Zhang, Y., Jordon, J., and van der Schaar, M. (2020). Vime: Extending the success of self-and semi-supervised learning to tabular domain. *Advances in Neural Information Processing Systems*, 33.
- Zhang, T., Wu, F., Katiyar, A., Weinberger, K. Q., and Artzi, Y. (2021). Revisiting few-sample BERT fine-tuning. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net.

Supplementary Material: Leveraging Time Irreversibility with Order-Contrastive Pre-training

A Theory details

In this section, we provide proofs for the simple model introduced in Section 3. In particular, we prove the finite sample bound for OCP by first proving that the driver features S are the *unique optimal representation*, and then we apply standard uniform convergence arguments to obtain the bound (3).

We also show how so-called *bias* (Chuang et al., 2020) in the distribution of negative samples can affect OCP. In particular, suppose that instead of choosing windows $(W, W + 1)$ as positive and $(W + 1, W)$ as negatives, we chose $(W, W + 1)$ as positives and a random pair (W, W') with $|W - W'| = 1$ as negatives. This is analogous to PCL’s negative sampling, where some of the negatives are still in the correct order. The difference with PCL is just that all negatives are still consecutive. For the model in Section 3, we prove that (i) in the infinite-data limit, this “biased” version recovers the same representation as S (so in fact, the estimator obtained by minimizing this objective is *not* biased in a statistical sense) and (ii) the bound on *unlabeled* sample complexity required to find a good representation is worse for this biased version of OCP than for the unbiased version. This gives more theoretical evidence for the value of *de-biasing* the negative distribution in contrastive learning, where possible: *even if it doesn’t change the representation learned with infinite unlabeled data, de-biasing the contrastive distribution can improve unlabeled sample complexity.*

A.1 Assumptions and example class of distributions

Here we provide simple assumptions under which a set of features S is the unique optimal solution to (1) for the model from Section 3. For simplicity in this section, we only consider consecutive windows and the case where τ (the sequence length) is the same for all sequences, but all of our results generalize (with suitable modifications to these assumptions) to the case where the positive and negative distributions over windows are symmetric up to ordering of the elements (correct versus incorrect), and to the case with a distribution over τ .⁵

Assumption A.1. For all $i \in S$ and all t , $\mathbb{P}[X_i^t = 1, X_i^{t+1} = 0] = 0$.

Assumption A.2. For all $v \in \{0, 1\}^d$ and $v' \in \{0, 1\}^d$ with $v_S = v'_S$, and for all t , $\mathbb{P}[X^t = v, X^{t+1} = v'] = \mathbb{P}[X^t = v', X^{t+1} = v]$.

Assumption A.3. For all U such that $S \not\subseteq U$, there exists t and $v \in \{0, 1\}^d$, $v' \in \{0, 1\}^d$, with $v_{U \cap S} = v'_{U \cap S}$ and:

$$\min(\mathbb{P}[X_S^t \subsetneq X_S^{t+1}, X_U^t = v_U, X_U^{t+1} = v'_U]; \mathbb{P}[X_S^t \subsetneq X_S^{t+1}, X_U^t = v'_U, X_U^{t+1} = v_U]) > 0.$$

Assumption A.4. The features S are suitable for the downstream classification task:

$$\operatorname{argmin}_{f \in \mathcal{F}} \mathbb{P}[f(X_S) \neq Y] = \operatorname{argmin}_{f \in \mathcal{F}, g \in \mathcal{G}} \mathbb{P}[(f \circ g)(X) \neq Y]$$

Now we prove that the example from Section 3 satisfies these assumptions (in particular, Assumptions A.1-A.3).

Theorem A.1. Partition the indices $[d]$ into three sets (S, \hat{S}, B) satisfying the following assumptions:

- *Time-irreversible features S satisfying Assumption 1.* We also assume that each $i \in S$ has a nonzero probability of activating on its own, without the other features in S . That is, for each $i \in S$ there exists $t \in [\tau - 1]$, with $\mathbb{P}[X_i^t = 0, X_i^{t+1} = 1, X_{S \setminus \{i\}}^t = X_{S \setminus \{i\}}^{t+1}] > 0$.

⁵For example, if for all T and $t \in \{0, \dots, T-1\}$, $\mathbb{P}[X^t, X^{t+1} | \tau = T] = \mathbb{P}[X^t, X^{t+1} | \tau \geq t+1]$ (i.e., the distribution of any window pair only depends on the fact that the trajectory is still active, and not on the actual length) then assumptions remain the same.

- *Noisy versions \hat{S} of S* : for each $j \in \hat{S}$, there exists $i \in S$ with $\mathbb{P}[X_j^t = X_i^t] = (1 - \epsilon_i)$, with $\epsilon_i > 0$, for all t . Additionally, X_i^t is conditionally independent of the other variables (for all times) given its driver X_i^t : $\mathbb{P}[X_j^t, X_{\mathcal{I}_1}^1, \dots, X_{\mathcal{I}_\tau}^\tau | X_i^t] = \mathbb{P}[X_j^t | X_i^t] \mathbb{P}[X_{\mathcal{I}_1}^1, \dots, X_{\mathcal{I}_\tau}^\tau | X_i^t]$, where $\mathcal{I}_i = [d] \setminus j$ if $\hat{t} = t$ and $[d]$ otherwise. **FiXme Note: this is not precisely the right assumption; be more careful notationally to include all times**
- *Background, reversible features B* : features such that for all t and all $v, v' \in \{0, 1\}^d$,

$$\begin{aligned} & \mathbb{P} \left[(X_B^t, X_{[d] \setminus B}^t) = (v_B, v_{[d] \setminus B}), (X_B^{t+1}, X_{[d] \setminus B}^{t+1}) = (v'_B, v'_{[d] \setminus B}) \right] = \\ & \mathbb{P} \left[(X_B^t, X_{[d] \setminus B}^t) = (v'_B, v'_{[d] \setminus B}), (X_B^{t+1}, X_{[d] \setminus B}^{t+1}) = (v_B, v_{[d] \setminus B}) \right] \end{aligned}$$

This class of distributions satisfies Assumption A.1-A.3.

