

# Evaluation of Active Feature Acquisition Methods for Time-varying Feature Settings

Henrik von Kleist<sup>1,2,3</sup>

HENRIK.VONKLEIST@HELMHOLTZ-MUNICH.DE

Alireza Zamanian<sup>2,4</sup>

ALIREZA.ZAMANIAN@IKS.FRAUNHOFER.DE

Ilya Shpitser<sup>3</sup>

ISHPITS1@JHU.EDU

Narges Ahmidi<sup>1,3,4</sup>

NARGES.AHMIDI@HELMHOLTZ-MUNICH.DE

<sup>1</sup>*Institute of AI for Health, Helmholtz Munich - German Research Center for Environmental Health, 85764 Neuherberg, Germany*

<sup>2</sup>*TUM School of Computation, Information and Technology, Technical University of Munich, 85748 Garching, Germany*

<sup>3</sup>*Department of Computer Science, Johns Hopkins University, Baltimore, MD 21218, USA*

<sup>4</sup>*Fraunhofer Institute for Cognitive Systems IKS, 80686 Munich, Germany*

**Editor:** Jin Tian

## Abstract

Machine learning methods often assume that input features are available at no cost. However, in domains like healthcare, where acquiring features could be expensive or harmful, it is necessary to balance a feature’s acquisition cost against its predictive value. The task of training an AI agent to decide which features to acquire is called active feature acquisition (AFA). By deploying an AFA agent, we effectively alter the acquisition strategy and trigger a distribution shift. To safely deploy AFA agents under this distribution shift, we present the problem of active feature acquisition performance evaluation (AFAPE). We examine AFAPE under i) a no direct effect (NDE) assumption, stating that acquisitions do not affect the underlying feature values; and ii) a no unobserved confounding (NUC) assumption, stating that retrospective feature acquisition decisions were only based on observed features. We show that one can apply missing data methods under the NDE assumption and offline reinforcement learning under the NUC assumption. When NUC and NDE hold, we propose a novel semi-offline reinforcement learning framework. This framework requires a weaker positivity assumption and introduces three new estimators: A direct method (DM), an inverse probability weighting (IPW), and a double reinforcement learning (DRL) estimator.

**Keywords:** active feature acquisition, semi-offline reinforcement learning, dynamic testing regimes, missing data, causal inference, semiparametric theory

## 1. Introduction

Machine learning methods typically assume that the full set of input features will be readily available after deployment, with little to no cost. This is, however, not always the case, as acquiring features may impose a significant cost. In such situations, the predictive value of a feature should be balanced against its acquisition cost. In the medical diagnostics context, the cost of feature acquisition (e.g., for a biopsy test) may include not only monetary cost but also the potential adverse harm to patients. This is why physicians acquire certain features,

e.g., via biopsies, MRI scans, or lab tests, only when their diagnostic values outweigh their costs or risks. The challenge is exacerbated when prediction must be made regarding a large number of diverse outcomes with different sets of informative features. Going back to the medical example, a typical emergency department is able to diagnose thousands of different diseases based on a large set of possible observations. For every new emergency patient with ambiguous symptoms, clinicians must narrow down their search for a proper diagnosis via step-by-step feature acquisitions.

Active feature acquisition (AFA) addresses this problem by designing two AI systems: i) a so-called *AFA agent*, deciding which features must be observed while balancing information gain vs. feature acquisition cost; ii) an ML prediction model, often a classifier, that solves the prediction task based on the acquired set of features. To elucidate the AFA process, we present a hypothetical and simplified scenario of diagnosing heart attacks.

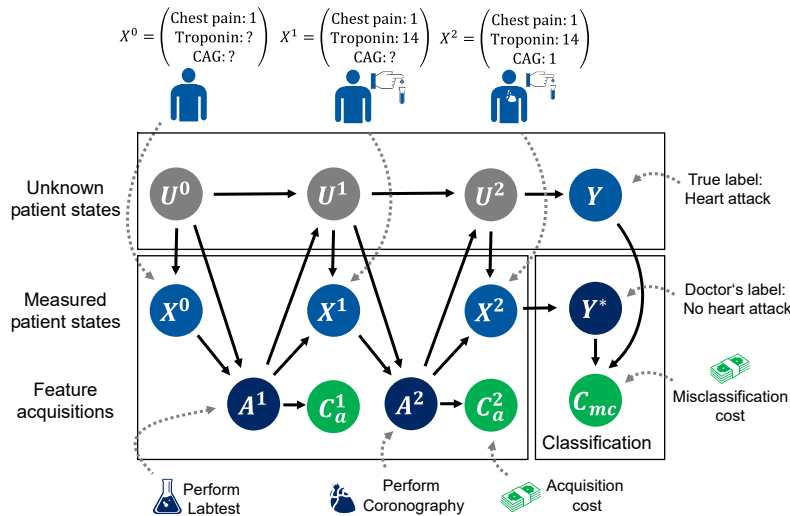


Figure 1: AFA process for a simplified hypothetical heart attack diagnosis example. A patient with chest pain ( $X^0$ ) prompts the doctor to first order a troponin lab test ( $A^1$ ) and, upon reviewing the result ( $X^1$ ), to also order a coronagraphy (CAG) ( $A^2$ ). The feature acquisitions  $A^1$  and  $A^2$  produce feature acquisition costs  $C_a^1$  and  $C_a^2$ . After the acquisition process concludes, the doctor makes a diagnosis  $Y^*$ , which, if different from the true underlying condition  $Y$ , produces a misclassification cost  $C_{mc}$ .

### 1.1 Heart Attack Diagnosis Example

Figure 1 presents the partially observable decision process that encapsulates the sequential decision-making aspect of the AFA problem for a heart attack diagnosis example. Upon arrival at the hospital, a patient with an unknown health state ( $U^0$ ) exhibits the symptom of chest pain ( $X^0 = \text{”chest pain”}$ ). At this stage, no additional information is available. The attending doctor decides to order a troponin lab test ( $A^1 = \text{”acquire troponin”}$ ) as

part of the feature acquisition process. The laboratory test incurs a feature acquisition cost ( $C_a^1 = \text{"\$100"}$ ). Subsequently, upon reviewing the results of the lab test ( $X^2$ ), the doctor decides that a coronography ( $A^2$ ), an invasive imaging procedure, is necessary. Notably, the feature acquisition cost ( $C_a^2$ ) for this procedure may be substantially higher due to the potential harm to the patient. After the completion of the feature acquisition process, a diagnosis of whether the patient is experiencing a heart attack is performed. A (hypothetical) misclassification cost  $C_{mc}$  arises if the diagnosis  $Y^*$  and the true condition  $Y$  differ.

In general, medical tests may also impact the patient’s health (illustrated by the edges  $A^t \rightarrow U^t$ ). The invasive coronography from our example may, for example, cause bleeding, infections, hypotension or other problems (Tavakol et al., 2012). Such an effect is denoted as a *direct effect*, and we refer to its absence as the *no direct effect (NDE)* assumption.

Furthermore, the decision to perform a clinical test  $A^t$  may not solely rely on past observed variables  $X^\tau$  ( $\tau < t$ ) but can also depend on past unobserved variables  $U^\tau$  (illustrated by the edges  $U^{t-1} \rightarrow A^t$ ) or even other factors. For example, a doctor might base acquisition decisions on information that is partially not recorded in the medical database. We refer to the assumption that acquisitions are only determined by past observed variables as the *no unobserved confounding (NUC)* assumption.

## 1.2 Paper Goal

We investigate the evaluation of AFA agents under the distribution shift that occurs since the AFA agent makes different acquisition decisions than the doctors who were responsible for collecting the retrospective data set. The focus of the paper is thus not to design new AFA agents and classifiers but to estimate the performance of *any* AFA agent and classifier at deployment. This means the doctor should be informed, for example, how many wrong diagnoses are to be expected or how much acquisition cost will be incurred on average if an AFA system is deployed.

The paper has two primary objectives: i) Identification, which involves determining the assumptions that enable the unbiased representation of costs from the retrospective data distribution and ii) Estimation, which focuses on turning the obtained identification strategy into point and interval estimators following well-grounded statistical principles. We specifically analyze scenarios that involve both adherence to and violation of the NDE and NUC assumptions. We formulate this problem of *active feature acquisition performance evaluation (AFAPE)* as the problem of estimating the expected counterfactual acquisition and misclassification costs using retrospective data.

## 1.3 Paper Outline and Contributions

The remainder of this paper is organized as follows. After reviewing the necessary background and related methods in Section 2, we formulate the AFAPE problem in Section 3. The general AFAPE problem is not identified—that is, it is not possible to estimate the counterfactual acquisition and misclassification costs from retrospective data when both the NDE and NUC assumptions are violated.

Therefore, we begin Section 4 by employing the NUC assumption and show that this makes AFAPE amenable to an offline reinforcement learning (RL) / dynamic treatment

regimes (DTR) view. This allows the application of known identification and estimation theory from the offline RL / DTR literature.

In Section 5, we make instead the NDE assumption and assume the NUC assumption can be violated. We demonstrate that under the NDE assumption, the AFA decision process depicted in Figure 1 transforms into a missing data graph (m-graph) (Mohan et al., 2013; Shpitser et al., 2015), a recognized graphical framework in the missing data literature. This enables us to apply established identification and estimation theory from the missing data literature. After solving the missing data problem, the AFAPE problem is transformed into an online RL setting where one can simulate different acquisition trajectories, leading to a trivial solution for AFAPE.

In Section 6, we assume both the NUC and the NDE assumptions hold. In this setting, one can apply either offline RL or missing data methods to solve AFAPE, but both require strong positivity assumptions and do not utilize the data optimally. Therefore, we propose a new viewpoint on AFA, which we denote as *semi-offline reinforcement learning*. Under the semi-offline RL viewpoint, the AFA agent engages with the environment in an online manner, but certain actions (where the underlying feature values are missing in the retrospective data) cannot be explored. The positivity assumption required for identification is drastically reduced under the new semi-offline RL viewpoint. We derive three novel estimators that can be denoted as semi-offline RL versions of known offline RL estimators, including the Q-function based direct method (DM) (Levine et al., 2020), inverse probability weighting (IPW) (Levine et al., 2020), and the double reinforcement learning (DRL) estimator (Kallus and Uehara, 2020). Notably, our DRL estimator is doubly robust, exhibiting consistency even if either the underlying Q-function or the propensity score model is misspecified.

In Section 7, we explore the estimation of AFAPE under the NUC and NDE assumptions using semiparametric theory. We demonstrate how all three viewpoints—the offline RL view (under the NDE assumption), the missing data view (under the NUC assumption), and the semi-offline RL view—interconnect within this theoretical framework. Unfortunately, there is no closed-form solution for an efficient estimator. However, we can enhance the efficiency of all estimators by applying established semiparametric techniques from related fields, such as standard missing data problems (Tsiatis, 2006) and dynamic testing and treatment regimes (Liu et al., 2021). These methods, though, come with significant computational costs, are challenging to implement, and require strong positivity assumptions.

In Section 8, we present synthetic data experiments that exemplify the improved data efficiency and reduced positivity requirements of the semi-offline RL estimators. Our experiments also show that biased evaluation methods commonly used in the AFA literature can lead to detrimental conclusions regarding the performance of AFA agents. Deploying such methods without caution may pose significant risks to patients’ lives. We end the paper with a Discussion (Section 9) and Conclusion (Section 10).

## 2. Background and Related Methods

In the following, we review some of the literature about AFA and provide some background on offline RL/ DTR, missing data, and semi-parametric theory.

## 2.1 Active Feature Acquisition (AFA)

Research on active feature acquisition (AFA) and related problem formulations have been published under various different names and in different, largely disjoint, research communities. Early research in economics and decision science literature addressed the problem of "Value of Information" (VoI) (LaValle, 1968a,b; Gould, 1974; Hilton, 1979; Hess, 1982; Keisler et al., 2014). Similar methods have also been applied in the medical field, often in terms of cost-effectiveness analysis of screening policies (Mushlin and Fintor, 1992; Krahn et al., 1994; Botteman et al., 2003; Force\*, 2009). AFA has further been studied under the name of "dynamic testing regimes" (Liu et al., 2021; Robins et al., 2008) or "dynamic monitoring regimes" (Neugebauer et al., 2017; Kreif et al., 2021) in the causal inference literature, often in combination with dynamic treatment regimes. In these settings, the goal of the feature acquisitions is not to enable better predictions/diagnoses but to enable better treatment decisions.

The name "active feature acquisition" (AFA) (An et al., 2006; Li and Oliva, 2021a; Li et al., 2021; Chang et al., 2019; Shim et al., 2018; Yin et al., 2020) is common in the machine learning literature, but other names are also frequently used. These include, but are not limited to, "active sensing" (Yoon et al., 2019, 2018; Tang et al., 2020; Jarrett and van der Schaar, 2020), "active feature elicitation" (Natarajan et al., 2018; Das et al., 2021), "dynamic feature acquisition" (Li and Oliva, 2021b), "dynamic active feature selection" (Zhang, 2019), "element-wise efficient information acquisition" (Gong et al., 2019), "classification with costly features" (Janisch et al., 2020) and "test-cost sensitive classification" (Xiaoyong Chai et al., 2004).

These diverse research fields share a common characteristic, which involves designing an agent to selectively acquire a subset of features to balance acquisition cost and information gain. The approaches used to design such agents range from simple greedy acquisition strategies to more complex RL-based strategies. However, the focus of this work is not on any specific AFA method but rather on evaluating the performance of *any* AFA method under the acquisition distribution shift. For a more comprehensive literature review of existing AFA methods and a distinction between AFA and other related fields, we direct interested readers to Appendix A.

## 2.2 (Offline) Reinforcement Learning (RL) / Dynamic Treatment Regimes (DTR)

We show in Section 4 that AFA can be analyzed from an offline RL/ DTR viewpoint. In Section 5, we show that AFAPE can also be analyzed from an online RL viewpoint (if NDE holds and after missingness has been resolved). Online RL allows the interaction of an agent with the environment and thus the simulation of outcomes under any desired policy, thereby leading to a trivial solution for the AFAPE problem. In offline RL, however, such a simulation is not possible due to missing knowledge about the environment. The AFAPE problem then becomes equivalent to the problem of off-policy policy evaluation (OPE) (Dudik et al., 2011; Thomas and Brunskill, 2016; Kallus and Uehara, 2020), in which the goal is to evaluate the performance of a "target" policy (here the AFA policy) from data collected under a different "behavior" policy (here the retrospective acquisition policy of, for example, a doctor). Several estimators have been developed for the OPE problem. These

include the plug-in based on the G-formula (Robins, 1986) (also referred to as model-based evaluation (Levine et al., 2020)), inverse probability weighting (IPW) (Levine et al., 2020) (also known as importance sampling or the Horvitz-Thompson estimator (Horvitz and Thompson, 1952)), the direct method (DM) (Levine et al., 2020), and double reinforcement learning (DRL) (Kallus and Uehara, 2020).

### 2.3 Missing Data

In this paper, we show that AFAPE can be viewed as a missing data (+ a trivial online RL) problem. Thus, known identification and estimation techniques from the missing data literature can be employed. We show that the NUC assumption described in this paper corresponds under NDE to a *missing-at-random (MAR)* assumption. Violations of the NUC assumption correspond, in our setting, to a special, identified *missing-not-at-random (MNAR)* scenario. Estimation strategies generally include inverse probability weighting (IPW) (Seaman and White, 2013), and multiple imputation (MI) (Sterne et al., 2009) (a special case of the plug-in of the G-formula).

### 2.4 Semiparametric Theory

The goal of AFAPE is to estimate the expected acquisition and misclassification costs that would arise when following the AFA system’s decisions. In more general terms, this corresponds to estimating a target parameter  $J = J(p)$  of some unknown distribution  $p$  given a set of observed samples from  $p$  (the retrospective data set). The goal in semiparametric theory is to find suitable estimators for such a target parameter  $J$  while leaving at least part of the data-generating process  $p$  unrestricted/ unspecified, thereby imposing fewer assumptions, which can lead to more credible estimates. Assumptions that can be taken with a reasonable level of confidence (such as in many AFA settings the NUC or NDE assumptions), can, however, be leveraged to derive more efficient estimators.