Proof. Assumption A.1 is satisfied by definition. For Assumption A.2, fix $v, v' \in \{0, 1\}^d$ with $v_S = v'_S$ and $t \in [\tau - 1]$. Then we have:

$$\mathbb{P}[X^t = v, X^{t+1} = v'] = \mathbb{P}[X_S^t = v_S, X_S^{t+1} = v'_S, X_{\hat{S}}^t = v_{\hat{S}}, X_{\hat{S}}^{t+1} = v'_{\hat{S}}, X_B^t = v_B, X_B^{t+1} = v'_B].$$

Using the conditional independence assumption for the features in \hat{S} , the right-hand-side factors to:

$$\mathbb{P}[X_S^t = v_S, X_S^{t+1} = v'_S, X_B^t = v_B, X_B^{t+1} = v'_B] \mathbb{P}[X_{\hat{S}}^t = v_{\hat{S}} | X_S^t = v_S] \mathbb{P}[X_{\hat{S}}^{t+1} = v'_{\hat{S}} | X_S^{t+1} = v'_S].$$

Because $v_S = v'_S$, trivially this is equal to:

$$\mathbb{P}[X_S^t = v'_S, X_S^{t+1} = v_S, X_B^t = v_B, X_B^{t+1} = v'_B] \mathbb{P}[X_{\hat{S}}^t = v_{\hat{S}} | X_S^t = v_S] \mathbb{P}[X_{\hat{S}}^{t+1} = v'_{\hat{S}} | X_S^{t+1} = v'_S].$$

The reversibility of the features in B (and summing over the full joint distribution) imply that the above is equal to:

$$\mathbb{P}[X_S^t = v'_S, X_S^{t+1} = v_S, X_B^t = v'_B, X_B^{t+1} = v_B] \mathbb{P}[X_{\hat{S}}^t = v_{\hat{S}} | X_S^t = v_S] \mathbb{P}[X_{\hat{S}}^{t+1} = v'_{\hat{S}} | X_S^{t+1} = v'_S].$$

Finally, note that because $v_S = v'_S$, X_S^t and X_S^{t+1} are identically distributed by the definition of features in \hat{S} , so:

$$\begin{aligned} & \mathbb{P}[X_{\hat{S}}^t = v_{\hat{S}} | X_S^t = v_S] \mathbb{P}[X_{\hat{S}}^{t+1} = v'_{\hat{S}} | X_S^{t+1} = v'_S] = \\ & \mathbb{P}[X_{\hat{S}}^t = v'_{\hat{S}} | X_S^t = v'_S] \mathbb{P}[X_{\hat{S}}^{t+1} = v_{\hat{S}} | X_S^{t+1} = v_S]. \end{aligned}$$

Combining with the previous equation and simplifying, we obtain

$$\mathbb{P}[X^t = v, X^{t+1} = v'] = \mathbb{P}[X^t = v', X^{t+1} = v],$$

which is Assumption A.2.

Finally, for Assumption A.3, fix $U \subset [d]$ with $S \not\subseteq U$, and fix some $i \in S \setminus U$. We assumed that each $i \in S$ has some probability of activating on its own, so there exists t with $\mathbb{P}[X_i^t = 0, X_i^{t+1} = 1, X_{S \setminus \{i\}}^t = X_{S \setminus \{i\}}^{t+1}] > 0$. In particular, this implies that there exist v, v' with $v_i = 0, v'_i = 1, v_{S \setminus \{i\}} = v'_{S \setminus \{i\}}$ and $\mathbb{P}[X^t = v, X^{t+1} = v'] > 0$. By summing over variables in this joint distribution, that implies $\mathbb{P}[X_i^t = 0, X_i^{t+1} = 1, X_U^t = v_U, X_U^{t+1} = v'_U] > 0$, so $\mathbb{P}[X_S^t \subsetneq X_S^{t+1}, X_U^t = v_U, X_U^{t+1} = v'_U] > 0$. Now we need to reverse v and v' for the U indices to account for the other term in the min of Assumption A.3. For each $j \in \hat{S}$, we can take $v_j = v'_j = 0$ while maintaining $\mathbb{P}[X^t = v, X^{t+1} = v'] > 0$ and $\mathbb{P}[X_S^t \subsetneq X_S^{t+1}, X_U^t = v_U, X_U^{t+1} = v'_U] > 0$, since we assumed that all values of $\epsilon_j > 0$. We know that $\mathbb{P}[X^t = v, X^{t+1} = v'] > 0$, so we will try to rearrange $\mathbb{P}[X_U^t = v'_U, X_U^{t+1} = v_U, X_S^t \subsetneq X_S^{t+1}]$ to make use of this fact by attempting to switch v for v' .

We have:

$$\mathbb{P}[X_U^t = v'_U, X_U^{t+1} = v_U, X_S^t \subsetneq X_S^{t+1}] \geq \mathbb{P}[X_U^t = v'_U, X_U^{t+1} = v_U, X_S^t = v_S, X_S^{t+1} = v'_S],$$

because $v_i = 0$, $v'_i = 1$. Let \mathcal{J} be the set of indices in \hat{S} corresponding to noisy indicators of feature i . Because we chose v, v' so that $v_{S \setminus \{i\}} = v'_{S \setminus \{i\}}$, all features in $\hat{S} \setminus \mathcal{J}$ are identically distributed at times t and $t+1$. Therefore, we can switch v and v' for some of the indices in U :

$$\begin{aligned} \mathbb{P}[X_U^t = v'_U, X_U^{t+1} = v_U, X_{S \setminus U}^t = v_{S \setminus U}, X_{S \setminus U}^{t+1} = v'_{S \setminus U}] &= \mathbb{P}[X_{U \cap S}^t = v_{U \cap S}, X_{U \cap S}^{t+1} = v'_U, \\ X_{U \cap \hat{S} \setminus \mathcal{J}}^t &= v_{U \cap \hat{S} \setminus \mathcal{J}}, X_{U \cap \hat{S} \setminus \mathcal{J}}^{t+1} = v'_{U \cap \hat{S} \setminus \mathcal{J}}, \\ X_{U \cap B}^t &= v'_{U \cap B}, X_{U \cap B}^{t+1} = v_{U \cap B}, \\ X_{U \cap \mathcal{J}}^t &= v'_{U \cap \mathcal{J}}, X_{U \cap \mathcal{J}}^{t+1} = v_{U \cap \mathcal{J}}, \\ X_{S \setminus U}^t &= v_{S \setminus U}, X_{S \setminus U}^{t+1} = v'_{S \setminus U}] \end{aligned}$$

By the definition of the features in B , we can also switch v and v' for $U \cap B$. The only remaining terms to switch are $X_{U \cap \mathcal{J}}^t$ and $X_{U \cap \mathcal{J}}^{t+1}$. Because we took $v_j = v'_j = 0$ for all $j \in \mathcal{J}$ without loss of generality, and $v_i = 0$, $v'_i = 1$, we have:

$$\begin{aligned} \mathbb{P}[X_{\mathcal{J}}^t = v'_{\mathcal{J}} | X_i^t = v_i = 0] &= \prod_{j \in \mathcal{J}} (1 - \epsilon_j) = \mathbb{P}[X_{\mathcal{J}}^t = v_{\mathcal{J}} | X_i^t = v_i = 0] \\ \mathbb{P}[X_{\mathcal{J}}^{t+1} = v_{\mathcal{J}} | X_i^{t+1} = v'_i = 1] &= \prod_{j \in \mathcal{J}} \epsilon_j = \mathbb{P}[X_{\mathcal{J}}^{t+1} = v'_{\mathcal{J}} | X_i^{t+1} = v'_i = 1]. \end{aligned}$$