A key focus of semiparametric theory is the identification of *influence functions*, which are used to construct estimators with desirable properties, such as a consistently estimable asymptotic variance. The estimator associated with the influence function that has the smallest asymptotic variance is the most efficient one. Often, these influence functions include nuisance functions—unknown components of the model that must be estimated from the data. For example, a nuisance function might model the probability of a doctor acquiring a particular feature. Estimating these nuisance functions typically involves parametric assumptions (e.g., using a logistic regression model). Consequently, the resulting estimator is only *locally efficient*, meaning it achieves efficiency only if the parametric assumptions for the nuisance function hold true.

Even if the parametric assumptions are incorrect, many influence function-based estimators remain consistent due to a property known as *multiple robustness* (Rotnitzky et al., 2017). For instance, the DRL estimators in this paper exhibit a form of double robustness (Scharfstein et al., 1999; Chernozhukov et al., 2018). These estimators rely on learning two nuisance functions and remain consistent as long as the parametric assumption holds for at least one of these functions.

For a more detailed review, please see Appendix B, which covers both the general principles of semiparametric theory and specific insights related to missing data problems.

A thorough understanding of semiparametric theory is primarily necessary for Section 7, which is intended for interested readers.

## 2.5 Active Feature Acquisition Performance Evaluation (AFAPE)

Although we believe to be the first to explicitly formulate and analyze the AFAPE problem, other AFA papers have reported performance metrics that can be seen as attempts to address it. The reported results, however, often lack assumption statements and justification for the chosen evaluation framework and are, in general, biased or inefficiently estimated. We categorize these results based on the viewpoints analyzed in this paper:

*Offline RL view:* The offline RL view has been utilized in the AFA context (Chang et al., 2019; Cheng et al., 2018). As we show in this paper, this approach is only valid under the NUC and strong positivity assumptions.

*Missing data + online RL view:* We show in this paper that one can apply, under the NDE assumption, a missing data + online RL viewpoint to solve AFAPE. While this viewpoint has been taken in the AFA literature before, the missing data part of it has, to our knowledge, only been solved using (conditional) mean imputation (An et al., 2022; Erion et al., 2021; Janisch et al., 2020). (Conditional) mean imputation leads, however, to biased estimation results, as we illustrate in Section 5.

*Semi-offline RL view:* Some AFA papers (Janisch et al., 2020; Yoon et al., 2018) have addressed the problem of missing data during the online RL simulations by simply blocking the corresponding feature acquisitions. This approach is similar to our proposed semi-offline RL view. However, unlike our approach, these papers did not correct for the distribution shift caused by blocking feature acquisitions, resulting in biased estimation results.

## 2.6 No Direct Effect (NDE) Assumption

The only work that, to the best of our knowledge, leverages the NDE assumption in a similar way to our semi-offline RL viewpoint is a series of publications from the causal inference literature around the slightly different problem of evaluation of joint dynamic testing and treatment regimes (Robins et al., 2008; Caniglia et al., 2019; Liu et al., 2021; Neugebauer et al., 2017; Kreif et al., 2021). In this setting, the agent is not only tasked with deciding which features to acquire, but also which treatments to give to the patient. The treatment assignment replaces the need for classification/diagnosis in these settings. Robins et al. (2008) introduced within this setting for the first time the term "no direct effect" (NDE) assumption. NDE stated that the feature acquisition decisions have no direct effect (or no long-term direct effect (Liu et al., 2021)) on the health status of the patient, except through their effect on the treatment decisions.

Caniglia et al. (2019) derived an IPW estimator for this context, which demonstrated a 50-fold increase in data efficiency compared to the offline RL IPW estimator, signaling the enormous benefits that can be achieved by leveraging the NDE assumption. We adapt this estimator to the AFA setting and show that it is equivalent to our proposed IPW estimator for a simple setting and a special positivity assumption. However, our IPW estimator can be applied in more general settings under weaker positivity assumptions and combined with our DM method to form the novel DRL estimator for semi-offline RL.

Liu et al. (2021) developed nearly semiparametrically efficient estimators for this problem by modifying the DRL estimator from the offline RL perspective to enhance efficiency under the NDE assumption. In Section 7, we demonstrate that this approach can be adapted for AFA, where it becomes a specialized method within established semiparametric techniques for missing data problems. Our proposed semi-offline RL estimators can also achieve greater data efficiency using this framework. However, this approach has limitations: it requires strong positivity assumptions, is only applicable in very simple settings with one acquisition action per time point, and the resulting augmentation necessitates complex approximations of function spaces, making implementation challenging and resulting in estimators that are only "nearly" efficient. Additionally, it has only been tested in an extremely simplified context with one acquisition action and one time step (Liu et al., 2021).

## 2.7 Distribution Shift Robust ML Models

Lastly, this work also relates to the general literature on distribution shift-robust ML models. A common problem with the deployment of ML models occurs if the model is trained, for example, on data from hospital 1 but should be deployed to hospital 2. The related literature aims at building robust models that retain their performances across deployment environments (Rockenschaub et al., 2023a,b). One part of the distribution that might change between hospital 1 and hospital 2 is the feature acquisition policy. If this is the case, and if the acquisition policy at hospital 2 is known, one may directly apply our methods to this scenario and treat the acquisition policy at hospital 2 as the AFA policy that is to be evaluated. However, we will not go into more detail about this scenario and will focus on the AFA setting.

## 3. Active Feature Acquisition Performance Evaluation (AFAPE) Problem Definition

We begin the section by introducing the mathematical notation for the AFA setting and AFAPE problem. A glossary containing all the variables and important terms can be found in Appendix C.

### 3.1 Feature Acquisition Process

The feature acquisition process, as illustrated in Figure 2, is modeled using the following variables:

- **Unobserved underlying features:**  $U^t \in \mathbb{R}^d$  for  $t \in \{0, \dots, T\}$ . The features are dynamic, meaning they change over time, so generally,  $U_i^t \neq U_i^{t+1}$ .
- **Feature acquisition decisions:**  $A^t \in \mathcal{A}^t = \{0, 1\}^d$  for  $t \in \{1, \dots, T\}$ , where  $A_i^t = 1$  indicates that feature  $U_i^t$  will be acquired.
- **Observed feature values:**  $X^t \in (\mathbb{R} \cup \{ "? " \})^d$  for  $t \in \{0, \dots, T\}$ , where "?" denotes a missing feature value that was not acquired. We assume no measurement error, so



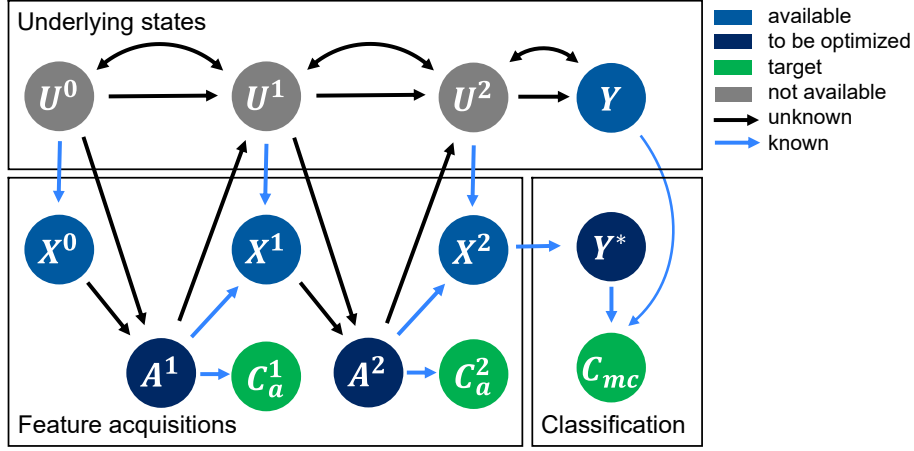


Figure 2: The causal graph depicting the AFA setting as a partially observable decision process consisting of unobserved underlying features  $U^t$ , feature acquisition actions  $A^t$ , feature measurements  $X^t = G_{A^t}(U^t)$ , and associated acquisition costs  $C_a^t$ . After a number of acquisition steps  $T$  (here  $T = 2$ ), a classification  $Y^*$  is to be performed. In the case of misclassification ( $Y^*$  is not equal to the true label  $Y$ ), a misclassification cost  $C_{mc}$  is produced. Edges showing long-term dependencies are omitted from the graph for visual clarity. These include:  $\underline{U}^{t-1}, \underline{X}^{t-1}, \underline{A}^{t-1} \rightarrow A^t$ ;  $\underline{X}^T, \underline{A}^T \rightarrow Y^*$ ;  $A^t \rightarrow \bar{U}^t$ ;  $\underline{U}^{t-1} \leftrightarrow U^t$ ;  $\underline{U}^{t-1} \rightarrow U^t$ ;  $\underline{U}^T \leftrightarrow Y$  and  $\underline{U}^T \rightarrow Y$  (where  $\leftrightarrow$  denotes unobserved confounding).

$X^t$  is deterministically determined by  $A^t$  and  $U^t$  according to:

$$X_i^t = G_{A_i^t}(U_i^t) = \begin{cases} U_i^t & \text{if } A_i^t = 1, \\ \text{"?"} & \text{if } A_i^t = 0. \end{cases} \quad (1)$$

$G_A$  denotes the observation function for a given feature acquisition decision  $A$ . Further, we assume for simplicity  $X^0 = U^0$ .

- **Acquisition costs:**  $C_a^t \in \mathbb{R}$  represents the known feature acquisition cost associated with  $A^t$ .

Additionally, let  $\underline{U}^t$  and  $\bar{U}^t$  denote the complete past and complete future of  $U^t$ , respectively, where  $\underline{U}^t = \{U^0, \dots, U^t\}$  and  $\bar{U}^t = \{U^t, \dots, U^T\}$ . Similarly, define  $\underline{A}^t, \bar{A}^t, \underline{X}^t = G_{\underline{A}^t}(\underline{U}^t)$ , and  $\bar{X}^t = G_{\bar{A}^t}(\bar{U}^t)$  for the variables  $A^t$  and  $X^t$ . Let  $U = \bar{U}^0$ ,  $A = \bar{A}^1$ , and  $X = \bar{X}^0$ , and denote the space of  $A$  by  $\mathcal{A}$ . The retrospective acquisition policy is given by  $\pi_\beta^t(A^t | \underline{X}^{t-1}, \underline{U}^{t-1}, \underline{A}^{t-1})$ .

### 3.2 Classification Process

At time  $T$ , the feature acquisition process concludes, and the classification of an underlying label is performed based on the acquired information. The classification process includes:

- **Underlying categorical label:**  $Y$ , the true label to be classified.
- **Predicted label:**  $Y^*$ , obtained using a deterministic classifier  $Y^* = f_{cl}(\underline{A}^T, \underline{X}^T)$ .
- **Misclassification cost:**  $C_{mc} \in \mathbb{R}$ , the cost incurred when  $Y^*$  differs from  $Y$ , defined by a predefined cost function  $C_{mc} = f_C(Y^*, Y)$ . We alternatively write  $C_{mc} = f_C(f_{cl}(\underline{A}^T, \underline{X}^T), Y) \equiv f_C(\underline{A}^T, \underline{X}^T, Y)$  directly as a function of  $\underline{A}^T, \underline{X}^T$  and  $Y$ .

We let  $g$  represent known deterministic distributions or densities, distinct from  $p$  used for other distributions or densities. For example, we let  $g_{cl}(Y^* | \underline{A}^T, \underline{X}^T)$  denote the deterministic distribution of the classifier  $f_{cl}$  and  $g_C(C_{mc} | Y^*, Y)$  denote the deterministic distribution of the cost function  $f_C$ .

We also assume  $Y$  is always available in the retrospective data set and allow for unobserved confounding among the unobserved underlying features and the label (represented by edges  $\underline{U}^t \leftrightarrow U^{t+1}, Y$ ), but no additional confounding with  $A^t$ . We denote the retrospective data set consisting of the variables  $A, X$ , and  $Y$  as  $\mathcal{D}$ .

### 3.3 Problem Definition: Active Feature Acquisition Performance Evaluation (AFAPE)

Given a target AFA policy  $\pi_\alpha^t(A^t | \underline{X}^{t-1}, \underline{A}^{t-1})$  (which is not allowed to depend on the unobserved underlying features  $\underline{U}^{t-1}$ ) and a target classifier  $f_{cl}(\underline{X}^T, \underline{A}^T)$ , the goal of AFAPE is to estimate the expected acquisition and misclassification costs that would arise, had the target policy  $\pi_\alpha$  and classifier  $f_{cl}$  been deployed. The estimation problem for this expected counterfactual cost can be expressed as estimating

$$J_a = \mathbb{E} \left[ \sum_{t=1}^T C_{a,(\pi_\alpha)}^t \right], \text{ and } J_{mc} = \mathbb{E} [C_{mc,(\pi_\alpha)}], \quad (2)$$

where  $C_{a,(\pi_\alpha)}^t$  and  $C_{mc,(\pi_\alpha)}$  denote the potential outcomes of the acquisition and misclassification costs under the AFA policy  $\pi_\alpha$ . Therefore,  $J_a$  and  $J_{mc}$  represent the expected acquisition and misclassification costs under a distribution induced by  $\pi_\alpha$  rather than by the retrospective acquisition policy  $\pi_\beta$ . Note that this assumes  $\pi_\alpha$  can be followed perfectly, which may not always hold in practice, for example, if patients refuse certain medical tests or miss appointments.

The goal of this paper is to i) perform identification, i.e., to determine under which assumptions it is possible to resolve this distribution shift and to obtain an unbiased estimate of  $J_a$  and  $J_{mc}$ ; and ii) to derive such unbiased estimators.

As the AFAPE problem is similar for  $J_a$  and  $J_{mc}$ , we will focus on  $J_{mc}$  throughout the main part of the paper. We abbreviate  $J_{mc} \equiv J$  and  $C_{mc} \equiv C$ . We provide the estimation formulas for  $J_a$  and for  $J_{mc}$  when a prediction is to be performed at each time step in the relevant appendices.

### 3.4 Problem Definition: Optimization of Active Feature Acquisition Methods

While the focus of the paper is on the AFAPE problem, we provide the definition of the AFA optimization problem for completeness. The goal in AFA is to find the optimal AFA policy  $\pi_\alpha^t(A^t | \underline{X}^{t-1}, \underline{A}^{t-1}; \phi_1^*)$  parameterized by  $\phi_1^*$ , and the optimal classifier  $f_{cl}(\underline{X}^T, \underline{A}^T; \phi_2^*)$

parameterized by  $\phi_2^*$ , such that their joint application minimizes the expected sum of counterfactual acquisition and misclassification costs:

$$\phi_1^*, \phi_2^* = \arg \min_{\phi_1, \phi_2} J_{\text{total}}(\phi_1, \phi_2) = \arg \min_{\phi_1, \phi_2} \mathbb{E} \left[ \sum_{t=1}^T C_{a,(\pi_\alpha)}^t + C_{mc,(\pi_\alpha)} \middle| \phi_1, \phi_2 \right].$$

### 3.5 Assumptions

Here, we provide an overview of the key assumptions in this paper. We start by stating the fixed assumptions that hold throughout the paper before stating assumptions that we vary within different sections.

#### 3.5.1 FIXED ASSUMPTIONS

We make the following assumptions throughout the paper:

**Assumption 1 (No measurement noise)** There is no noise in feature measurements, as expressed by Eq. 1.

**Assumption 2 (Consistency)** If  $\underline{A}^t = \underline{a}^t$ , then  $U_{(\underline{a}^t)}^t = U^t$ .

This standard consistency assumption from the causal inference literature states that an individual’s observed outcomes align with their potential outcomes under the observed acquisition decisions. Here,  $U_{(\underline{A}^t = \underline{a}^t)}^t$  represents the potential outcome of  $U^t$  under potential acquisition decisions  $\underline{A}^t = \underline{a}^t$ .

**Assumption 3 (No interference)** The acquisition decisions for one individual do not affect other individuals.

One prominent example of interference in medical settings is allocation interference, which can occur when a hospital’s resources or staff are overwhelmed by a high volume of medical test requests for multiple patients simultaneously, resulting in the inability to fulfill all feature acquisition requests.

#### 3.5.2 INVESTIGATED ASSUMPTIONS

In this paper, we analyze how the following assumptions affect identification and estimation of the target  $J$  in the AFAPE problem.

**Assumption 4 (No direct effect (NDE))** The unobserved underlying features are not influenced by feature acquisitions (i.e.,  $A^t \not\rightarrow \bar{U}^t$ ). Equivalently:  $U^t \perp\!\!\!\perp \underline{A}^t \mid \underline{U}^{t-1}$ .

This is a standard assumption in missing data problems, but it may not hold in all medical settings, as some medical tests can alter certain features of the patient. The NDE assumption is relaxed in Section 4, and made in Sections 5, 6 and 7.

**Assumption 5 (No unobserved confounding (NUC))** Acquisition decisions are independent of the unobserved underlying features given past acquisition decisions and measured features:  $A^t \perp\!\!\!\perp \underline{U}^{t-1} \mid \underline{X}^{t-1}, \underline{A}^{t-1}$ . This is graphically expressed by the missing arrow  $\underline{U}^{t-1} \not\rightarrow A^t$ .

The no unobserved confounding assumption may, for example, be violated in medical settings if certain feature values are seen by the physician and influence their decision-making for further tests but are not recorded in the database. We assume NUC in Sections 4, 6 and 7 and allow certain violations in Section 5. Note that when referring to NUC, we only assume no unobserved confounding of the acquisition actions. Potential unobserved confounding within  $U$  and between  $U$  and  $Y$  is allowed throughout the paper.

**Assumption 6 (Positivity/ experimental treatment assignment/ overlap)** Certain feature sets have a positive probability of being acquired under the retrospective acquisition policy  $\pi_\beta$ .

The positivity assumption (also known as experimental treatment assignment assumption or overlap assumption) relates to how much exploration was done under the retrospective acquisition policy  $\pi_\beta$ . Positivity requirements are crucial for identification and vary between the discussed views. Hence, we derive and discuss them separately for each view.

#### 4. Offline Reinforcement Learning View

*Assumptions in this section: Assumption 5 (NUC)*

Firstly, we consider the scenario where the NUC assumption (Assumption 5) holds (i.e.  $\underline{U}^{t-1} \not\rightarrow A^t$ ), but the NDE assumption (Assumption 4) does not necessarily hold (i.e.  $A^t \rightarrow \bar{U}^t$ ). This scenario can be addressed using the offline reinforcement learning (RL) view. The NUC assumption allows us to perform a latent projection (Verma and Pearl, 1990) to project out the unknown variables  $U^t$  (along with  $Y$  and  $Y^*$ ) from the causal graph in Figure 2 and obtain the graph in Figure 3 which contains only observed variables. The projected graph allows us to apply established identification and estimation methods from the offline RL literature.




Figure 3: Updated causal graph of the AFA setting under the NUC assumption (Assumption 5) and a latent projection. The graph depicts a standard, identified offline RL setting. Long-term dependencies are omitted from the graph for visual clarity. These include edges  $\underline{X}^{t-1}, \underline{A}^{t-1} \rightarrow A^t$ ;  $\underline{X}^T, \underline{A}^T \rightarrow C$ ;  $\underline{X}^{t-1} \leftrightarrow X^t$  and  $\underline{X}^T \leftrightarrow C$ .

##### 4.1 Identification

Under the offline RL view, solving the AFAPPE problem is equivalent to performing off-policy policy evaluation (OPE). Identification for OPE requires sequential exchangeability (also known as sequential ignorability), which implies that adjusting for  $\underline{A}^{t-1}$  and  $\underline{X}^{t-1}$  eliminates any confounding factors affecting  $A^t$ . The graph in Figure 3 satisfies this requirement. Note

that sequential exchangeability would not hold under violations of the NUC assumption (Assumption 5) because the latent projection under the presence of edges  $\underline{U}^{t-1} \rightarrow A^t$  would produce confounding edges  $\bar{X}^t \leftrightarrow A^t$  and  $C \leftrightarrow A^t$ .

Identification of  $J$  further requires Assumptions 1-3 (no measurement noise, consistency, no interference) and the following (sequential) positivity assumption:

**Assumption 6.1 (Positivity for offline RL)**

if 
$$p(\underline{X}_{(\pi_\alpha)}^{t-1} = \underline{x}^{t-1}, \underline{A}_{(\pi_\alpha)}^{t-1} = \underline{a}^{t-1})\pi_\alpha^t(a^t|\underline{x}^{t-1}, \underline{a}^{t-1}) > 0,$$
 then 
$$p(\underline{x}^{t-1}, \underline{a}^{t-1})\pi_\beta^t(a^t|\underline{x}^{t-1}, \underline{a}^{t-1}) \geq \mathcal{O}$$
 
$$\forall t, a^t, \underline{x}^{t-1}, \underline{a}^{t-1}, \text{ and some constant } \mathcal{O} > 0$$

where we introduced the following notation:  $\pi_\alpha^t(a^t|\underline{x}^{t-1}, \underline{a}^{t-1}) \equiv \pi_\alpha^t(A^t = a^t|\underline{X}^{t-1} = \underline{x}^{t-1}, \underline{A}^{t-1} = \underline{a}^{t-1})$ .

The positivity assumption states that, for every set of actions and observations  $\underline{a}^{t-1}, \underline{x}^{t-1}$  reachable under  $\pi_\alpha$  and a desired next action  $a^t$  (i.e., an action with positive support under  $\pi_\alpha$ ), we require also positive support for  $a^t$  under  $\pi_\beta$ . A violation of this assumption may occur if the acquisition decisions under the AFA policy  $\pi_\alpha$  differ significantly from the decisions made by doctors ( $\pi_\beta$ ).

**4.2 Estimation**

Estimation can be performed using well-known techniques from the offline RL / DTR literature. The following are common estimators:

1) *Model-based evaluation:*

The model-based estimator, also known in the causal inference literature as the noniterative conditional expectation (NICE) estimator of the G-formula (Rein et al., 2024; Wen et al., 2021), estimates the target cost  $J$  as follows:

$$\hat{J}_{MB-Off} = \sum_{X,A} \hat{\mathbb{E}}[C|\underline{X}^T, \underline{A}^T] \prod_{t=1}^T \hat{p}(X^t|\underline{X}^{t-1}, \underline{A}^t)\pi_\alpha^t(A^t|\underline{X}^{t-1}, \underline{A}^{t-1}) \tag{3}$$

where the integration over  $X$  and  $A$  can be solved using Monte Carlo integration. Note that we use sums to denote the integration over  $X$ . All results in this paper do, however, also hold for continuous  $X$ , by replacing the sums with proper integrals.

This estimator requires learning the state transition function  $p(X^t|\underline{X}^{t-1}, \underline{A}^t)$  and the expected cost  $\mathbb{E}[C|\underline{X}^T, \underline{A}^T]$ . We denote the learned nuisance functions as  $\hat{p}(X^t|\underline{X}^{t-1}, \underline{A}^t)$  and  $\hat{\mathbb{E}}[C|\underline{X}^T, \underline{A}^T]$ .

2) *Inverse probability weighting (IPW):*

The target cost that is estimated by IPW (Levine et al., 2020) is

$$\hat{J}_{IPW-Off} = \hat{\mathbb{E}}_n [\hat{\rho}_{Off}^T C], \text{ where } \hat{\rho}_{Off}^T = \prod_{t=1}^T \frac{\pi_\alpha^t(A^t|\underline{X}^{t-1}, \underline{A}^{t-1})}{\hat{\pi}_\beta^t(A^t|\underline{X}^{t-1}, \underline{A}^{t-1})}. \tag{4}$$

where  $\hat{\mathbb{E}}_n[\cdot]$  denotes the empirical average over the data set  $\mathcal{D}$ . This estimator requires learning the retrospective acquisition policy/propensity score model  $\pi_\beta^t(A^t|\underline{X}^{t-1}, \underline{A}^{t-1})$ .

**Remark 1 (Cross-fitting estimators)** When nuisance functions are estimated using flexible machine learning methods, the training of these functions and the evaluation of the estimator must be conducted on separate data splits to avoid introducing bias. To enhance efficiency, this process can be performed by alternating the training and evaluation splits, a method known as cross-fitting (Chernozhukov et al., 2018). Throughout this paper, we assume that cross-fitting is used, though this is not explicitly noted in the proposed estimators' notation, leading to a slight abuse of notation.

3) *Direct method (DM):*

The DM estimator, also known as the iterative conditional expectation (ICE) estimator of the G-formula (Wen et al., 2021), estimates the target cost  $J$  as:

$$\hat{J}_{DM-Off} = \hat{\mathbb{E}}_n[\hat{V}_{Off}^0]. \quad (5)$$

This estimator relies on learning a state-action value function  $Q_{Off}^t$  or state value function  $V_{Off}^t$ :

$$\begin{aligned} Q_{Off}^t &\equiv Q_{Off}(\underline{X}^{t-1}, \underline{A}^t) \equiv \mathbb{E}[C_{(\pi_\alpha^{t+1})} | \underline{X}^{t-1}, \underline{A}^t], \\ V_{Off}^t &\equiv V_{Off}(\underline{X}^t, \underline{A}^t) \equiv \mathbb{E}[C_{(\pi_\alpha^{t+1})} | \underline{X}^t, \underline{A}^t]. \end{aligned}$$

where  $C_{(\pi_\alpha^{t+1})}$  denotes the potential outcome of  $C$  under a policy intervention  $\pi_\alpha$  applied only from time step  $t + 1$  onwards.  $Q_{Off}$  and  $V_{Off}$  can be learned using, for example, the dynamic programming (DP) algorithm, which is based on the recursive property of the Bellman equation (Bertsekas, 2012):

$$Q_{Off}(\underline{X}^{t-1}, \underline{A}^t) = \sum_{X^t} V_{Off}(\underline{X}^t, \underline{A}^t) p(X^t | \underline{X}^{t-1}, \underline{A}^t) \quad (6)$$

$$V_{Off}(\underline{X}^t, \underline{A}^t) = \sum_{A^{t+1}} Q_{Off}(\underline{X}^t, \underline{A}^{t+1}) \pi_\alpha^{t+1}(A^{t+1} | \underline{X}^t, \underline{A}^t) \quad (7)$$

In practice, one only needs to learn  $Q_{Off}$  such that  $V_{Off}$  can be simply computed as  $V_{Off}^t = \mathbb{E}_{\pi_\alpha}[Q_{Off}^{t+1}]$  using, for example, Monte Carlo integration over the known AFA policy  $\pi_\alpha$ .

4) *Double reinforcement learning (DRL):*

The target cost that is estimated by DRL (Kallus and Uehara, 2020) is

$$\hat{J}_{DRL-Off} = \hat{\mathbb{E}}_n \left[ \hat{\rho}_{Off}^T C + \sum_{t=1}^T \left( -\hat{\rho}_{Off}^t \hat{Q}_{Off}^t + \hat{\rho}_{Off}^{t-1} \hat{V}_{Off}^{t-1} \right) \right]. \quad (8)$$

The DRL estimator combines approaches 2) and 3) by using both the learned propensity score  $\hat{\pi}_\beta$  and the state action value function  $\hat{Q}_{Off}$  (and the derived  $\hat{V}_{Off}$ ). This estimator is (locally) efficient and doubly robust, in the sense that it is consistent if either the propensity score model  $\hat{\pi}_\beta$ , or the state action value function  $\hat{Q}_{Off}$  is correctly specified (Kallus and Uehara, 2020).

## 5. Missing Data (+ Online Reinforcement Learning) View

*Assumptions in this section: Assumption 4 (NDE)*

In this section, we assume that the NDE assumption holds (i.e.,  $A^t \not\leftrightarrow \bar{U}^t$ ), but do not require the NUC assumption (Assumption 5). We observe that the general AFA graph from Figure 2 transforms under NDE into the graph shown in Figure 4A). This new graph represents a temporal missing data graph (m-graph) (Mohan et al., 2013; Shpitser et al., 2015) from the missing data literature. The unobserved underlying feature values  $U^t$  are replaced by counterfactuals of the measured feature values  $X_{(1)}^t$  since  $U_{(\underline{a}^t)}^t = U_{(a^t)}^t = U_{(1)}^t = X_{(1)}^t$  for all potential acquisitions  $\underline{a}^t$  and thus also  $a^t = \vec{1}$ . Due to the temporal restrictions  $\bar{X}_{(1)}^t \not\leftrightarrow A^t$ , the shown graph can be more precisely specified as the known block-conditional missing data model (Zhou et al., 2010). The graph depicting the counterfactual distribution is shown in Figure 4B).

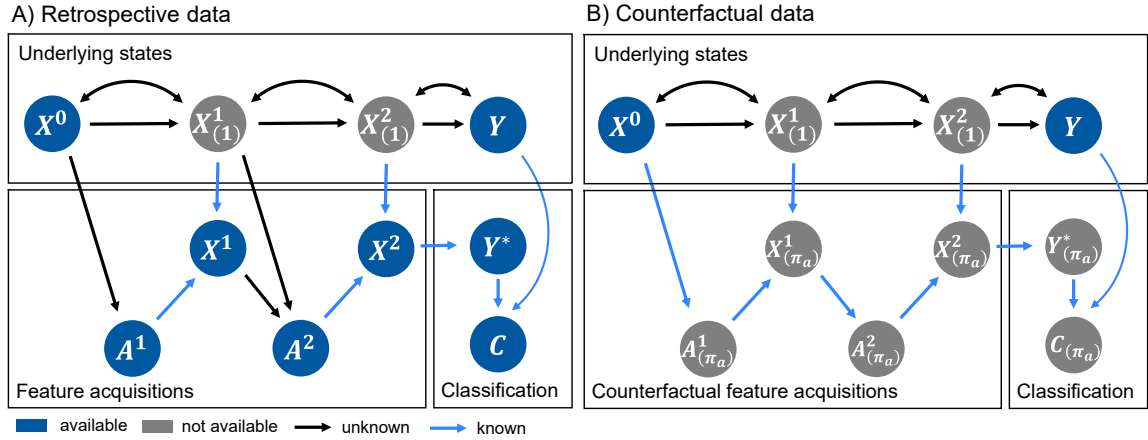


Figure 4: A) Updated causal graph of the AFA process under the NDE assumption (Assumption 4). Unknown state variables  $U^t$  are replaced with the counterfactual feature values  $X_{(1)}^t$ , which represent the values  $X^t$  would have taken if  $A^t$  was  $\vec{1}$  (i.e., the decision to observe all feature values). This graph describing the feature acquisition process is known as a missing data graph (m-graph). B) Graph showing the counterfactual distribution under  $\pi_\alpha$ . Edges showing long-term dependencies are omitted for visual clarity. These include for both graphs  $\underline{X}_{(1)}^{t-1} \leftrightarrow X_{(1)}^t$  and  $\underline{X}_{(1)}^T \leftrightarrow Y$ ; for A)  $\underline{X}^{t-1}, \underline{X}_{(1)}^{t-1}, \underline{A}^{t-1} \rightarrow A^t$ , and  $\underline{X}^T, \underline{A}^T \rightarrow Y^*$ ; and for B)  $X^0, \underline{X}_{(\pi_\alpha)}^{t-1}, \underline{A}_{(\pi_\alpha)}^{t-1} \rightarrow A^t_{(\pi_\alpha)}$  and  $X^0, \underline{X}_{(\pi_\alpha)}^T, \underline{A}_{(\pi_\alpha)}^T \rightarrow Y^*_{(\pi_\alpha)}$ .

### 5.1 Problem Reformulation

Now, we can establish the following theorem. Under the NDE assumption, the AFAP problem becomes equivalent to a standard missing data problem, for which one can apply known identification and estimation theory.