Hence, we can finally combine:

$$\begin{aligned} \mathbb{P}[X_U^t = v'_U, X_U^{t+1} = v_U, X_{S \setminus U}^t = v_{S \setminus U}, X_{S \setminus U}^{t+1} = v'_{S \setminus U}] &= \mathbb{P}[X_{U \cap S}^t = v_{U \cap S}, X_{U \cap S}^{t+1} = v'_U, \\ X_{U \cap \hat{S} \setminus \mathcal{J}}^t &= v_{U \cap \hat{S} \setminus \mathcal{J}}, X_{U \cap \hat{S} \setminus \mathcal{J}}^{t+1} = v'_{U \cap \hat{S} \setminus \mathcal{J}}, \\ X_{U \cap B}^t &= v_{U \cap B}, X_{U \cap B}^{t+1} = v'_{U \cap B}, \\ X_{U \cap \mathcal{J}}^t &= v_{U \cap \mathcal{J}}, X_{U \cap \mathcal{J}}^{t+1} = v'_{U \cap \mathcal{J}}, \\ X_{S \setminus U}^t &= v_{S \setminus U}, X_{S \setminus U}^{t+1} = v'_{S \setminus U}] \end{aligned}$$

So we have shown for these v, v' that:

$$\begin{aligned} \mathbb{P}[X_U^t = v'_U, X_U^{t+1} = v_U, X_{S \setminus U}^t = v_{S \setminus U}, X_{S \setminus U}^{t+1} = v'_{S \setminus U}] &= \mathbb{P}[X_U^t = v_U, X_U^{t+1} = v'_U, \\ X_{S \setminus U}^t &= v_{S \setminus U}, X_{S \setminus U}^{t+1} = v'_{S \setminus U}]. \end{aligned}$$

But we know the RHS is positive because v, v' had $\mathbb{P}[X^t = v, X^{t+1} = v'] > 0$. Therefore,

$$\begin{aligned} \mathbb{P}[X_U^t = v'_U, X_U^{t+1} = v_U, X_S^t \subsetneq X_S^{t+1}] &\geq \mathbb{P}[X_U^t = v'_U, X_U^{t+1} = v_U, X_{S \setminus U}^t = v_{S \setminus U}, X_{S \setminus U}^{t+1} = v_{S \setminus U}] \\ &> 0. \end{aligned}$$

So we have shown that both $\mathbb{P}[X_S^t \subsetneq X_S^{t+1}, X_U^t = v_U, X_U^{t+1} = v'_U]$ and $\mathbb{P}[X_S^t \subsetneq X_S^{t+1}, X_U^t = v'_U, X_U^{t+1} = v_U]$ are strictly positive, which gives Assumption A.3. \square

A.2 Optimal representations and finite-sample guarantee

This first lemma shows that the feature set S is an optimal representation for the OCP objective.

Lemma A.1. *Let*

$$h^*(v, v') = \underset{y \in \{-1, 1\}}{\operatorname{argmax}} \mathbb{P}[Y = y | X^W = v, X^{W'} = v']$$

be the Bayes-optimal classifier for the OCP objective. There exists a classifier $h(v_S, v'_S)$, depending only on the S coordinates of v and v' , achieving the same OCP error as h^ .*

Proof. First we compute the error of the Bayes-optimal classifier. For $A \subset [d]$, $v, v' \in \{0, 1\}^d$, we say $v_A \subsetneq v'_A$ if $\{i \in A : v_i = 1\} \subsetneq \{i \in A : v'_i = 1\}$. Assumption 1 implies that:

$$\mathbb{P}[Y = 1 | X^W, X^{W'}] = \begin{cases} 1 & X_S^W \subsetneq X_S^{W'} \\ 0 & X_S^{W'} \subsetneq X_S^W. \end{cases}$$

That is, given windows $(X^W, X^{W'})$, if the set of S variables active in X^W is properly contained in the set of S variables active in $X^{W'}$, we know the windows must be in the correct order. Likewise, if the active S variables in $X^{W'}$ are properly contained in those of X^W , the windows must be in the wrong order. Assumption 1 also implies that the only possible cases are $X_S^W \subsetneq X_S^{W'}$, $X_S^{W'} \subsetneq X_S^W$, and $X_S^W = X_S^{W'}$. For the last case, we have:

$$\mathbb{P}[Y = 1|X^W, X^{W'}] = \frac{\mathbb{P}[X^W, X^{W'}|Y = 1]}{\mathbb{P}[X^W, X^{W'}|Y = 1] + \mathbb{P}[X^W, X^{W'}|Y = -1]}$$

Assumption 2 implies that the two terms in the denominator are equal when $X_S^W = X_S^{W'}$. This implies that when $X_S^W = X_S^{W'}$, $\mathbb{P}[Y = 1|X^W, X^{W'}] = 1/2$, so the Bayes error is:

$$\mathbb{P}[Y \neq h^*(X^W, X^{W'})] = \frac{1}{2}\mathbb{P}[X_S^W = X_S^{W'}].$$

Hence, we can define:

$$h(v_S, v'_S) = \begin{cases} 1 & v_S \subsetneq v'_S \\ -1 & \text{otherwise,} \end{cases} \quad (4)$$

where we've made the arbitrary choice of -1 when $v_S = v'_S$. Then we have:

$$\mathbb{P}[Y \neq h(X^W, X^{W'})] = \frac{1}{2}\mathbb{P}[X_S^W = X_S^{W'}].$$

So we can achieve the Bayes error using only the coordinates in S . Additionally, note that $h(v_S, v'_S)$ can easily be written as a linear function over $2|S|$ coordinates. If we let $z = (v_{S_1}, \dots, v_{S_{|S|}}, v'_{S_1}, \dots, v'_{S_{|S|}})$ (i.e., the first $|S|$ coordinates represent v_S and the second $|S|$ v'_S), then $h(v_S, v'_S) = h(z) = \text{Sign}(\sum_{i=|S|+1}^{2|S|} z_i - \sum_{i=1}^{|S|} z_i)$, where we break ties with -1 in accordance with (4). This means that we can safely take \mathcal{C} to be the set of linear functions (in fact, we've just shown one linear function will always suffice), so $\text{VC}(\mathcal{C})$ is bounded. \square

Recall that we are searching over hypotheses $h = (c, g)$ with $h(v, v') = h(g(v), g(v'))$. When g is a feature selector (i.e., g picks a certain set of features) and c is unconstrained, the previous lemma immediately implies we can take $g = S$.

Corollary A.1. *The feature selector g that selects the features S is an optimal representation for the OCP objective.*

Now we show that S is the *unique* optimal choice of features. This lemma gives a clean expression for the error incurred by choosing a set of features U potentially different from S .

Lemma A.2. *For $U \subset [d]$, define the error of U as the best OCP loss achievable when using the features in the set U :*

$$\text{err}(U) = \inf_{c \text{ measurable}} \mathbb{P}[Y \neq c(X_U^W, X_U^{W'})].$$

For any $p, q \in \{0, 1\}^{|U|}$ with $p_{U \cap S} = q_{U \cap S}$, define:

$$m_U(p, q) = \min \left(\mathbb{P}[X_S^W \subsetneq X_S^{W'} | X_U^W = p, X_U^{W'} = q]; \mathbb{P}[X_S^{W'} \subsetneq X_S^W | X_U^W = p, X_U^{W'} = q] \right).$$

The expected value of $m_U(p, q)$ is given by:

$$\mathbb{E}_{p,q}[m_U(p, q)] = \sum_{p,q \in \{0,1\}^{|U|}} \mathbb{I}[p_{U \cap S} = q_{U \cap S}] \mathbb{P}[X_U^W = p, X_U^{W'} = q] m_U(p, q)$$

Then for any U ,

$$\begin{aligned} \text{err}(U) &= \frac{1}{2}\mathbb{P}[X_S^W = X_S^{W'}] + \mathbb{E}_{p,q}[m_U(p, q)] \\ &= \text{err}(S) + \mathbb{E}_{p,q}[m_U(p, q)] \end{aligned}$$

Proof. As with S , we know the optimal classifier using the features U is given by:

$$h_U(p, q) = \underset{y}{\operatorname{argmax}} \mathbb{P}[Y = y | X_U^W = p, X_U^{W'} = q].$$

So we have:

$$\begin{aligned} \operatorname{err}(U) &= \mathbb{P}[Y \neq h_U(p, q)] = \sum_{p, q} \mathbb{P}[Y \neq h_U(p, q) | X_U^W = p, X_U^{W'} = q] \mathbb{P}[X_U^W = p, X_U^{W'} = q] \\ &= \sum_{p, q} \mathbb{P}[Y \neq h_U(p, q), X_S^W = X_S^{W'} | X_U^W = p, X_U^{W'} = q] \mathbb{P}[X_U^W = p, X_U^{W'} = q] \\ &\quad + \sum_{p, q} \mathbb{P}[Y \neq h_U(p, q), X_S^W \subsetneq X_S^{W'} | X_U^W = p, X_U^{W'} = q] \mathbb{P}[X_U^W = p, X_U^{W'} = q] \\ &\quad + \sum_{p, q} \mathbb{P}[Y \neq h_U(p, q), X_S^{W'} \subsetneq X_S^W | X_U^W = p, X_U^{W'} = q] \mathbb{P}[X_U^W = p, X_U^{W'} = q], \end{aligned} \quad (5)$$

where we've used Assumption 1 to narrow down to these three cases for the relationship between X_S^W and $X_S^{W'}$. The first term of (5) is the easiest to handle:

$$\begin{aligned} &\sum_{p, q} \mathbb{P}[Y \neq h_U(p, q), X_S^W = X_S^{W'} | X_U^W = p, X_U^{W'} = q] \mathbb{P}[X_U^W = p, X_U^{W'} = q] \\ &= \sum_{p, q} \mathbb{P}[Y \neq h_U(p, q) | X_S^W = X_S^{W'}] \mathbb{P}[X_S^W = X_S^{W'} | X_U^W = p, X_U^{W'} = q] \mathbb{P}[X_U^W = p, X_U^{W'} = q] \\ &= \frac{1}{2} \sum_{p, q} \mathbb{P}[X_S^W = X_S^{W'} | X_U^W = p, X_U^{W'} = q] \mathbb{P}[X_U^W = p, X_U^{W'} = q] \\ &= \frac{1}{2} \mathbb{P}[X_S^W = X_S^{W'}]. \end{aligned}$$

In the first equality, we used that only the values of X_S^W and $X_S^{W'}$ affect the posterior distribution of Y , as we saw in the proof of Lemma A.1. The second equality used that $\mathbb{P}[Y \neq h_U(p, q) | X_S^W = X_S^{W'}] = \frac{1}{2}$ regardless of the value of $h_U(p, q)$, since we showed in the proof of Lemma A.1 that $\mathbb{P}[Y = 1 | X_S^W = X_S^{W'}] = \frac{1}{2}$.

Now we consider the second term of (5):

$$\begin{aligned} &\sum_{p, q} \mathbb{P}[Y \neq h_U(p, q), X_S^W \subsetneq X_S^{W'} | X_U^W = p, X_U^{W'} = q] \mathbb{P}[X_U^W = p, X_U^{W'} = q] \\ &= \sum_{p, q} \mathbb{P}[Y \neq h_U(p, q) | X_S^W \subsetneq X_S^{W'}] \mathbb{P}[X_S^W \subsetneq X_S^{W'} | X_U^W = p, X_U^{W'} = q] \mathbb{P}[X_U^W = p, X_U^{W'} = q] \end{aligned}$$

When $p_{U \cap S} \subsetneq q_{U \cap S}$, it can be easily verified that $h_U(p, q) = 1$ and $Y = 1$. Likewise, when $q_{U \cap S} \subsetneq p_{U \cap S}$, $h_U(p, q) = Y = 0$. So this term is equal to:

$$\sum_{p, q : p_{U \cap S} = q_{U \cap S}} \mathbb{P}[Y \neq h_U(p, q) | X_S^W \subsetneq X_S^{W'}] \mathbb{P}[X_S^W \subsetneq X_S^{W'} | X_U^W = p, X_U^{W'} = q] \mathbb{P}[X_U^W = p, X_U^{W'} = q]$$

Applying the same trick to the third term of (5) yields:

$$\begin{aligned} &\sum_{p, q} \mathbb{P}[Y \neq h_U(p, q), X_S^{W'} \subsetneq X_S^W | X_U^W = p, X_U^{W'} = q] \mathbb{P}[X_U^W = p, X_U^{W'} = q] \\ &= \sum_{p, q : p_{U \cap S} = q_{U \cap S}} \mathbb{P}[Y \neq h_U(p, q) | X_S^{W'} \subsetneq X_S^W] \mathbb{P}[X_S^{W'} \subsetneq X_S^W | X_U^W = p, X_U^{W'} = q] \mathbb{P}[X_U^W = p, X_U^{W'} = q] \end{aligned}$$

So we're left with:

$$\begin{aligned} &\sum_{p, q : p_{U \cap S} = q_{U \cap S}} \mathbb{P}[Y \neq h_U(p, q) | X_S^W \subsetneq X_S^{W'}] \mathbb{P}[X_S^W \subsetneq X_S^{W'} | X_U^W = p, X_U^{W'} = q] \mathbb{P}[X_U^W = p, X_U^{W'} = q] + \\ &\sum_{p, q : p_{U \cap S} = q_{U \cap S}} \mathbb{P}[Y \neq h_U(p, q) | X_S^{W'} \subsetneq X_S^W] \mathbb{P}[X_S^{W'} \subsetneq X_S^W | X_U^W = p, X_U^{W'} = q] \mathbb{P}[X_U^W = p, X_U^{W'} = q] \end{aligned}$$

Now consider a fixed p, q with $p_{U \cap S} = q_{U \cap S}$. If $X_S^W \subsetneq X_S^{W'}$, then $Y = 1$, so we only pay if $h_U(p, q) = -1$. Likewise, if $X_S^{W'} \subsetneq X_S^W$, $Y = -1$, so we only pay if $h_U(p, q) = 1$. Since $h_U(p, q)$ is the Bayes-optimal classifier, it pays the minimal loss for each p, q . So the above sum is equal to:

$$\sum_{p, q : p_{U \cap S} = q_{U \cap S}} \min(\mathbb{P}[X_S^W \subsetneq X_S^{W'} | X_U^W = p, X_U^{W'} = q], \mathbb{P}[X_S^{W'} \subsetneq X_S^W | X_U^W = p, X_U^{W'} = q]) \mathbb{P}[X_U^W = p, X_U^{W'} = q],$$

which is precisely $\mathbb{E}_{p, q}[m_U(p, q)]$. \square

Lemma A.3. For any $U \subset [d]$ with $S \not\subset U$, $\mathbb{E}_{p, q}[m_U(p, q)] > 0$.