**Theorem 1.** (AFAPE problem reformulation and identification under the missing data view). The AFAPE problem of estimating  $J$  (Equation 2) is under Assumption 1 (no measurement noise), Assumption 2 (consistency), Assumption 3 (no interference), and Assumption 4 (NDE) equivalent to estimating

$$J = \sum_{X_{(1)}, Y} \underbrace{\mathbb{E}[C_{(\pi_\alpha)} | X_{(1)}, Y]}_{\text{online RL}} \underbrace{p(X_{(1)}, Y)}_{\text{missing data}}. \quad (9)$$

Furthermore,  $J$  is identified if  $p(X_{(1)}, Y)$  is identified.

**Proof** The decomposition of  $J$  into the two expected values follows from the law of iterated expectations and the independence of  $X_{(1)}, Y$  from a policy intervention  $\pi_\alpha$ . The fact that the inner expected value is identified can be easily verified by examining the graph representing the counterfactual distribution shown in Figure 4B). The graph shows that all functional relationships between variables that are part of the feature acquisition and classification processes are known (represented as blue edges). In particular, this implies the following factorization:

$$\mathbb{E}[C_{(\pi_\alpha)} | X_{(1)}, Y] = \sum_{\underline{X}_{(\pi_\alpha)}^T, \underline{A}_{(\pi_\alpha)}^T, Y_{(\pi_\alpha)}^*, C_{(\pi_\alpha)}} C_{(\pi_\alpha)} q(C_{(\pi_\alpha)}, Y_{(\pi_\alpha)}^*, X_{(\pi_\alpha)}, A_{(\pi_\alpha)} | X_{(1)}, Y) \quad (10)$$

with the identifying distribution

$$\begin{aligned} q(C_{(\pi_\alpha)}, Y_{(\pi_\alpha)}^*, X_{(\pi_\alpha)}, A_{(\pi_\alpha)} | X_{(1)}, Y) &= \\ &= \prod_{t=0}^T \underbrace{g(X_{(\pi_\alpha)}^t | A_{(\pi_\alpha)}^t, X_{(1)}^t)}_{\text{feature revelations}} \prod_{t=1}^T \underbrace{\pi_\alpha^t(A_{(\pi_\alpha)}^t | \underline{X}_{(\pi_\alpha)}^{t-1}, \underline{A}_{(\pi_\alpha)}^{t-1})}_{\text{acquisition decisions}} \underbrace{g_{cl}(Y_{(\pi_\alpha)}^* | \underline{X}_{(\pi_\alpha)}^T, \underline{A}_{(\pi_\alpha)}^T)}_{\text{label prediction}} \underbrace{g_C(C_{(\pi_\alpha)} | Y, Y_{(\pi_\alpha)}^*)}_{\text{cost computation}} \end{aligned}$$

which is identified since all (deterministic) distributions  $g(\cdot)$  and  $\pi_\alpha$  are known functions. ■

Since all densities  $g(\cdot)$  are deterministic and known, we can further simplify Eq. 10 using simpler notation. We use  $X = G_A(X_{(1)})$ ,  $\pi_\alpha(a | G_a(X_{(1)})) \equiv \prod_{t=0}^T \pi_\alpha^t(a^t | G_{\underline{a}^{t-1}}(X_{(1)}), \underline{a}^{t-1})$  and  $C = f_C(A, G_A(X_{(1)}), Y)$  to obtain:

$$\begin{aligned} \mathbb{E}[C_{(\pi_\alpha)} | X_{(1)}, Y] &= \sum_{a \in \mathcal{A}} f_C(\underline{a}^T, G_{\underline{a}^T}(X_{(1)}), Y) \prod_{t=0}^T \pi_\alpha^t(a^t | G_{\underline{a}^{t-1}}(X_{(1)}), \underline{a}^{t-1}) \quad (11) \\ &= \sum_{a \in \mathcal{A}} f_C(a, G_a(X_{(1)}), Y) \pi_\alpha(a | G_a(X_{(1)})). \end{aligned}$$

We denote  $\mathbb{E}[C_{(\pi_\alpha)} | X_{(1)}, Y]$  as the online RL part of the problem because it involves the evaluation of a policy in a known environment. We refer to the outer expected value of  $p(X_{(1)}, Y)$  as the missing data problem, as it requires the identification of the counterfactual feature distribution.



## 5.2 Identification

As established in Theorem 1, the AFAPE problem is identified if the missing data problem (i.e.,  $p(X_{(1)}, Y)$ ) is identified. The following positivity assumption is required to allow identification of  $p(X_{(1)}, Y)$  and therefore for the target parameter  $J$  from Eq. 9:

### Assumption 6.2 (Positivity for missing data)

$$\begin{aligned} \text{if} \quad & p(\underline{X}_{(1)}^{t-1} = \underline{x}^{t-1}, \underline{A}^{t-1} = \vec{1}) > 0, \\ \text{then} \quad & \pi_{\beta}^t(A^t = \vec{1} | \underline{X}_{(1)}^{t-1} = \underline{x}^{t-1}, \underline{A}^{t-1} = \vec{1}) \geq \mathcal{O} \\ & \forall t, \underline{x}^{t-1}, \text{ and some constant } \mathcal{O} > 0 \end{aligned}$$

This positivity assumption is very different from the positivity assumption assumed under the offline RL view (Assumption 6.1). It requires the "acquire everything" action trajectory  $A^t = \vec{1} \forall t$  to have positive support for all possible feature values. In other words, this is a requirement for complete cases among all subpopulations.

Given the positivity assumption, the block-conditional model describing  $p(X_{(1)}, Y)$  is identified. We show how identification can be achieved in Appendix D.

Note that Theorem 1 holds even in the more general case without the temporal restriction  $\overline{X}_{(1)}^t \not\rightarrow A^t$ . In this case,  $p(X_{(1)}, Y)$  may or may not be identified, depending on what assumptions can be made. There exists a vast literature on identification theory for missing data problems (Bhattacharya et al., 2020; Nabi et al., 2020; Mohan and Pearl, 2021) that can be applied.

Here, we briefly discuss how the common missing data scenarios of missing-completely-at-random (MCAR), missing-at-random (MAR), and missing-not-at-random (MNAR) apply within the AFA framework. MNAR scenarios arise when the missingness of a feature is influenced by the value of another feature which may itself be missing. This situation occurs when the NUC assumption is violated, indicated by the presence of arrows  $X_{(1)} \rightarrow A$ . On the other hand, MAR scenarios occur when the missingness of a feature is dependent only on observed features. This scenario aligns with our setting if the NUC assumption (Assumption 5) holds. Lastly, MCAR represents the simplest case, where the missingness of a feature is independent of all other feature values. In our framework, this corresponds to the absence of any edges  $X \rightarrow A$ .

## 5.3 Estimation

An estimate of  $\mathbb{E}[C_{(\pi_{\alpha})} | X_{(1)}, Y]$ , denoted as  $\hat{\mathbb{E}}[C_{(\pi_{\alpha})} | X_{(1)}, Y]$ , can be readily computed from Eq. 11 using Monte Carlo integration:

$$\hat{\mathbb{E}}[C | X_{(1)}, Y] = \sum_i^{n_{MC}} f_C(a_i, G_{a_i}(X_{(1)}), Y) \quad (12)$$

with  $n_{MC}$  samples  $a_i \sim \prod_{t=1}^T \pi_{\alpha}(A^t | G_{\underline{A}^{t-1}}(X_{(1)}), \underline{A}^{t-1})$ .

This approach is common in online RL settings, where the agent interacts with the environment. Monte Carlo integration is performed in the following way: First, one takes

a sample  $x, y$  from  $p(X_{(1)}, Y)$  (obtained through either of the methods shown below) and reveals the initial features  $x^0$  to the agent. The agent’s first action, denoted by  $a^1$ , is then sampled from  $\pi_\alpha(A^1|x^0)$ . Depending on  $a^1$ , the corresponding feature set amongst  $x^1$  is revealed to the agent and the next action  $a^2$  is sampled. This continues until step  $T$ , when the classifier is applied (using the acquired subset of features from the simulation as input) and a resulting misclassification cost  $c'$  is computed. This process is repeated multiple times, and the costs are averaged to obtain  $\hat{\mathbb{E}} [C_{(\pi_\alpha)}|X_{(1)}, Y]$ .

This online RL process requires samples from  $p(X_{(1)}, Y)$ . These samples can be obtained using standard missing data estimation methods, which result in the following estimators:

1) *Inverse probability weighting (IPW):*

The target cost that is estimated by IPW (Seaman and White, 2013) is

$$\hat{J}_{IPW-Miss} = \hat{\mathbb{E}}_n \left[ \hat{\rho}_{Miss} \hat{\mathbb{E}} [C_{(\pi_\alpha)}|X_{(1)}, Y] \right], \text{ where } \hat{\rho}_{Miss} = \prod_{t=1}^T \frac{\mathbb{I}(A^t = \vec{1})}{\hat{\pi}_\beta^t(A^t = \vec{1} | \underline{X}_{(1)}^{t-1}, \underline{A}^{t-1} = \vec{1})}, \tag{13}$$

and where  $\mathbb{I}(\cdot)$  denotes the indicator function. This IPW approach requires learning the propensity score  $\pi_\beta$ , but only for the scenario of full data acquisition (where  $A = \vec{1}$ ). Because of the indicator function  $\mathbb{I}(A^t = \vec{1})$ , only the complete cases are selected for reweighting.

2) *Multiple Imputation (MI):*

The target cost that is estimated by the multiple imputation (MI) (Sterne et al., 2009) estimator is:

$$\hat{J}_{MI-Miss} = \hat{\mathbb{E}}_n \left[ \sum_{X_m} \hat{\mathbb{E}} [C_{(\pi_\alpha)}|X_{(1)}, Y] \hat{p}(X_m|X_o, Y) \right]. \tag{14}$$

Missing part of  $X_{(1)}$ 
↓
↑
Observed part of  $X_{(1)}$

This estimator is based on the G-formula, which requires the counterfactual data distribution  $p(X_{(1)}, Y)$ . In the MI estimator, this density is not modeled fully, but the empirical distribution of the available data is augmented with samples (i.e., imputations) from a model for the missing data. The MI estimator is based on the decomposition  $\hat{p}(X_{(1)}, Y) = \hat{p}(X_m|X_o, Y)p(X_o, Y)$ , where  $X_o$  denotes the observed part and  $X_m$  the missing part of  $X_{(1)}$ . The sampling of the missing part is then usually repeated multiple times to increase the precision of the estimate, hence the name "multiple imputation".

In certain scenarios, MI can outperform the estimators from the semi-offline RL view, which will be described next. Appendix E discusses the advantages and disadvantages of the MI estimator in more detail. The appendix also highlights why using (conditional) mean imputation, which has been previously employed in AFA settings (An et al., 2022; Erion et al., 2021; Janisch et al., 2020) generally leads to biased estimation results.

## 6. Semi-offline Reinforcement Learning View

*Assumptions in this section: Assumption 4 (NDE), Assumption 5 (NUC)*

In this section, we assume both the NDE and NUC assumptions hold. In this context,

either the offline RL or the missing data view can be applied to solve AFAPE, but both have limitations, as illustrated in the following scenarios.

We assume an available data point without  $X_{(1)}^0$ , but with two time-steps and univariate  $X_{(1)}^1 = 0.6$  and  $X_{(1)}^2 = 0.9$ . Initially, we assume the retrospective acquisition decisions were  $A^1 = 1$  and  $A^2 = 1$  (denoted as scenario 1). The data point in scenario 1 corresponds to a complete case where both feature values  $X_{(1)}^1$  and  $X_{(1)}^2$  are known.

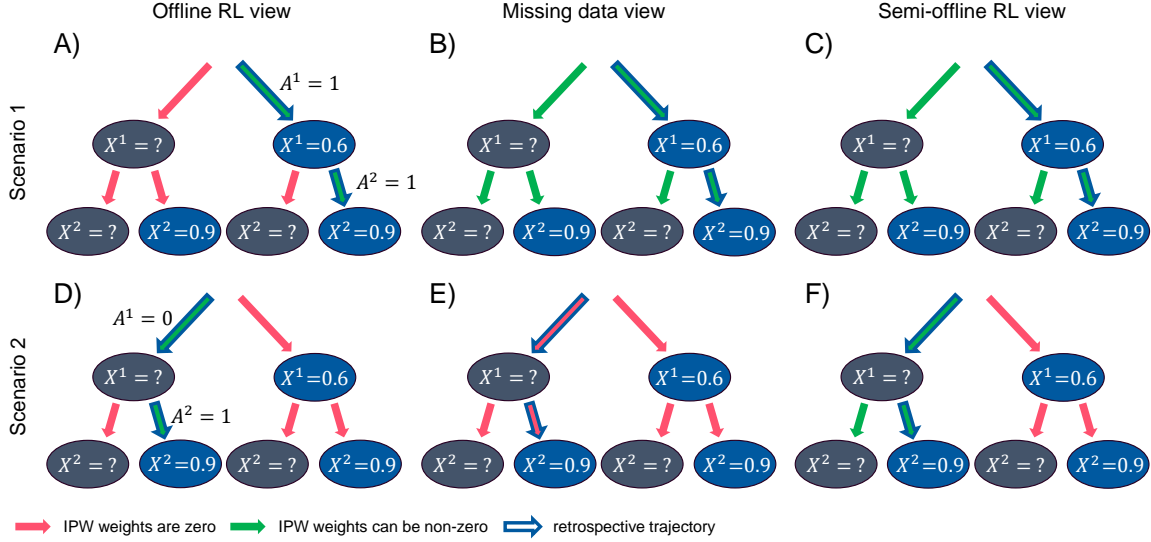


Figure 5: Visualization of data utilization by IPW estimators under different views. Each graph shows the four possible target acquisition trajectories for two exemplary retrospective acquisition scenarios and highlights which target trajectories can receive non-zero IPW weights under the respective views. A), D) The IPW estimator from the offline RL viewpoint: only target trajectories that match the retrospective trajectory can be evaluated. B), E) The IPW estimator from the missing data viewpoint: all trajectories can be evaluated if the datapoint is a complete case; otherwise, no evaluation is possible. C), F) The IPW estimator from the semi-offline RL viewpoint: all trajectories with equal or fewer acquisitions than the retrospective trajectory can be evaluated.

We analyze how this data point is utilized by the IPW estimators within different views. Figures 5A) and 5B) show the four possible target trajectories:  $a^1 = 0, a^2 = 0$ ;  $a^1 = 1, a^2 = 0$ ;  $a^1 = 0, a^2 = 1$ ; and  $a^1 = 1, a^2 = 1$  under the offline RL and missing data views, respectively. In the offline RL view, this data point can only be used to evaluate the target trajectory  $a^1 = 1$  and  $a^2 = 1$ , matching the retrospective trajectory. The missing data IPW estimator, however, can evaluate any of the four target trajectories using this data point since it is a complete case.

Now consider scenario 2 with hypothetical retrospective acquisition decisions  $A^1 = 0$  and  $A^2 = 1$ . Figures 5D) and 5E) depict the corresponding trajectories for both views. The

offline RL IPW estimator can use this datapoint to evaluate the matching trajectory  $a^1 = 0$  and  $a^2 = 1$  but none others. The missing data IPW estimator cannot evaluate any target trajectories as this is not a complete case.

In scenario 2, we only know the value of  $X_{(1)}^2$ , not  $X_{(1)}^1$ , since  $A^1 = 0$ . However, the target trajectories  $a^1 = 0, a^2 = 0$  and  $a^1 = 0, a^2 = 1$  do not require the value of  $X_{(1)}^1$  and can still be simulated. This motivates our novel semi-offline RL viewpoint, where we consider all target trajectories with the same or fewer acquisitions, compared with the retrospective data point, as simulatable.

Figures 5C) and 5F) show the simulatable trajectories under the semi-offline RL view for both scenarios. Similar to online RL, a policy can sample different trajectories, but under the semi-offline RL view, only among the simulatable ones, while trajectories involving non-simulatable actions (like  $a^1 = 1$  in scenario 2) should not be sampled.