Proof. Fix $U \subset [d]$ with $S \not\subset U$. Assumption 3 states that there exists t and p, q with $p_{U \cap S} = q_{U \cap S}$ satisfying:

$$\mathbb{P}[X_S^t \subsetneq X_S^{t+1}, X_U^t = p, X_U^{t+1} = q] > 0 \quad (6)$$

$$\mathbb{P}[X_S^t \subsetneq X_S^{t+1}, X_U^t = q, X_U^{t+1} = p] > 0 \quad (7)$$

Fix t, p, q from this assumption. Observe that:

$$\begin{aligned} \mathbb{P}[X_S^W \subsetneq X_S^{W'} | X_U^W = p, X_U^{W'} = q] &\geq \mathbb{P}[X_S^W \subsetneq X_S^{W'}, W = t, W' = t + 1 | X_U^W = p, X_U^{W'} = q] \\ &= \mathbb{P}[X_S^W \subsetneq X_S^{W'} | W = t, W' = t + 1, X_U^W = p, X_U^{W'} = q] \mathbb{P}[W = t, W' = t + 1 | X_U^W = p, X_U^{W'} = q] \\ &= \mathbb{P}[X_S^t \subsetneq X_S^{t+1} | X_U^t = p, X_U^{t+1} = q] \mathbb{P}[W = t, W' = t + 1 | X_U^W = p, X_U^{W'} = q] \end{aligned} \quad (8)$$

The first term is strictly positive because (6) guaranteed $\mathbb{P}[X_S^t \subsetneq X_S^{t+1}, X_U^t = p, X_U^{t+1} = q] > 0$.

$$\begin{aligned} \mathbb{P}[W = t, W' = t + 1 | X_U^W = p, X_U^{W'} = q] &= \frac{\mathbb{P}[X_U^W = p, X_U^{W'} = q | W = t, W' = t + 1] \mathbb{P}[W = t, W' = t + 1]}{\mathbb{P}[X_U^W = p, X_U^{W'} = q]} \\ &= \frac{\mathbb{P}[X_U^t = p, X_U^{t+1} = q] \mathbb{P}[W = t, W' = t + 1]}{\mathbb{P}[X_U^W = p, X_U^{W'} = q]} \\ &= \frac{\mathbb{P}[X_U^t = p, X_U^{t+1} = q] \mathbb{P}[W = t, W' = t + 1]}{\sum_{w, w'} \mathbb{P}[X_U^w = p, X_U^{w'} = q | W = w, W' = w'] \mathbb{P}[W = w, W' = w']} \end{aligned}$$

Equation (6) guarantees the first term of the numerator is positive. The second term of the numerator is positive by the definition of OCP sampling. The denominator is positive since the numerator is a term in the denominator. Hence, $\mathbb{P}[W = t, W' = t + 1 | X_U^W = p, X_U^{W'} = q] > 0$, so (8) implies $\mathbb{P}[X_S^W \subsetneq X_S^{W'} | X_U^W = p, X_U^{W'} = q] > 0$. A symmetric analysis setting $W' = t$ and $W = t + 1$ and appealing to (7) implies $\mathbb{P}[X_S^{W'} \subsetneq X_S^W | X_U^W = p, X_U^{W'} = q] > 0$.

We've shown that for any U with $S \not\subset U$ there exists p, q with $p_{U \cap S} = q_{U \cap S}$, $m_U(p, q) > 0$, and $\mathbb{P}[X_U^W = p, X_U^{W'} = q] > 0$. This implies that $\mathbb{E}_{p, q}[m(p, q)] > 0$. \square

Combining the previous two lemmas immediately implies that S is the unique optimal set of features for OCP (of size $|S|$).

Corollary A.2. S is the unique optimal set of features of size $|S|$.

Notation. Recall that the OCP pretraining objective R_{ord} is defined over pairs (c, g) of an order classifier c with a representation g . For $h = (c, g)$, we equate $R_{ord}(h) = R_{ord}(g, c)$. For a fixed g , we let $R_{ord}(g)$ be the ordering loss of the best classifier c on top of g : $R_{ord}(g) = \inf_{c \in \mathcal{C}} R_{ord}(c, g)$. Similarly, the downstream loss $R(f, g)$ is a function of a downstream classifier f on top of a representation g . We refer to the empirical versions of R_{ord} and R as \hat{R}_{ord} and \hat{R} , respectively. Recall from the proof of Lemma A.1 that in our case, we can take \mathcal{C} to consist of one (linear) function without loss of generality.

Finally, we can prove Theorem 1.

Theorem (Theorem 1 (formal)). Let $\epsilon_0 = \min_{U: S \not\subset U} \text{err}(U) - \text{err}(S)$. Suppose we have an unlabeled dataset of m i.i.d. pretraining points $\{(Z_i, Y_i)\}_{i=1}^m$, with:

$$m > \frac{2 \left(\log \binom{d}{|S|} + \log \frac{4}{\delta} \right)}{\epsilon_0^2},$$

and a labeled dataset of n downstream training points $\{(X_i^t, Y_i^t)\}_{i=1}^n$. Let \mathcal{G} be all sets of size $|S|$ features chosen from the full set of d features. Let \hat{g} be the minimizer of the empirical OCP pretraining objective:

$$\hat{g} = \operatorname{argmin}_{g \in \mathcal{G}} \hat{R}_{ord}(g)$$

Let \hat{f} be the minimizer of the empirical downstream objective when using the fixed representation \hat{g} :

$$\hat{f} = \operatorname{argmin}_{f \in \mathcal{F}} \hat{R}(f, \hat{g})$$

Then for any $\delta > 0$, with probability at least $1 - \delta$, (\hat{f}, \hat{g}) has excess risk:

$$R(\hat{f}, \hat{g}) - \inf_{(f, g) \in \mathcal{F} \times \mathcal{G}} R(f, g) \leq O\left(\sqrt{\frac{\operatorname{VC}(\mathcal{F}) + \log(\frac{2}{\delta})}{n}}\right)$$

Proof. First, observe that Lemma A.2 implies that $\epsilon_0 > 0$. We identify the functions $g \in \mathcal{G}$ with the sets that they select. Mohri et al. (Theorem 2.13, p.20, 2018) implies that the amount of unlabeled data is sufficient so that $R_{ord}(\hat{g}) < R_{ord}(S) + \epsilon_0$ with probability at least $1 - \delta/2$. By Lemma A.2, this implies $\hat{g} = S$. Then we can learn downstream without incurring any cost compared to an optimal pair (f^*, g^*) , since $\hat{g} = g^* = S$. Hence, the standard uniform convergence argument for VC-classes (e.g., Shalev-Shwartz and Ben-David (2014, Theorem 6.8)) implies the excess risk bound for (\hat{f}, \hat{g}) holds with probability at least $1 - \delta/2$. Combining the probabilities of failure for each step gives total failure probability at most δ . □