Since not all trajectories are simulatable, we restrict the simulation policy to block actions that would result in non-simulatable trajectories. A simple Monte Carlo estimator averaging the costs within the simulated trajectories would be biased due to the blocking of actions, necessitating post-simulation bias correction.

The remainder of this section formalizes the semi-offline RL viewpoint and is organized as follows. First, we introduce the simulation policy, referred to as the semi-offline sampling policy, and illustrate that simulations using this policy do not require information about  $X_{(1)}$  not already contained in  $X = G_A(X_{(1)})$ . Next, we explain how the AFAPE target  $J$  can be equivalently formulated based on the semi-offline sampling distribution. We then prove that this reformulated  $J$  is identified under a new, weaker positivity assumption, and finally, we derive novel estimators for  $J$ .

### 6.1 The Semi-offline Sampling Policy

To introduce the semi-offline RL view, we first revisit the problem formulation from the missing data view (Eqs. 9 and 11):

$$J = \sum_{X_{(1)}, Y} \underbrace{\sum_{a \in \mathcal{A}} f_C(a, G_a(X_{(1)}), Y) \pi_\alpha(a | G_a(X_{(1)}))}_{\text{online RL}} \underbrace{p(X_{(1)}, Y)}_{\text{missing data}}. \quad (15)$$

The online RL part involves integrating over all possible trajectories  $a \in \mathcal{A}$ , representing subsets of the features. Notably, many terms in this inner integral do not require complete knowledge of  $X_{(1)}$ . Specifically, any summand with  $a \leq A$  (where we let  $\leq$  denote an element-wise comparison) does not need information from  $X_{(1)}$  beyond what is already in  $X = G_A(X_{(1)})$ .

This missing data + online RL approach effectively says: "Solve all missingness, then integrate over all possible feature subsets, including many that didn't require addressing the missingness to begin with." This approach is clearly suboptimal. Instead, we propose a novel semi-offline RL viewpoint: "Integrate over all subsets of the observed data and adjust for the bias introduced by excluding subsets where data was missing." We call this the semi-offline RL viewpoint because some subsets/trajectories can be simulated (the online part) while others cannot (the offline part).

We define the semi-offline sampling distribution to include the observed data trajectories. To avoid integrating over unobserved feature subsets, we replace  $\pi_\alpha$  in the distribution of Eq. 15 with a policy  $\pi'_{sim}$  that has no support for trajectories where the corresponding  $X_{(1)}$  are missing. A policy  $\pi'$  that enforces this exclusion is defined formally as a blocked policy:

**Definition 2.** (*Blocked Policy*). A policy  $\pi^t(A^t|G_{\underline{A}^{t-1}}(X_{(1)}), \underline{A}^{t-1}, A^t)$  is called a 'blocked policy' of the policy  $\pi^t(A^t|G_{\underline{A}^{t-1}}(X_{(1)}), \underline{A}^{t-1})$  if it satisfies the following conditions:

1) *Blocking of acquisitions of non-available features:*

$$\text{if } a^t \not\leq \underline{a}^t, \quad \text{then } \pi^t(a^t|G_{\underline{a}^{t-1}}(X_{(1)}) = \underline{x}^{t-1}, \underline{a}^{t-1}, a^t) = 0 \quad \forall t, a^t, \underline{a}^t, \underline{x}^{t-1}, \underline{a}^{t-1}$$

2) *No blocking of acquisitions of available features:*

$$\text{if } a^t \leq \underline{a}^t \text{ and } \pi^t(a^t|\underline{x}^{t-1}, \underline{a}^{t-1}) > 0, \quad \text{then } \pi^t(a^t|\underline{x}^{t-1}, \underline{a}^{t-1}, a^t) > 0 \quad \forall t, a^t, \underline{a}^t, \underline{x}^{t-1}, \underline{a}^{t-1}$$

Condition 1 ensures that sampling does not depend on values of  $X_{(1)}$  that are not contained in  $X = G_A(X_{(1)})$ . Condition 2 ensures that the online exploration part is utilized by forcing the blocked policy  $\pi'$  to have positive support whenever  $\pi$  has positive support and the desired features are available. A practical choice for  $\pi'$  just sets  $\pi'(A^t = a^t|G_{\underline{A}^{t-1}}(X_{(1)}), \underline{A}^{t-1}, A^t = a^t) = 0$  if  $a^t \not\leq \underline{a}^t$  and rescales the other probabilities accordingly.

Having defined the blocked policy, we can now introduce the semi-offline sampling distribution:

$$\begin{aligned} p'(A', G_A(X_{(1)}), Y, A) &= \prod_{t=1}^T \pi_{sim}^t(A^t|G_{\underline{A}^{t-1}}(X_{(1)}), \underline{A}^{t-1}, A^t) p(A, G_A(X_{(1)}), Y) \quad (16) \\ &\equiv \underbrace{\pi'_{sim}(A'|G_{A'}(X_{(1)}), A)}_{\text{semi-offline RL}} \underbrace{p(A, G_A(X_{(1)}), Y)}_{\text{observed data}} \end{aligned}$$

which only involves observed data because  $A' \leq A$  due to the blocking.

Similar to online RL, we can sample from this distribution to construct a simulated data set  $\mathcal{D}'$ , consisting of  $X = G_A(X_{(1)}), Y, A, A'$ , and  $C' = f_C(A', G_{A'}(X_{(1)}), Y)$ . In Figure 6, we show the full causal graph of how such sampling is performed.

Note that although the number of observed data samples is limited by the data set size  $n$ , the semi-offline RL variables  $A'$  and  $C'$  can be sampled multiple times beyond this constraint. While  $\pi'_{sim}$  can be chosen to be a blocked AFA policy  $\pi'_\alpha$ , this is not required; different choices of  $\pi'_{sim}$  introduce an off-policy aspect. The only requirement is that  $\pi_{sim}$  (the unblocked version of  $\pi'_{sim}$ ) meets the positivity assumption from the offline RL view (Assumption 6.1).

Since we replace  $\pi_\alpha$  with  $\pi'_{sim}$ , the resulting cost samples  $C'$  cannot simply be averaged to estimate  $J$ . Instead, we reformulate the AFAPe problem as a causal inference problem by intervening on the semi-offline sampling distribution  $p'$ , reversing  $\pi'_{sim}$  back to  $\pi_\alpha$ . This reformulation is formalized as follows.

## 6.2 Problem Reformulation

The AFAPe problem can be reformulated under the semi-offline RL view (i.e., under the proposed distribution  $p'$ ) as the following theorem states.

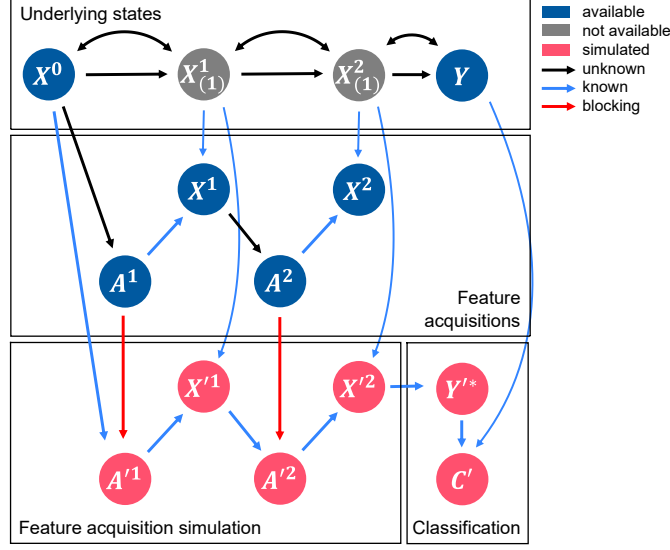


Figure 6: Causal graph for the semi-offline sampling distribution  $p'$ . Simulated acquisition actions  $A^t$  and observations  $X^t = G_{A^t}(X^t_{(1)})$  follow a blocked simulation policy  $\pi'_{sim}$ . The simulation policy is restricted by  $A^t$  such that actions  $A^t$  are blocked if  $A^t_i > A^t_i$  for any  $i$ . The simulated cost  $C'$  can be computed from  $A'$ ,  $X'$  and  $Y$ :  $C' = f_C(Y'^*, Y)$  with  $Y'^* = f_{cl}(X', A')$  being the predicted label under the simulated acquisitions. Edges showing long-term dependencies are omitted from the graph for visual clarity. These include  $\underline{X}^{t-1}, \underline{A}^{t-1} \rightarrow A^t$ ;  $\underline{X}^{t-1}, \underline{A}^{t-1} \rightarrow A^t$ ;  $\underline{X}^{tT}, \underline{A}^{tT} \rightarrow Y^*$ ;  $\underline{X}^{t-1} \leftrightarrow X^t_{(1)}$ ; and  $\underline{X}^T_{(1)} \leftrightarrow Y$ .

**Theorem 3.** (AFAPE problem reformulation under the semi-offline RL view). The AFape problem of estimating  $J$  (Eq. 2 or Eq. 9) is under Assumption 1 (no measurement noise), Assumption 2 (consistency), Assumption 3 (no interference), Assumption 4 (NDE) and Assumption 5 (NUC) equivalent to estimating

$$J = \mathbb{E}_{p'}[C'_{(\pi_\alpha)}]. \quad (17)$$

$C'_{(\pi_\alpha)}$  denotes the potential outcome of  $C'$ , had, instead of the blocked simulation policy  $\pi'_{sim}$ , the AFA policy  $\pi_\alpha$  been employed.

**Proof** Starting from Eq. 15, we find:

$$\begin{aligned} J &= \sum_{X_{(1)}, Y} \sum_{a \in \mathcal{A}} f_C(a, G_a(X_{(1)}), Y) \pi_\alpha(a | G_a(X_{(1)})) p(X_{(1)}, Y) \\ &= \sum_{X_{(1)}, Y, A_{(\pi_\alpha)}} f_C(A_{(\pi_\alpha)}, G_{A_{(\pi_\alpha)}}(X_{(1)}), Y) p(A_{(\pi_\alpha)}, X_{(1)}, Y) \\ &= \sum_{X_{(1)}, Y, A_{(\pi_\alpha)}, A} f_C(A_{(\pi_\alpha)}, G_{A_{(\pi_\alpha)}}(X_{(1)}), Y) p(A_{(\pi_\alpha)}, X_{(1)}, Y, A) \\ &\stackrel{*1}{=} \sum_{X_{(1)}, Y, A'_{(\pi_\alpha)}, A} f_C(A'_{(\pi_\alpha)}, G_{A'_{(\pi_\alpha)}}(X_{(1)}), Y) p'(A'_{(\pi_\alpha)}, X_{(1)}, Y, A) = \mathbb{E}_{p'}[C'_{(\pi_\alpha)}]. \end{aligned}$$

where the change in \*1) from  $A_{(\pi_\alpha)}$  to  $A'_{(\pi_\alpha)}$  is possible since an intervention on the simulated actions leads to the same distribution as an intervention on the observed actions. ■

**Remark 2 (Comparison of AFape under offline RL and semi-offline RL)** The AFape problem formulation under the semi-offline RL view (Eq. 17) closely resembles the original AFape formulation under the offline RL view (Eq. 2):  $J = \mathbb{E}[C_{(\pi_\alpha)}]$ . This similarity might raise the question of what is gained by this reformulation if one still needs to adjust for the intervention  $\pi_\alpha$ . The key difference is that, in the offline RL view, one must account for a significant distribution shift from  $\pi_\beta$  to  $\pi_\alpha$ . In contrast, under the semi-offline RL view, the distribution shift is much smaller, from a blocked  $\pi'_{sim}$  to  $\pi_\alpha$ .

### 6.3 Identification

In the following section, we focus on identifying the reformulated target  $J$  from Eq. 17. Similar to the other views, identification under the semi-offline RL view requires a positivity assumption. However, under the semi-offline RL view, this assumption is considerably less stringent compared to the positivity requirements in the offline RL (Assumption 6.1) and missing data views (Assumption 6.2).

In the offline RL view, the positivity assumption requires support for retrospective trajectories (i.e., action sequences under  $\pi_\beta$ ) that match any target trajectory (i.e., action sequences under  $\pi_\alpha$ ). The missing data view demands support for "acquire all" retrospective trajectories, meaning complete cases. In contrast, the semi-offline RL view only necessitates that for any target trajectory, there is positive support for at least one retrospective trajectory with equal or more acquisitions. We now formalize this positivity assumption rigorously. Due to its inherent complexity, we divide the formalization into multiple definitions:

**Definition 4.** (*Local positivity assumption and local admissible set  $\mathcal{A}_{adm}$  for semi-offline RL*). Let the local admissible set  $\mathcal{A}_{adm}^t(\underline{x}^{t-1}, \underline{a}^{t-1}, a^t)$ , defined for all  $a^t$  and all  $\underline{x}^{t-1}, \underline{a}^{t-1}$  s.t.  $p(G_{\underline{a}^{t-1}}(X_{(1)}) \equiv \underline{X}^{t-1} = \underline{x}^{t-1}, \underline{A}^{t-1} = \underline{a}^{t-1}) > 0$ , be the non-empty set containing all values of  $a^t$  for which

$$\begin{aligned} (1) \quad & a^t \geq a^{tt} \\ (2) \quad & \pi_\beta^t(a^t | \underline{x}^{t-1}, \underline{a}^{t-1}) \geq \mathcal{O} \end{aligned}$$

for some constant  $\mathcal{O} > 0$ . We further say that the local positivity assumption holds at  $\underline{x}^{t-1}, \underline{a}^{t-1}, a^{tt}$  if  $\mathcal{A}_{adm}^t(\underline{x}^{t-1}, \underline{a}^{t-1}, a^{tt})$  exists.

The local positivity assumption states that the observed data allows the simulation of a desired action  $a^{tt}$  (i.e., there is positive support for at least one value  $a^t$  s.t.  $a^t \geq a^{tt}$ ). The local positivity assumption is, however, not enough, which leads to the following definition of regional positivity.

**Definition 5.** (*Regional positivity assumption and regional admissible set  $\tilde{\mathcal{A}}_{adm}$  for semi-offline RL*). Let the regional admissible set  $\tilde{\mathcal{A}}_{adm}^t(\underline{x}^{t-1}, \underline{a}^{t-1}, \underline{a}^{tt}) \subseteq \mathcal{A}_{adm}^t(\underline{x}^{t-1}, \underline{a}^{t-1}, a^{tt})$ ,

defined for all  $\underline{x}^{t-1}, \underline{a}^{t-1}, \underline{a}^t$  s.t.  $p'(G_{\underline{a}^{t-1}}(X_{(1)}) = \underline{x}^{t-1}, \underline{A}^{t-1} = \underline{a}^{t-1}, \underline{A}^t = \underline{a}^t) > 0$ , be the non-empty set containing all values of  $a^t$  such that  $\tilde{\mathcal{A}}_{adm}^{t+1}(\underline{x}^t, \underline{a}^t, \underline{a}^{t+1})$  exists for all  $\underline{x}^t, \underline{a}^t, \underline{a}^{t+1}$  such that the following holds for  $a^{t+1}$ , and  $x^t$ :

$$p(x^t | \underline{x}^{t-1}, \underline{a}^t) \pi_\alpha^t(a^{t+1} | \underline{x}^{t-1}, \underline{a}^t) > 0 \quad (18)$$

where  $\underline{x}^{t-1}$  denotes the acquired features under  $\underline{a}^{t-1}$  and thus a subset of  $\underline{x}^{t-1}$ . We further say that the regional positivity assumption holds at  $\underline{x}^{t-1}, \underline{a}^{t-1}, \underline{a}^t$  if  $\tilde{\mathcal{A}}_{adm}^t(\underline{x}^{t-1}, \underline{a}^{t-1}, \underline{a}^t)$  exists.

Regional positivity states that there is not only a value  $a^t$  with positive support in the observed data that allows the simulation of a desired action  $a^t$  (i.e., local positivity), but it also ensures for such an  $a^t$ , that the simulations of all possible future desired actions are also possible. As regional positivity is still limited to a given data point, we also make the following global positivity assumption.

**Assumption 6.3 (Global positivity assumption for semi-offline RL)** We say that the global positivity assumption holds if the regional positivity assumption holds for all datapoints  $a^1, x^0$  s.t.  $p(x^0) \pi_\alpha(a^1 | x^0) > 0$ .

The following lemma (proved in Appendix F) establishes that the global positivity assumption for semi-offline RL is weaker than both the positivity assumptions required by the offline RL and the missing data views:

**Lemma 6.** (Sufficiency conditions for global positivity). *The global positivity assumption for semi-offline RL (Assumption 6.3) holds if the positivity assumption from offline RL (Assumption 6.1) or from missing data (Assumption 6.2) holds.*

After having defined positivity for semi-offline RL, we can now perform identification for  $J$ :

**Theorem 7.** (Identification of  $J$  for the semi-offline RL view). *The reformulated AFAPF problem of estimating  $J$  under the semi-offline RL view (Eq. 17) is under Assumption 1 (no measurement noise), Assumption 2 (consistency), Assumption 3 (no interference), Assumption 4 (NDE), Assumption 5 (NUC) and Assumption 6.3 (global positivity) identified by*

$$J = \mathbb{E}_{p'}[C'_{(\pi_\alpha)}] = \sum_{A', A, G_A(X_{(1)}), Y} f_C(A', X', Y) q'(A', A, X, Y) \quad (19)$$

with the distribution

$$q'(A', A, X, Y) = \prod_{t=1}^T \underbrace{\pi_{id}^t(A^t | \underline{A}^t, \underline{X}^{t-1}, \underline{A}^{t-1})}_{\text{distr. subject to constraints}} \underbrace{\pi_\alpha^t(A^t | \underline{X}^{t-1}, \underline{A}^{t-1})}_{\text{target policy}} \prod_{t=0}^T p(X^t | \underline{X}^{t-1}, \underline{A}^t, Y) p(Y) \quad (20)$$

where

$$\pi_{id}^t(A^t | \underline{A}^t, \underline{X}^{t-1}, \underline{A}^{t-1}) = \mathbb{I}(A^t \in \underbrace{\tilde{\mathcal{A}}_{adm}^t(\underline{X}^{t-1}, \underline{A}^{t-1}, \underline{A}^t)}_{\text{support restriction}}) f_{id}^t(A^t, \underline{X}^{t-1}, \underline{A}^{t-1}) \quad (21)$$

for any function  $f_{id}^t$  s.t.  $\pi_{id}^t$  is a valid density.



Here, we used again  $X \equiv G_A(X_{(1)})$  and  $X' \equiv G_{A'}(X_{(1)})$  to simplify notation. In fact, since  $A' \leq A$ , we can also write  $X' = G_{A'}(X)$ . Additionally, we can define semi-offline RL versions of the Bellman equation:

**Theorem 8.** *(Bellman equation for semi-offline RL). The semi-offline RL view admits under Assumption 1 (no measurement noise), Assumption 2 (consistency), Assumption 3 (no interference), Assumption 4 (NDE), Assumption 5 (NUC) and the local positivity assumption at datapoint  $\underline{x}^{t-1}, \underline{a}^{t-1}, a^t$  (from Definition 4), the following semi-offline RL version of the Bellman equation:*

$$Q_{Semi}(\underline{A}^t, \underline{X}^{t-1}, \underline{A}^{t-1}, \Xi) = \sum_{X^t} V_{Semi}(\underline{A}^t, \underline{X}^t, \underline{A}^{t-1}, A^t = a^t, \Xi) p(X^t | \underline{X}^{t-1}, \underline{A}^{t-1}, A^t = a^t, \Xi) \quad (22)$$

for any  $a^t \in \mathcal{A}_{adm}^t(\underline{X}^{t-1}, \underline{A}^{t-1}, A^t)$

$$V_{Semi}(\underline{A}^t, \underline{X}^t, \underline{A}^t, \Xi) = \sum_{A^{t+1}} Q_{Semi}(\underline{A}^{t+1}, \underline{X}^t, \underline{A}^t, \Xi) \pi_\alpha^{t+1}(A^{t+1} | \underline{X}^t, \underline{A}^t) \quad (23)$$

with semi-offline RL versions of the state-action value function  $Q_{Semi}$  and state value function  $V_{Semi}$ :

$$\begin{aligned} Q_{Semi}^t &\equiv Q_{Semi}(\underline{A}^t, \underline{X}^{t-1}, \underline{A}^{t-1}, \Xi) \equiv \mathbb{E}_{p'}[C'_{(\bar{\pi}_\alpha^{t+1})} | \underline{A}^t, \underline{X}^{t-1}, \underline{A}^{t-1}, \Xi] \\ V_{Semi}^t &\equiv V_{Semi}(\underline{A}^t, \underline{X}^t, \underline{A}^t, \Xi) \equiv \mathbb{E}_{p'}[C'_{(\bar{\pi}_\alpha^{t+1})} | \underline{A}^t, \underline{X}^t, \underline{A}^t, \Xi] \end{aligned}$$

where  $C'_{(\bar{\pi}_\alpha^{t+1})}$  denotes the potential outcome of  $C'$  under interventions from time step  $t+1$  onwards.  $\Xi \subseteq \{Y, O\}$ , with  $O$  denoting all features that are always available, denotes an optional subset of additional variables that can be conditioned on. Furthermore,  $Q_{Semi}^t$  and  $V_{Semi}^t$  are identified if the regional positivity assumption (from Definition 5) holds at  $\underline{X}^{t-1}, \underline{A}^{t-1}, \underline{A}^t$  and  $a^t \in \tilde{\mathcal{A}}_{adm}^t(\underline{X}^{t-1}, \underline{A}^{t-1}, \underline{A}^t)$ .

The proofs for Theorems 7 and 8 are shown in Appendix G.

The functions  $Q_{Semi}$  and  $V_{Semi}$  are very similar to their counterparts from the offline RL view ( $Q_{Off}$  and  $V_{Off}$ ), with two differences: i) they are learned from a curated data set  $\mathcal{D}'$  which arises from sampling  $p'$ ; and ii) they contain the simulated acquisitions  $\underline{A}^t$ , but also the real features and actions  $\underline{X}^t, \underline{A}^t$  which are needed to adjust for confounding of the blocking operation.

The identification steps so far have been very specific to knowledge about  $\tilde{\mathcal{A}}_{adm}^t$  that needs to be assessed from the data. We now look more closely at a specific, stronger positivity assumption (where  $\mathcal{A}_{adm} = \tilde{\mathcal{A}}_{adm}$ ) that allows the use of the maximum amount of datapoints and therefore leads to easier to use estimators.

**Definition 9.** *(Maximal regional positivity assumption for semi-offline RL): We say that the maximal regional positivity assumption holds for a datapoint  $\underline{x}^{t-1}, \underline{a}^{t-1}, \underline{a}^t$  if  $\tilde{\mathcal{A}}_{adm}^t(\underline{x}^{t-1}, \underline{a}^{t-1}, \underline{a}^t) = \mathcal{A}_{adm}^t(\underline{x}^{t-1}, \underline{a}^{t-1}, a^t)$  and the maximal regional positivity assumption further holds for all  $\underline{x}^t, \underline{a}^t, \underline{a}^{t+1}$  such that  $a^t \in \mathcal{A}_{adm}^t(\underline{x}^{t-1}, \underline{a}^{t-1}, a^t)$  and the following holds for  $a^{t+1}$ , and  $x^t$ :*

$$p(x^t | \underline{x}^{t-1}, \underline{a}^t) \pi_\alpha(a^{t+1} | \underline{x}^t, \underline{a}^t) > 0.$$

**Assumption 6.4 (Maximal global positivity assumption for semi-offline RL)** We say that the maximal global positivity assumption holds if the maximal regional positivity assumption holds for all datapoints  $x^0, a^1$  s.t.  $p(x^0)\pi_\alpha(a^1|\underline{x}^0) > 0$ .

The maximal regional positivity and maximal global positivity assumptions ensure that we can use all available data points where  $A^t \geq A^{t-1}$  without running into positivity problems in later time steps. This makes the identification and estimation significantly easier, as shown next.

We can now propose the following corollary of Theorem 7, which states identification under the maximal global positivity assumption.

**Corollary 10.** *(Identification of  $J$  for the semi-offline RL view under maximal global positivity). The reformulated AFAPPE problem of estimating  $J$  under the semi-offline RL view (Eq. 17) is under Assumption 1 (no measurement noise), Assumption 2 (consistency), Assumption 3 (no interference), Assumption 4 (NDE), Assumption 5 (NUC) and Assumption 6.4 (maximum global positivity) identified by Eqs. 19 and 20 where*

$$\begin{aligned} \pi_{id}^t(A^t|\underline{A}^{t-1}, A^t = a^t, \underline{X}^{t-1}, \underline{A}^{t-1}) &= \\ &= \mathbb{I}(A^t \geq a^t)\pi_\beta^t(A^t|\underline{X}^{t-1}, \underline{A}^{t-1})f_{id}^t(\underline{A}^{t-1}, A^t = a^t, \underline{X}^{t-1}, \underline{A}^{t-1}) \end{aligned}$$

for any function  $f_{id}^t$  s.t.  $\pi_{id}^t$  is a valid density. This holds in particular for the choice of a truncated  $\pi_\beta$ :

$$\begin{aligned} \pi_{id}^t(A^t|\underline{A}^{t-1}, A^t = a^t, \underline{X}^{t-1}, \underline{A}^{t-1}) &= \pi_\beta^t(A^t|A^t \geq a^t, \underline{X}^{t-1}, \underline{A}^{t-1}) = \\ &= \frac{\mathbb{I}(A^t \geq a^t)\pi_\beta^t(A^t|\underline{X}^{t-1}, \underline{A}^{t-1})}{\pi_\beta^t(A^t \geq a^t|\underline{X}^{t-1}, \underline{A}^{t-1})}. \end{aligned}$$

The proof for Corollary 10 is shown in Appendix H.

Lastly, we also provide the following Remark that brings the factorization of the observational distribution  $p'$  from Eq. 16 into a comparable form to the identifying distribution  $q'$ .

**Remark 3 (Factorization of the semi-offline sampling distribution)** The semi-offline sampling distribution  $p'$  can be alternatively written as:

$$\begin{aligned} p'(A', G_A(X_{(1)}), Y, A) &= \prod_{t=1}^T \pi_{sim}^{t'}(A^t|G_{\underline{A}^{t-1}}(X_{(1)}), \underline{A}^{t-1}, A^t)p(A, G_A(X_{(1)}), Y) \\ &= \prod_{t=1}^T \underbrace{\pi_{sim}^{t'}(A^t|\underline{X}^{t-1}, \underline{A}^{t-1}, A^t)}_{\text{known simulation policy}} \underbrace{\pi_\beta(A^t|\underline{X}^{t-1}, \underline{A}^{t-1})}_{\text{retro. acquisition policy}} \prod_{t=0}^T p(X^t|\underline{X}^{t-1}, \underline{A}^t, Y)p(Y) \quad (24) \end{aligned}$$

## 6.4 Estimation

We propose the following novel estimators for  $J$ , which arise from the semi-offline RL viewpoint. We differentiate between estimators derived under the global positivity assumption (Assumption 6.3) and under the (stronger) maximal global positivity assumption (Assumption 6.4).

1) *Inverse probability weighting (IPW):*

The target cost that is estimated by the semi-offline IPW estimator is

$$\hat{J}_{IPW-Semi} = \hat{\mathbb{E}}_n \left[ \hat{\mathbb{E}}_{n'} \left[ \hat{\rho}_{Semi}^T C' | A, X, Y \right] \right]. \quad (25)$$

The inner expectation, denoted by  $\hat{\mathbb{E}}_{n'}[\cdot]$ , represents the empirical average over the simulated values of  $A'$ , which can involve many more samples compared to the outer expectation, taken over the observed data. The inverse probability weights are under the global positivity assumption:

$$\hat{\rho}_{Semi}^T = \hat{\rho}_{Semi}^T(\pi_{id}) = \prod_{t=1}^T \frac{\pi_{\alpha}^t(A^t | \underline{X}^{t-1}, \underline{A}^{t-1})}{\pi_{sim}^{tt}(A^t | \underline{X}^{t-1}, \underline{A}^{t-1}, A^t)} \frac{\pi_{id}^t(A^t | \underline{A}^t, \underline{X}^{t-1}, \underline{A}^{t-1})}{\hat{\pi}_{\beta}^t(A^t | \underline{X}^{t-1}, \underline{A}^{t-1})} \quad (26)$$

or under the maximal global positivity assumption (by choosing  $\pi_{id}^t = \pi_{\beta}^t(A^t | A^t \geq a^t, \underline{X}^{t-1}, \underline{A}^{t-1})$ ):

$$\hat{\rho}_{Semi}^T = \prod_{t=1}^T \frac{\pi_{\alpha}^t(A^t | \underline{X}^{t-1}, \underline{A}^{t-1})}{\pi_{sim}^{tt}(A^t | \underline{X}^{t-1}, \underline{A}^{t-1}, A^t)} \frac{\mathbb{I}(A^t \geq a^t)}{\hat{\pi}_{\beta}^t(A^t \geq a^t | \underline{X}^{t-1}, \underline{A}^{t-1})}. \quad (27)$$

The following remarks state that the IPW estimators from the offline RL and missing data viewpoints are special cases of the proposed estimator:

**Remark 4 (Offline RL IPW estimator as a special version of the semi-offline RL IPW estimator)** The IPW estimator from the offline RL view,  $\hat{J}_{IPW-Off}$ , is, for the choice  $\pi'_{sim} = \pi'_{\alpha}$ , equal to  $\hat{J}_{IPW-Semi}$  with:

$$\pi_{id}^t(A^t | \underline{A}^{t-1}, A^t = a^t, \underline{X}^{t-1}, \underline{A}^{t-1}) = \mathbb{I}(A^t = a^t).$$

**Remark 5 (Missing data IPW estimator as a special version of the semi-offline RL IPW estimator)** The IPW estimator from the missing data view,  $\hat{J}_{IPW-Miss}$ , is, for the choice  $\pi'_{sim} = \pi'_{\alpha}$ , equal to  $\hat{J}_{IPW-Semi}$  with:

$$\pi_{id}(A^t | \underline{A}^t, \underline{X}^{t-1}, \underline{A}^{t-1}) = \mathbb{I}(A^t = \bar{1}).$$

The IPW estimator  $\hat{J}_{IPW-Semi}$  under the maximal global positivity assumption demonstrates the large benefits of the semi-offline RL view over both the offline RL and missing data views. Its second fraction shows that not only datapoints where  $A^t = A^t$  are used (i.e., have positive weight), as in the offline RL view, neither only datapoints where  $A^t = \bar{1}$  are used, as in the missing data view, but all datapoints where  $A^t \geq A^t$  can be used. However, this benefit diminishes as the target policy becomes more "data-hungry"—that is, as it acquires more features. In fact, there is no difference between all three estimators for an "acquire all features" policy:

**Remark 6 (Equality of IPW estimators for an "acquire all features" policy)** The IPW estimators from the missing data view,  $\hat{J}_{IPW-Miss}$ , from the offline RL view  $\hat{J}_{IPW-Off}$ , and from the semi-offline RL view  $\hat{J}_{IPW-Semi}$  are identical if  $\pi_{\alpha}^t(A^t = \bar{1} | \underline{X}^{t-1} = \underline{x}^{t-1}, \underline{A}^{t-1} = \bar{1}) = 1 \forall t, \underline{x}^{t-1}$ .

We show in Appendix I that  $\hat{J}_{IPW-Semi}$  (under maximal global positivity) is equivalent in simple AFA settings to an adapted version of the IPW estimator by Caniglia et al. (2019). Our IPW estimators  $\hat{J}_{IPW-Semi}$  can, however, be applied in more general AFA settings.