A.3 PCL versus OCP

Consider a feature $X_i^t \in \{0, 1\}$ such that $\mathbb{P}[X_i^1 = 1] = 1/2$, and $X_i^{t+1} = 1 - X_i^t$. This is a periodic feature with period 1. X_i^t qualifies as a background feature, since it is independent of all other features X_j , and $P[X_i^t = v, X_i^{t+1} = v'] = P[X_i^t = v', X_i^{t+1} = v]$ for all $(v, v') \in \{0, 1\}^2$. By the results from the previous section, inclusion of this feature therefore does not affect whether OCP finds the correct representation S . In OCP with consecutive windows, $X_i^W \neq X_i^{W'}$ with probability 1 regardless of Y , so that feature is not useful for distinguishing between correctly and incorrectly-ordered pairs. However, this feature is very helpful for PCL. Because the positives for PCL are always consecutive, but the negatives are pairs of random elements, $\mathbb{P}[Y = -1 | X_i^W = X_i^{W'}] = 1$. This biases PCL towards inclusion of feature i over inclusion of more rarely-occurring features in S , so there are simple examples where PCL provably never recovers S even in the infinite-sample limit. Additionally, even when PCL works well in the infinite data limit, features that are *weakly* predictive of whether elements are consecutive or not can affect the unlabeled sample complexity of PCL, as we show empirically in Appendix B.

A.4 Biased OCP

The OCP algorithm samples windows $(W, W + 1)$ when $Y = 1$ and $(W + 1, W)$ when $Y = -1$. Consider instead a version of OCP with “bias” in the negative distribution. When $Y = 1$, we still sample $(W, W + 1)$ as the window pair. But when $Y = -1$, we sample a *random* pair (W, W') with $|W - W'| = 1$. As with PCL, half of the negatives for this “OCP-biased” sampling are actually in the correct order. Does this affect the representation learned during pretraining, and does it affect the unlabeled sample complexity required for learning?

Now we give two lemmas analogous to lemmas A.1 and A.2.

Lemma A.4. *Let*

$$h^*(v, v') = \operatorname{argmax}_{y \in \{-1, 1\}} \mathbb{P}[Y = y | X^W = v, X^{W'} = v']$$

be the Bayes-optimal classifier for the OCP-biased objective. There exists a classifier $h(v_S, v'_S)$, depending only on the S coordinates of v and v' , achieving the same OCP-biased error as h^ .*

Lemma A.5. For $U \subset [d]$, define the error of U as the best OCP-biased loss achievable when using the features in the set U :

$$\text{err}(U) = \inf_{c \text{ measurable}} \mathbb{P}[Y \neq c(X_U^W, X_U^{W'})].$$

For any $p, q \in \{0, 1\}^{|U|}$ with $p_{U \cap S} = q_{U \cap S}$, define:

$$m'_U(p, q) = \min \left(\frac{1}{3} \mathbb{P}[X_S^W \subsetneq X_S^{W'} | X_U^W = p, X_U^{W'} = q]; \mathbb{P}[X_S^{W'} \subsetneq X_S^W | X_U^W = p, X_U^{W'} = q] \right).$$

The expected value of $m'_U(p, q)$ is given by:

$$\mathbb{E}_{p, q}[m'_U(p, q)] = \sum_{p, q \in \{0, 1\}^{|U|}} \mathbb{I}[p_{U \cap S} = q_{U \cap S}] \mathbb{P}[X_U^W = p, X_U^{W'} = q] m'_U(p, q)$$

Then for any U ,

$$\begin{aligned} \text{err}(U) &= \frac{1}{2} \mathbb{P}[X_S^W = X_S^{W'}] + \frac{1}{3} \mathbb{P}[X_S^W \subsetneq X_S^{W'}] + \mathbb{E}_{p, q}[m'_U(p, q)] \\ &= \text{err}(S) + \mathbb{E}_{p, q}[m'_U(p, q)] \end{aligned}$$

The proofs of these lemmas are entirely analogous to the proofs of lemmas A.1, A.2 with some extra handling of the case of the $X_S^W \subsetneq X_S^{W'}$, since in OCP-biased those samples could also be negatives.

Now we can prove Theorem 2:

Theorem (Theorem 2 (formal)). *Suppose Assumptions 1-3 are satisfied. Then S is the unique optimal representation for the OCP-biased objective of size $|S|$.*

Proof. By direct comparison of the definition of $m_U(p, q)$ in Lemma A.2 and $m'_U(p, q)$ in A.5, we see that $\mathbb{E}_{p, q}[m_U(p, q)] > 0 \implies \mathbb{E}_{p, q}[m'_U(p, q)] > 0$. Therefore, Lemma A.3 shows that Assumption 3 implies $\mathbb{E}_{p, q}[m'_U(p, q)] > 0$ for all U with $S \not\subset U$. Hence, S is the unique optimal representation for OCP-biased. \square

Since S is the unique optimal representation for OCP-biased, an entirely similar analysis to the proof of Theorem 1 yields a finite-sample bound for OCP-biased. However, recall that the bound on the *unlabeled* sample complexity depends on $1/\epsilon_0^2$, where $\epsilon_0 = \min_{U: S \not\subset U} \text{err}(U) - \text{err}(S)$. By comparing the error formulae in Lemmas A.2 and A.5 (in particular, the difference between m_U and m'_U), we see that for any fixed U with $S \not\subset U$, $\text{err}_{\text{biased}}(U) - \text{err}_{\text{biased}}(S) < \text{err}_{\text{ocp}}(U) - \text{err}_{\text{ocp}}(S)$. Therefore, for the same distribution over X , the ϵ_0 for OCP is larger than the ϵ_0 for OCP-biased. This results in a better upper bound on the unlabeled sample complexity required to find a good representation. In the following section, we show empirically that OCP-biased indeed requires more samples to find a good representation, but ultimately finds the same representation as OCP.

B Synthetic Experiments

B.1 Synthetic Distributions

In this section we describe the distributions for the synthetic experiments in Section 4.1. Both of these distributions are members of the class of distributions from Section 3.1. We proved in Appendix A that all of these distributions satisfy Assumptions 1-3, and that S is the optimal representation for OCP on these distributions.

Distribution 1 (Figure 2a)

Our first synthetic distribution includes trajectories X of length 10 with $d = 8$ features that are generated as described below:

- The set S consists of 4 time-irreversible features. Each feature in S has a fixed probability of activating (switching from ‘0’ to ‘1’) over the entire trajectory, independent of the other features in S . Activation time was chosen uniformly over the whole trajectory and independently per feature, and once activated, features remained on. In our synthetic data, probabilities were 0.4 for the first two features and 0.6 for the next two features.
- The set \hat{S} consisted of noisy versions of the first three variables in S . ϵ_i was set to 0.7 for all variables.
- The last feature X_{-1} was a background, reversible feature in B . It was a periodic function, alternating between 0 and 1, with uniform initialization over $\{0, 1\}$.