2) *Direct method (DM):*

The target cost that is estimated by the semi-offline DM estimator is

$$\hat{J}_{DM-Semi} = \hat{\mathbb{E}}_n[\hat{V}_{Semi}^0] \quad (28)$$

This estimator is based on learning a semi-offline RL version of the state-action value function  $Q_{Semi}$  using the semi-offline version of the Bellman equation (Eqs. 22 and 23). Using  $Q_{Semi}$ , one can compute the state value function:  $V_{Semi}^t = \mathbb{E}_{\pi_\alpha}[Q_{Semi}^{t+1}]$ .

The training process of  $Q_{Semi}$  can benefit from using the off-policy aspect of the proposed semi-offline sampling distribution  $p'$  (i.e., from using a simulation policy  $\pi'_{sim}$  that is different from  $\pi'_\alpha$ ). This is because a deterministic AFA policy, for example, will only generate one exact trajectory of simulated actions  $A'$  and costs  $C'$  per datapoint  $X, Y, A$ . A stochastic simulation policy  $\pi'_{sim}$  can instead be used to generate multiple such trajectories. Usually, a parametric working model, for example a multi-layer perceptron (MLP), is assumed for the nuisance function  $\hat{Q}_{Semi}$ . The estimation under such a working model will benefit from the additional datapoints generated under a policy  $\pi'_{sim} \neq \pi'_\alpha$ .

3) *Double reinforcement learning (DRL):*

The target cost that is estimated by the semi-offline DRL estimator is

$$\hat{J}_{DRL-Semi} = \hat{\mathbb{E}}_n \left[ \hat{\mathbb{E}}_{n'} \left[ \hat{\rho}_{Semi}^T C' + \sum_{t=1}^T \left( -\hat{\rho}_{Semi}^t \hat{Q}_{Semi}^t + \hat{\rho}_{Semi}^{t-1} \hat{V}_{Semi}^{t-1} \right) \middle| A, X, Y \right] \right]. \quad (29)$$

which holds for both choices for  $\rho_{Semi}$ , given that the respective positivity assumption holds. Similar to the DRL estimator from the offline RL view, this approach combines the other two estimators (Eqs. 25 and 28).

The following theorems state some notable facts about these estimators.

**Theorem 11.** (*Consistency of  $\hat{J}_{IPW-Semi}$* ). *The estimator  $\hat{J}_{IPW-Semi}$  is consistent if the propensity score model  $\hat{\pi}_\beta$  is correctly specified.*

**Proof** We apply the standard inverse probability weighting approach  $\mathbb{E}_{q'}[C'] = \mathbb{E}_{p'}[\frac{q'}{p'}C']$  and use the factorizations for  $q'$  and  $p'$  from Eqs. 20 (Theorem 7) and 24 (Remark 3), respectively, to obtain Eq. 26 for the weights. The weights from Eq. 27 arise from inserting the special choice for  $\pi_{id}$ . ■

**Theorem 12.** (*Consistency of  $\hat{J}_{DM-Semi}$* ). *The estimator  $\hat{J}_{DM-Semi}$  is consistent if the  $Q$ -function  $\hat{Q}_{Semi}$  is correctly specified.*

**Proof** The proof of the consistency of  $\hat{J}_{DM-Semi}$  follows simply from the semi-offline Bellman equation (Theorem 8) and the law of total expectation. ■

**Theorem 13.** (*Double robustness of  $\hat{J}_{DRL-Semi}$* ). *The estimator  $\hat{J}_{DRL-Semi}$  is doubly robust, in the sense that it is consistent if either the  $Q$ -function  $\hat{Q}_{Semi}$  or the propensity score model  $\hat{\pi}_\beta$  is correctly specified.*

The proof is shown in Appendix J. The estimator  $\hat{J}_{DRL-Semi}$  is a 1-step estimator based on an influence function derived for  $J$  under  $p'$ . Therefore, the DRL estimator is regular and asymptotically linear (RAL). The influence function is given by the following theorem:

**Theorem 14.** (*An influence function under the semi-offline RL view*). *An influence function of  $J$  is:*

$$\varphi_{Semi} = -J + \mathbb{E} \left[ \rho_{Semi}^T C' + \sum_{t=1}^T (-\rho_{Semi}^t Q_{Semi}^t + \rho_{Semi}^{t-1} V_{Semi}^{t-1}) \mid A, X, Y \right]. \quad (30)$$

We will prove this theorem in the next section, addressing semiparametric estimation from all three views.

**Remark 7 (Efficiency of  $\varphi_{Semi}$  as a function of  $\Xi$ )** In Theorem 8, we showed that the semi-offline RL version of the Bellman equation holds for any subset  $\Xi \subseteq \{Y, O\}$ , and the same applies to Theorem 14. The choice of  $\Xi$  affects the efficiency of the DRL estimator: a bigger set  $\Xi$  corresponds to a higher efficiency of the corresponding DRL estimator. However, the class of influence functions presented here does not include the efficient influence function. In fact, as we will discuss in the next section, the efficient influence function lacks a closed-form expression.

In Appendix K, we extend the estimators discussed in this section to other settings. These include: i) the estimation of  $J_a$ , and ii) scenarios where a prediction  $Y^{*t}$  is required at each time step  $t$ .

## 7. Semiparametric Theory under NUC and NDE

*Assumptions in this section: Assumption 4 (NDE), Assumption 5 (NUC)*

In this section, we explore semiparametric estimation approaches for  $J$  under both NDE and NUC assumptions. Readers not interested in the detailed semiparametric theory can skip this section. We demonstrate that all three views can be unified within an established semiparametric theory framework for MAR missing data problems. Using this framework, we prove Theorem 14, which defines a class of influence functions derived under the semi-offline RL view. Although no closed-form efficient influence function exists in this setting, efficiency improvements from Tsiatis (2006) and Liu et al. (2021) can be applied. However, these methods pose significant challenges, including strong positivity assumptions, applicability to a limited set of problems, and implementation complexity.

To discuss semiparametric approaches to AFAPE, we first remind the reader of the semiparametric theory review in Appendix B. The following section draws extensively on the foundational framework of Tsiatis (2006) to establish essential context.

Let the observed data influence function be denoted as  $\varphi \equiv \varphi(A, G_A(X_{(1)}), Y)$ , and the observed data tangent space as  $\Lambda$ , with the corresponding nuisance tangent space denoted by  $\Lambda_{nuis}$ . The full data influence function - an influence function given the counterfactual

variables  $X_{(1)}$  - is denoted  $\varphi^F \equiv \varphi^F(X_{(1)}, Y)$ , with full data tangent space  $\Lambda^F$  and nuisance tangent space  $\Lambda_{nuis}^F$ . The two key relationships between these spaces and the influence functions are as follows:

- An observed data influence function must lie in the orthocomplement of the observed data nuisance tangent space:  $\varphi \in \Lambda_{nuis}^\perp$ .
- The observed data efficient influence function must be in the observed data tangent space:  $\varphi_{eff} \in \Lambda$ .

Known semiparametric theory for MAR missing data methods can be applied to the AFAPe problem under the NUC assumption. We begin by defining the space of full data influence functions. Since we assume no restrictions on the full data, the full data influence function is unique, efficient, and given by the online RL part of the missing data + online RL view:

$$\varphi^F(X_{(1)}, Y) = \underbrace{\mathbb{E}[C_{(\pi_\alpha)} | X_{(1)}, Y]}_{\text{online RL}} - J = \sum_{a \in \mathcal{A}} f_C(a, G_a(X_{(1)}, Y)) \pi_\alpha(a | G_a(X_{(1)})) - J. \quad (31)$$

While the full data influence function is unique, the space of observed data influence functions is generally not. To find it, we first define the orthogonal complement of the nuisance tangent space, which can be expressed as (Theorem 8.3 from Tsiatis (2006)):

$$\Lambda_{nuis}^\perp = \left\{ [h^*(A, G_A(X_{(1)}, Y) \oplus \Lambda_2) - \Pi \left( [h^*(A, G_A(X_{(1)}, Y) \oplus \Lambda_2) \mid \Lambda_{nuis, \psi} \right)] \right\}, \quad (32)$$

where  $h^*$  belongs to the inverse probability weighting (IPW) space  $\Lambda_{IPW}^*$ ,  $\Lambda_2$  is the augmentation space, and  $\Lambda_{nuis, \psi} = \Lambda_\psi$  is the nuisance tangent space, or equivalently the tangent space of the acquisition process. We explain these three spaces now in more detail.

**The inverse probability weighting space  $\Lambda_{IPW}^*$ :**

The function  $h^*(A, G_A(X_{(1)}, Y))$  in Eq. 32 can be any function in the IPW space  $\Lambda_{IPW}^*$  which is defined as:

$$\Lambda_{IPW}^* \equiv \left\{ h^*(A, G_A(X_{(1)}, Y)) \in \mathcal{H} : \mathbb{E}[h^*(A, G_A(X_{(1)}, Y)) | X_{(1)}, Y] = \varphi^{*F}(X_{(1)}, Y) \right\}.$$

where  $\mathcal{H}$  denotes the space of random functions with zero mean and finite variance and  $\varphi^{*F}(X_{(1)}, Y)$  denotes an element of the orthocomp of the full data nuisance tangent space:  $\varphi^{*F} \in \Lambda_{nuis}^\perp$ .

In fact, if we further restrict the IPW space such that we don't allow any element of the orthocomp of the full data nuisance tangent space  $\varphi^{*F}(X_{(1)}, Y)$ , but only the full data influence function  $\varphi^F(X_{(1)}, Y)$ , then we also obtain only observed data influence functions by Eq. 32 (Theorem 8.3 from Tsiatis (2006)). In the following, we thus restrict the IPW space to:

$$\Lambda_{IPW} \equiv \left\{ h(A, G_A(X_{(1)}, Y)) \in \mathcal{H} : \mathbb{E}[h(A, G_A(X_{(1)}, Y)) | X_{(1)}, Y] = \varphi^F(X_{(1)}, Y) \right\}.$$

In that case, the IPW space contains functions that, when taking the conditional expected value with respect to the full data, equal the full data influence function. The space is denoted as the IPW space because IPW-based identifying functions can be chosen to construct

elements in this space. In fact, as will be shown in this section, all the IPW estimators - from the offline RL, missing data, and semioffline RL views - are applicable and form valid elements in  $\Lambda_{IPW}$ .

**The augmentation space  $\Lambda_2$ :**

The augmentation space  $\Lambda_2$  is defined as follows (Lemma 7.4 from Tsiatis (2006)):

$$\Lambda_2 = \left\{ b(A, G_A(X_{(1)}), Y) \in \mathcal{H} : \mathbb{E} [b(A, G_A(X_{(1)}), Y) | X_{(1)}, Y] = 0 \right\}. \quad (33)$$

This space contains functions that, when taken in conditional expectation with respect to the full data, equal zero. This provides intuition behind the decomposition of the space of observed data influence functions: the space consists of one function that, in conditional expectation, equals the full data influence function, plus all functions that become zero in conditional expectation. One must, however, still obtain the residual of the projection of these functions onto the nuisance tangent space of the acquisition process,  $\Lambda_{nuis, \psi}$ , introduced below.

**The nuisance tangent space of the acquisition process  $\Lambda_{nuis, \psi}$ :**

The space  $\Lambda_{nuis, \psi}$  corresponds to the observed data nuisance tangent space of the acquisition process and is a subspace of  $\Lambda_2$  (Theorem 8.1 from Tsiatis (2006)). In our AFA setting, future feature values do not influence past acquisition decisions, resulting in the following conditional independences:  $A^t \perp\!\!\!\perp \bar{X}_{(1)}^t, Y | G_{\underline{A}^{t-1}}(\bar{X}_{(1)}^{t-1}), \underline{A}^{t-1}$  which can be translated into tangent space restrictions such that:

$$\Lambda_{nuis, \psi} = \Lambda_{nuis, \psi}^1 \oplus \Lambda_{nuis, \psi}^2 \oplus \dots \oplus \Lambda_{nuis, \psi}^T \quad (34)$$

where each subspace  $\Lambda_{nuis, \psi}^t$  is defined as:

$$\Lambda_{nuis, \psi}^t \equiv \left\{ \gamma^t(A^t, \underline{A}^{t-1}, G_{\underline{A}^{t-1}}(X_{(1)})) \in \mathcal{H} : \mathbb{E} \left[ \gamma^t(A^t, \underline{A}^{t-1}, G_{\underline{A}^{t-1}}(X_{(1)})) | \underline{A}^{t-1}, G_{\underline{A}^{t-1}}(X_{(1)}) \right] = 0 \right\}.$$

As these subspaces are orthogonal, projections onto them are available in closed form and are known to be (as derived in Appendix B):

$$\Pi([\cdot] | \Lambda_{nuis, \psi}^t) = \mathbb{E} \left[ [\cdot] | A^t, \underline{A}^{t-1}, G_{\underline{A}^{t-1}}(X_{(1)}) \right] - \mathbb{E} \left[ [\cdot] | \underline{A}^{t-1}, G_{\underline{A}^{t-1}}(X_{(1)}) \right].$$

Amongst the class of observed data influence functions, one may further be interested in finding the one with the smallest asymptotic variance, i.e., the efficient observed data influence function  $\varphi_{eff}$ . It can be found via the following projection of  $h$  (Theorem 10.1 from (Tsiatis, 2006)):

$$\varphi_{eff} = h(A, G_A(X_{(1)}), Y) - \Pi(h(A, G_A(X_{(1)}), Y) | \Lambda_2).$$

Hence, to construct an efficient influence function, one needs to find an element  $h$ , construct the space  $\Lambda_2$ , and project onto it. We now embed the three viewpoints of this work—missing data view, offline RL view, and semi-offline RL view—into this framework. We begin with the traditional missing data approach based on complete cases.

### 7.1 Missing data view

Now, we discuss the standard, traditional missing data approach to choosing an element  $h$  of the IPW space  $\Lambda_{IPW}$  and constructing the augmentation space  $\Lambda_2$ . Traditional semiparametric estimators for missing data problems rely on the missing data positivity assumption (Assumption 6.2), which requires complete cases. The corresponding estimators are referred to as augmented inverse probability weighting complete case (AIPWCC) estimators.

These estimators are called complete case estimators because they choose the missing data IPW estimator for  $h$  (see Tsiatis (2006) for more details):

$$h_{Miss}(A, G_A(X_{(1)}), Y) = \rho_{Miss} \mathbb{E}[C_{(\pi_\alpha)} | X_{(1)}, Y] - J \in \Lambda_{IPW}. \quad (35)$$

Furthermore, they also require the missing data positivity assumption (Assumption 6.2) for the construction of  $\Lambda_2$ , given as:

$$\Lambda_2 = \left\{ \sum_{a \in \mathcal{A} \setminus \bar{1}} \left[ \mathbb{I}(A = a) - \prod_{t=1}^T \frac{\mathbb{I}(A^t = \bar{1}) \pi_\beta^t(A^t = a^t | G_{\underline{a}^{t-1}}(X_{(1)}), \underline{a}^{t-1})}{\pi_\beta^t(A^t = \bar{1} | G_{\underline{A}^{t-1}}(X_{(1)}), \underline{A}^{t-1} = \bar{1})} \right] b_a(G_A(X_{(1)}), Y) : b_a(G_A(X_{(1)}), Y) \in \mathcal{H} \right\}. \quad (36)$$

For the interested reader, we show how both of these choices are derived under the missing data positivity assumption in Appendix L.

In addition to the strong positivity requirements, a key challenge is that general projections onto  $\Lambda_2$  are not available in closed form (Tsiatis, 2006). Alternatives to still perform such projections and achieve full efficiency include iterative numerical methods. However, these methods are difficult to implement and involve significant computational challenges, which have hindered their practical application (Tsiatis, 2006). Therefore, such methods are beyond the scope of this work, but we direct interested readers to Tsiatis (2006) for further details.