Distribution 2 (Figure 2b)

Similarly, our second synthetic distribution includes trajectories X of length 10 with $d = 7$ features that are generated as described below:

- The set S consists of 4 time-irreversible features, identical to before. Each feature in S has a fixed probability of activating (switching from ‘0’ to ‘1’) over the entire trajectory, independent of the other features in S . Activation time was chosen uniformly over the whole trajectory and independently per feature, and once activated, features remained on. In our synthetic data, probabilities were 0.4 for the first two features and 0.6 for the next two features.
- The set \hat{S} consisted of noisy versions of the first two variables in S . ϵ_i was set to 0.55 for all variables.
- The last feature X_{-1} was a background, reversible feature in B . X_{-1}^0 was sampled uniformly from $\{0, 1\}$ and the rest of X_{-1} was set such that for all $t \geq 1$, $\mathbb{P}[X_{-1}^t = X_{-1}^{t-1}] = 0.3$.

B.2 Experimental Setup

Synthetic Data Creation

In order to quantify how many unlabeled pre-training samples are required to recover S , we created data sets $\{X_i\}_{i=1}^m$ with varying m , taking on values 50, 100, 200, 400, 600, 800, 1000, 2000, 4000, 8000, and 16,000. For each dataset size m , we generated 100 independently drawn sets, according to the distributions described previously.

For each created set, a single pair was sampled from each trajectory according to OCP, PCL, or OCP-biased sampling. I.e. an unlabeled dataset of m trajectories produced a pre-training dataset with m pairs.

Selection of Optimal Representation

Now, for each of our pre-training datasets and each sampling scheme (OCP, PCL, OCP-biased), we now determine what feature representation \hat{g} would be selected via pre-training for each one.

In order to do so, we iterate over each possible $g \in \mathcal{G}$ where $\mathcal{G} = \{U \subset [d] : |U| = |S| = d_0\}$. After subselecting to the features in g , each pair of data points $(X^t, X^{t'})$ was featurized as input as $[X^t; X^{t'}; X^t - X^{t'}; |X^t - X^{t'}|]$. The

loss is then minimized via the LogisticRegression implementation from scikit-learn with the ‘liblinear’ optimizer. We then select the \hat{g} that minimizes the empirical pre-training loss and calculate the overlap with the true features in S .

B.3 Explanation of Observed Behavior

Distribution 1

In Distribution 1, we see that OCP essentially always converges to the optimal representation with 8,000 data points, and OCP-biased also converges, albeit slower. However, PCL is never able to break the barrier of 3 features, since it opts to choose the periodic feature X_{-1} instead. This feature (which alternates between 0 and 1) is highly discriminative for the PCL pre-training task, since $|X_{-1}^t - X_{-1}^{t+1}| = 1$ whenever $Y = 1$, and the same is true only half the time when $Y = 0$. Therefore, while it is a background, reversible feature that may not be useful for a downstream task, it is useful for the PCL task, and therefore PCL fails to find the optimal representation of time-irreversible features. In contrast, for OCP, $|X_{-1}^t - X_{-1}^{t+1}| = 1$ is true across all examples, and therefore would not be chosen as a discriminative feature.

Distribution 2

In Distribution 2, we again see that all methods are able to identify the optimal representation. However, OCP requires fewer pre-training samples in order to reach the optimal representation. While not periodic as in Distribution 1, X_{-1} is again weakly predictive of whether or not a window is consecutive. While X_{-1} is not as strongly predictive as in Distribution 1, PCL still opts to select it in \hat{g} in the lower-data regime. However, with sufficient data, PCL overcomes the “false” signal to opt for the “correct” feature instead.

C Real-world Experiments

C.1 Pre-training for feature selection

Implementation Details

For the linear pre-training for feature selection experiment, we utilized the LogisticRegression implementation from scikit-learn with the ‘liblinear’ optimizer and balanced class reweighting (Pedregosa et al., 2011). Hyperparameter tuning was conducted on the validation set independently for each feature subset, dataset size, and fold number. Hyperparameters included regularization scheme (‘l1’ or ‘l2’) and regularization constant (10^{-3} to 10^6). All performance reported is on the held-out test sets using the best hyperparameter setting from the validation set.

Granular Experimental Results

Below, we present a more granular view on results including standard deviations, focusing on the difference between direct downstream prediction and using order contrastive pre-training to select features.

Table 2: AUC \pm standard deviations (as calculated over five folds), comparing the use of all features to just those selected by OCP. The final row displays the difference in AUC between OCP-selected features and all features and the standard deviation of that difference.

| Features | Fraction of training data | | | | |
|---------------------|---------------------------|------------------|------------------|------------------|------------------|
| | 1 | 1/2 | 1/4 | 1/8 | 1/16 |
| OCP subset | 0.864 \pm .022 | 0.860 \pm .026 | 0.847 \pm .024 | 0.808 \pm .023 | 0.786 \pm .068 |
| All features | 0.856 \pm .021 | 0.851 \pm .021 | 0.818 \pm .036 | 0.723 \pm .057 | 0.726 \pm .074 |
| OCP - All | 0.008 \pm .014 | 0.008 \pm .039 | 0.029 \pm .035 | 0.082 \pm .058 | 0.054 \pm .050 |

Average Precision Results

In addition to the note-level AUC, we also provide results via a different patient-level precision metric; this provides another interpretable view on performance, since each patient has multiple notes in the test set, and the labels are imbalanced. We define *Average Precision at 80% recall* in the following manner. We first find the threshold at which 80% of positive labels (displays progression) would be recovered. Then per patient, we calculate the precision of the retrieved notes, assuming use of that threshold, which we then average over patients. If no notes are surfaced for a patient, we set the precision to 1 if no positive note exists, and 0 otherwise.

We then assess performance using this precision-level metric. We follow the same procedure as before, except that hyperparameter settings are now chosen on the basis of this precision metric on the validation set, instead of AUC. In the table below, we note the same trends are present with average precision as with AUC; namely, there is essentially no difference when all training data can be used, but a much larger difference when the model is restricted to only a fraction.

Table 3: Average precision at 80% recall \pm standard deviations (as calculated over five folds), comparing the use of all features to just those selected by OCP. The final row displays the difference in average precision between OCP-selected features and all features.

| Available features | Fraction of training data | | | | |
|---------------------|---------------------------|-----------------|-----------------|-----------------|-----------------|
| | 1 | 1/2 | 1/4 | 1/8 | 1/16 |
| OCP subset | 0.54 \pm 0.09 | 0.54 \pm 0.07 | 0.51 \pm 0.08 | 0.48 \pm 0.09 | 0.42 \pm 0.10 |
| All features | 0.54 \pm 0.12 | 0.50 \pm 0.12 | 0.49 \pm 0.05 | 0.37 \pm 0.07 | 0.37 \pm 0.07 |
| OCP - All | 0 \pm 0.08 | 0.05 \pm 0.10 | 0.02 \pm 0.08 | 0.10 \pm 0.07 | 0.05 \pm 0.05 |

C.2 Nonlinear representations

Implementation Details

The masked language modeling was conducted using the BertForMaskedLM implementation from Wolf et al. (2020) with a 15% masking rate and a learning rate of 5e-5. The checkpoint used downstream was selected as the one in which the model had the lowest validation loss on a held-out set of notes. The contrastive pre-training across all 3 objectives was conducted using the BertForNextSentencePrediction implementation from Wolf et al. (2020) with a learning rate of 1e-5 and weight decay of 0.01. For each contrastive approach, the model checkpoint with the highest validation accuracy on a held-out set of 3200 contrastive pairs was chosen for use downstream. The L2-regularized linear layer was implemented using scikit-learn, using class-balanced reweighting and the ‘liblinear’ optimizer. As before, the regularization (1e-2 to 1e5) was chosen for each seed, fold, model, and dataset size using the best performance on the validation set.

Computational Burden

Each pre-training method was trained on a single NVIDIA GeForce GTX 1080 Ti GPU with 12GB of memory. Convergence required about 6 hours for the masked language modeling pre-training and 3-4 hours for each of the contrastive pre-training methods. Since pre-training was conducted over 3 different seeds, it collectively required about 54 GPU hours. The forward pass to extract frozen embeddings from the different models across seeds and folds was minimal in time (less than 15 minutes). All downstream experiments and hyperparameter tuning involved solely a linear layer and were conducted on a CPU.

Granular Experimental Results

Below, we display more granular experimental results, splitting performance by seed. We display both the mean and standard deviation AUC across the 5 folds at each training data size. While there is some variation between seeds, we find that OCP consistently outperforms at smaller dataset sizes.

Table 4: AUC average and standard deviation (as calculated over five folds) over each of the three seeds trained per method.

| <i>AUC ± std dev</i> | <i>Fraction of training data</i> | | | | |
|------------------------------|----------------------------------|------------|------------|------------|-------------|
| | 1 | 1/2 | 1/4 | 1/8 | 1/16 |
| OCP Seed A | 0.87 ± .04 | 0.86 ± .04 | 0.84 ± .03 | 0.82 ± .03 | 0.83 ± .04 |
| OCP Seed B | 0.86 ± .03 | 0.85 ± .04 | 0.84 ± .04 | 0.81 ± .02 | 0.82 ± .03 |
| OCP Seed C | 0.88 ± .02 | 0.86 ± .04 | 0.85 ± .03 | 0.81 ± .03 | 0.79 ± .02 |
| BERT | 0.79 ± .06 | 0.74 ± .06 | 0.72 ± .04 | 0.63 ± .05 | 0.60 ± .06 |
| Fine-Tuned LM Seed A | 0.87 ± .05 | 0.85 ± .06 | 0.85 ± .02 | 0.77 ± .06 | 0.75 ± .06 |
| Fine-Tuned LM Seed B | 0.84 ± .03 | 0.82 ± .04 | 0.79 ± .02 | 0.70 ± .07 | 0.67 ± .09 |
| Fine-Tuned LM Seed C | 0.83 ± .02 | 0.80 ± .01 | 0.77 ± .01 | 0.73 ± .06 | 0.71 ± .05 |
| Pt-Contrastive Seed A | 0.82 ± .04 | 0.79 ± .03 | 0.77 ± .04 | 0.69 ± .07 | 0.66 ± .10 |
| Pt-Contrastive Seed B | 0.85 ± .02 | 0.83 ± .04 | 0.78 ± .04 | 0.72 ± .03 | 0.69 ± .1 |
| Pt-Contrastive Seed C | 0.84 ± .05 | 0.85 ± .05 | 0.83 ± .03 | 0.77 ± .04 | 0.72 ± .06 |
| PCL Seed A | 0.88 ± .02 | 0.88 ± .03 | 0.85 ± .02 | 0.83 ± .04 | 0.78 ± .03 |
| PCL Seed B | 0.86 ± .01 | 0.85 ± .04 | 0.80 ± .06 | 0.77 ± .02 | 0.76 ± .04 |
| PCL Seed C | 0.86 ± .03 | 0.84 ± .03 | 0.80 ± .04 | 0.75 ± .04 | 0.71 ± .08 |

Average Precision Results

Below we show results on the test set using the patient-level average precision metric introduced in Appendix B.1. As before, hyperparameter settings are chosen based on the model with best performance on the validation set, per this precision metric. Again, we find that OCP is relatively consistent in performance even with minimal training data, and outperforms other methods, particularly in the low data regime.

Table 5: Performance in terms of average precision at 80% recall as calculated across 5 folds and 3 seeds for fine-tuned methods. The first row contains the mean average precision of OCP \pm its standard deviation. The following rows contain the mean precision advantage of OCP over each comparison method, and the percentage of time OCP outperforms that method, computed over all seeds and folds.

| <i>Prec diff (OCP Win %)</i> | <i>Fraction of training data</i> | | | | |
|------------------------------|----------------------------------|----------------|----------------|----------------|----------------|
| | 1 | 1/2 | 1/4 | 1/8 | 1/16 |
| OCP Prec | 0.57 \pm .08 | 0.58 \pm .07 | 0.55 \pm .06 | 0.50 \pm .08 | 0.50 \pm .09 |
| OCP - BERT | 0.16 (100%) | 0.19 (100%) | 0.19 (100%) | 0.19 (100%) | 0.22 (100%) |
| OCP - FT LM | 0.09 (80%) | 0.09 (80%) | 0.09 (84%) | 0.11 (87%) | 0.13 (93%) |
| OCP - Pt-Contrastive | 0.08 (73%) | 0.09 (80%) | 0.14 (100%) | 0.15 (96%) | 0.17 (98%) |
| OCP - PCL | 0.02 (56%) | 0.04 (69%) | 0.07 (73%) | 0.10 (84%) | 0.13 (91%) |

Qualitative Analysis

We now conduct a qualitative analysis to understand whether the OCP pre-trained model is operating as expected, namely whether it is attending to those features we would expect to be most crucial both in ordering and for downstream progression extraction.

The BertForNextSentencePrediction model used for contrastive pre-training is implemented such that sequence classification is based off of the representation of the CLS classifier token. Therefore, we examine those tokens which are most highly attended to by the CLS token in the last BERT layer, as a proxy signal for what the model is attending to for its final classification representation. Over all examples in the validation set, we find the average attention each token contributed in the final layer. There are then 21 tokens with an average attention above a threshold of 0.1; they include *increased/increase/increasing*, *change*, *unchanged*, *no*, *stable*, *negative*, *persistent*, and *resolved*. All such features elucidate disease stage, and therefore, it seems qualitatively plausible that OCP learns useful downstream representations in the nonlinear case. Other highly ranked features include the SEP token used to split the contrastive pairs, as well as *prior*, which is a possible leaky feature that indicates order information, but is not as useful for downstream analysis.