### 7.2 Offline RL view

An alternative to applying traditional semiparametric theory for missing data problems is to start from the offline RL view. In this approach, one projects the influence function  $\varphi_{Off}$ —associated with the DRL estimator—onto the restricted tangent space under the NDE assumption, denoted by  $\Lambda_{NDE}$ :

$$\varphi_{eff} = \Pi(\varphi_{Off} | \Lambda_{NDE}) = \varphi_{Off} - \Pi(\varphi_{Off} | \Lambda_{NDE}^\perp).$$

This approach was first introduced by Liu et al. (2021) in the context of dynamic testing and treatment regimes, though it can be adapted to the AFAPPE setting. However, their derivation was limited to a single acquirable feature per time step ( $A^t \in \{0, 1\}$ ).

The space  $\Lambda_{NDE}^\perp$  is adapted to AFAPPE, given as:

$$\Lambda_{NDE}^\perp = \Lambda_* - \Pi(\Lambda_* | \Lambda_\psi)$$



with

$$\Lambda_* \equiv \left\{ b_{\underline{i}^{t-1}, \bar{i}^{t+1}}^t(\underline{A}^{t-1}, G_{\underline{A}^{t-1}}(X_{(1)}), X_{(1)}^{\underline{i}^{t-1}}, X_{(1)}^{\bar{i}^{t+1}}, Y) \left( \frac{A^t}{\pi_\beta^t} - 1 \right) \prod_{t'=1}^{t-1} \left( \frac{A^{t'}}{\pi_\beta^{t'}} \right)^{i^{t'}} \prod_{t'=t+1}^T \left( \frac{A^{t'}}{\pi_\beta^{t'}} \right)^{i^{t'}} : \right. \\ \left. b_{\underline{i}^{t-1}, \bar{i}^{t+1}}^t(\underline{A}^{t-1}, G_{\underline{A}^{t-1}}(X_{(1)}), X_{(1)}, Y) \in \mathcal{H} \right\}$$

where  $\underline{i}^{t-1} \equiv \{i^1, \dots, i^{t-1}\}$  and  $\bar{i}^{t+1} \equiv \{i^{t+1}, \dots, i^T\}$  index subsets of  $\{0, 1\}^{t-1}$  and  $\{0, 1\}^{T-t}$  respectively and  $X_{(1)}^{\underline{i}^{t-1}} \equiv \{X_{(1)}^{t'} : 1 \leq t' \leq t-1, i^{t'} = 1\}$  and  $X_{(1)}^{\bar{i}^{t+1}} \equiv \{X_{(1)}^{t'} : t+1 \leq t' \leq T, i^{t'} = 1\}$ . Furthermore, we denote  $\pi_\beta^t \equiv \pi_\beta^t(A^t = 1 | G_{\underline{A}^{t-1}}(X_{(1)}), \underline{A}^{t-1})$  for this one acquisition per time-point setting.

The space  $\Lambda_*$  includes a term for each combination of observed features (indexed by the subsets  $i$ ). This construction also relies on a stringent positivity assumption, which is even stronger than what is required for identification under the offline RL view (Assumption 6.1). Specifically, it demands that  $\pi_\beta^t > 0$  for all  $t$ ,  $G_{\underline{A}^{t-1}}(X_{(1)})$ , and  $\underline{A}^{t-1}$ , irrespective of the target policy. Furthermore, since the target parameter doesn't depend on the acquisition process, we have  $\Lambda_\psi = \Lambda_{\text{nuis}, \psi}$ .

In the following lemmas, we demonstrate the equivalence between this offline RL approach and the semiparametric theory for MAR missing data problems:

**Lemma 15.** *(Relating the offline RL IPW estimator to the IPW space). The functional  $h_{\text{off}} \equiv \rho_{\text{off}} C - J$ , based on the IPW estimator from the offline RL view, is a valid element of the IPW space:  $h_{\text{off}} \in \Lambda_{\text{IPW}}$ .*

**Lemma 16.** *( $\Lambda_*$  is equal to the augmentation space). The augmentation space  $\Lambda_2$  is equal to  $\Lambda_*$ .*

Both lemmas are proven in Appendix M.

We now demonstrate that both approaches—whether derived from the offline RL or the missing data view—yield the same influence function (as expected):

$$\begin{aligned} \varphi_{\text{eff}} &= \varphi_{\text{off}} - \Pi(\varphi_{\text{off}} | \Lambda_{\text{NDE}}^\perp) \\ &\stackrel{*1}{=} h_{\text{off}} - \Pi(h_{\text{off}} | \Lambda_{\text{nuis}, \psi}) - \Pi(h_{\text{off}} - \Pi(h_{\text{off}} | \Lambda_{\text{nuis}, \psi}) | \Lambda_* - \Pi(\Lambda_* | \Lambda_{\text{nuis}, \psi})) \\ &= h_{\text{off}} - \Pi(h_{\text{off}} | \Lambda_{\text{nuis}, \psi}) - \Pi(h_{\text{off}} | \Lambda_2 - \Pi(\Lambda_2 | \Lambda_{\text{nuis}, \psi})) \\ &\stackrel{*2}{=} h_{\text{off}} - \Pi(h_{\text{off}} | \Lambda_2) \end{aligned}$$

where we use in \*1) that the influence function can be decomposed:  $\varphi_{\text{off}} = h_{\text{off}} - \Pi(h_{\text{off}} | \Lambda_{\text{nuis}, \psi})$ . In \*2), we used that  $\Lambda_{\text{nuis}, \psi} \subset \Lambda_2$ .

This shows that a separate projection onto  $\Lambda_{\text{nuis}, \psi}$  is unnecessary. However, as noted earlier, a projection onto  $\Lambda_{\text{NDE}}^\perp$  (or  $\Lambda_2$ ) is not available in closed form (Liu et al., 2021), as previously discussed in the missing data approach. Liu et al. (2021) suggest instead constructing an arbitrarily large subspace  $\Omega$  onto which a projection is feasible. In this case, it becomes helpful to first project onto  $\Lambda_{\text{nuis}, \psi}$  (where closed-form projections are available), ensuring that the resulting functional remains a valid influence function, even when the second projection is onto the approximated space  $\Omega$ .

The resulting estimator by Liu et al. (2021) is termed "nearly efficient." It is, however, still difficult to implement and has so far only been tested only for a one time point, one acquisition setting.

### 7.3 Semi-offline RL view

In this section, we explore how the semi-offline RL view integrates with semiparametric theory and derive the corresponding influence function for the semi-offline DRL estimator given in Theorem 14.

We begin by establishing that the IPW estimator derived from the semi-offline RL framework can be used to construct an element of the IPW space  $\Lambda_{IPW}$ :

**Lemma 17.** *(Relating the semi-offline RL IPW estimator to the IPW space).  $h_{Semi} \equiv h_{Semi}(A, G_A(X_{(1)}), Y) = \hat{\mathbb{E}}_{n'}[\rho_{Semi}^T C' | A, G_A(X_{(1)}), Y] - J$  is an element of the IPW space  $\Lambda_{IPW}$ .*

The proof of this lemma can be found in Appendix N.

Since constructing the augmentation space  $\Lambda_2$  often involves strong positivity assumptions and projections onto  $\Lambda_2$  are typically not available in closed form, we propose an alternative approach. Specifically, we suggest projecting onto subspaces of  $\Lambda_2$  where closed-form projections are feasible. One such subspace is  $\Lambda_{2,Semi}$ , a tractable subspace of  $\Lambda_2$ , defined as:

$$\Lambda_{2,Semi}(\Xi) = \Lambda_{2,Semi}^1(\Xi) \oplus \Lambda_{2,Semi}^2(\Xi) \oplus \dots \oplus \Lambda_{2,Semi}^T(\Xi)$$

with each

$$\Lambda_{2,Semi}^t(\Xi) = \left\{ b^t(A^t, \underline{A}^{t-1}, G_{\underline{A}^{t-1}}(X_{(1)}), \Xi) \in \mathcal{H} : \mathbb{E}[b(A^t, \underline{A}^{t-1}, G_{\underline{A}^{t-1}}(X_{(1)}), \Xi) | \underline{A}^{t-1}, G_{\underline{A}^{t-1}}(X_{(1)}), \Xi] = 0 \right\}.$$

Notably, we have  $\Lambda_\psi \subseteq \Lambda_{2,Semi}(\Xi) \subset \Lambda_2$  with equality  $\Lambda_\psi = \Lambda_{2,Semi}$  if  $\Xi = \emptyset$ .

This yields the following class of influence functions, proving Theorem 14:

$$\begin{aligned} \varphi_{Semi}(A, G_A(X_{(1)}), Y; \Xi) &= h_{Semi} - \Pi(h_{Semi} | \Lambda_{2,Semi}(\Xi)) \\ &= h_{Semi} - \sum_{t=1}^T \mathbb{E} \left[ h_{Semi} \middle| A^t, \underline{A}^{t-1}, G_{\underline{A}^{t-1}}(X_{(1)}), \Xi \right] + \sum_{t=1}^T \mathbb{E} \left[ h_{Semi} \middle| \underline{A}^{t-1}, G_{\underline{A}^{t-1}}(X_{(1)}), \Xi \right] \\ &= \mathbb{E} \left[ \rho_{Semi}^T f_C(A', G_{A'}(X_{(1)}), Y) | Y, G_A(X_{(1)}), A \right] \\ &\quad - \sum_{t=1}^T \mathbb{E} \left[ \rho_{Semi}^t Q_{Semi}^t \middle| A^t, \underline{A}^{t-1}, G_{\underline{A}^{t-1}}(X_{(1)}), \Xi \right] \\ &\quad + \sum_{t=1}^T \mathbb{E} \left[ \rho_{Semi}^{t-1} V_{Semi}^{t-1} \middle| \underline{A}^{t-1}, G_{\underline{A}^{t-1}}(X_{(1)}), \Xi \right] - J \\ &= \mathbb{E} \left[ \rho_{Semi}^T f_C(A', G_{A'}(X_{(1)}), Y) - \sum_{t=1}^T \rho_{Semi}^t Q_{Semi}^t + \sum_{t=1}^T \rho_{Semi}^{t-1} V_{Semi}^{t-1} \middle| A, G_A(X_{(1)}), Y \right] - J. \end{aligned}$$

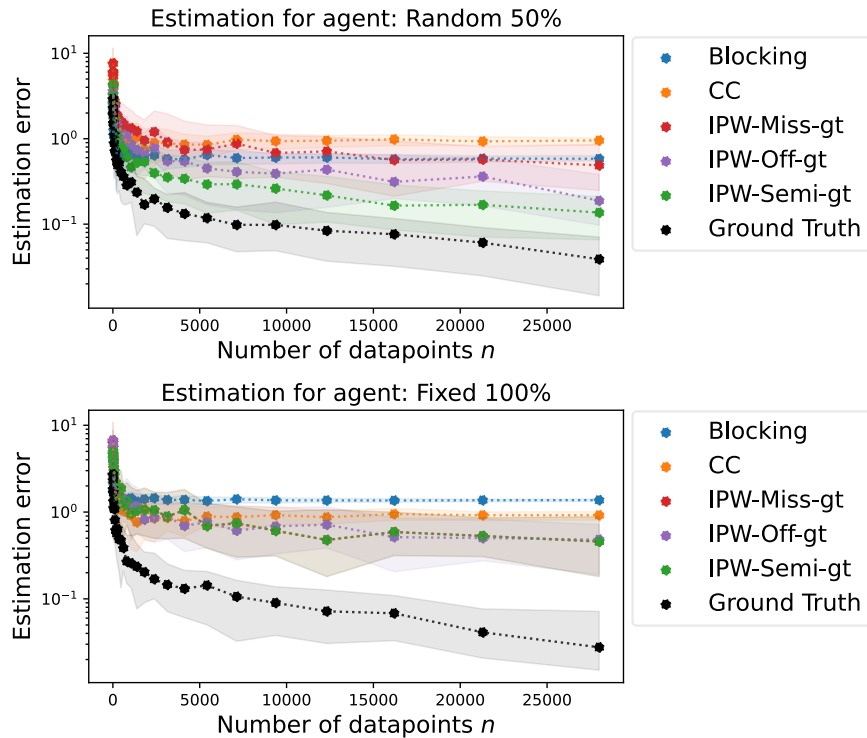


Figure 7: The plots depict convergence as a function of data set size  $n$  for sampling-based estimators in Experiment 1. Two agents are shown: one that acquires each costly feature with a probability of 50% and another with 100%. The estimation error is measured as the absolute difference between the estimate and the ground truth, which is calculated on the full data set ( $n = 40,000$ ). The semi-offline RL IPW estimator converges the fastest for the 'Random 50%' agent, while for the 'Fixed 100%' agent, all IPW estimators perform identically, converging at a slower rate.

We provide full details of this derivation in Appendix O.

From the semiparametric viewpoint, it holds that a larger set  $\Xi$  will increase the efficiency of the corresponding influence function as mentioned in Remark 7. This is the case since employing a larger set  $\Xi$  will result in a larger subspace  $\Lambda_{2,Semi}(\Xi) \subset \Lambda_2$ , which in turn implies a higher efficiency.

## 8. Experiments

We evaluate the different estimators on synthetic data sets where the missingness is artificially induced to allow the comparison with the ground truth.

### 8.1 Experiment Design

We perform five experiments with different violations of the identifying assumptions:

- **Experiment 1:** Assumptions 4 (NDE), 5 (NUC), 6.1 (offline RL positivity), 6.2 (missing data positivity), and 6.4 (maximal global positivity) all hold.

- **Experiment 2:** Assumption 6.2 (missing data positivity) is violated
- **Experiment 3:** Assumption 6.1 (offline RL positivity) is violated for some agents.
- **Experiment 4:** Assumptions 5 (NUC) is violated.
- **Experiment 5:** Assumptions 4 (NDE) is violated.

We evaluate random AFA policies and a proximal policy optimization (PPO) RL agent (Schulman et al., 2017) as AFA agents and use impute-then-regress classifiers (Le Morvan et al., 2021) with unconditional mean imputation and a logistic regression classifier. Nuisance functions ( $\hat{Q}_{Semi}$  and  $\hat{\pi}_\beta$ ) are fitted using multi-layer perceptrons and logistic regression models, respectively. The assumed logistic regression model for the propensity score correctly matches the ground truth. We compare the following estimators:

- *Imp-Mean:* Mean imputation (biased estimator)
- *Blocking:* Blocks the acquisitions of not available features but offers no correction. This corresponds to the estimate  $\hat{\mathbb{E}}_{n'}[C']$  (with  $\pi'_{sim} = \pi'_\alpha$ ) which is biased.
- *CC:* Complete case analysis (only unbiased under MCAR).
- *IPW-Miss/IPW-Miss-gt:* The IPW estimator from the missing data view. The weights were normalized to reduce the variance of the estimator. *IPW-Miss-gt* uses the ground truth propensity score model  $\pi_\beta$  instead of its estimate  $\hat{\pi}_\beta$ .
- *IPW-Off/IPW-Off-gt:* The IPW estimator from the offline RL view with normalized weights and with and without the ground truth propensity score model.
- *IPW-Semi/IPW-Semi-gt:* The IPW estimator (for the maximal global positivity assumption) from the semi-offline RL view with normalized weights and with and without the ground truth propensity score model.
- *DM-Semi:* The semi-offline RL version of the direct method.
- *DRL-Semi/DRL-Semi-gt:* The semi-offline RL version of the double reinforcement learning estimator with normalized weights under the maximal global positivity assumption, with and without the ground truth propensity score model.
- *Ground Truth:* In the experiments where the NDE assumptions hold, the agent is evaluated on the fully observed data set. This corresponds to estimating  $J$  using a Monte Carlo estimate,  $\hat{\mathbb{E}} [C_{(\pi_\alpha)}|X_{(1)}, Y]$ , derived from samples of the ground truth data without any missingness (i.e., samples from  $p(X_{(1)}, Y)$ ). Conversely, in experiments where the NDE assumption is violated, the ground truth is obtained by running the agent in an environment that is continuously sampled from the true data-generating process.

Complete experiment details are given in Appendix P.

## 8.2 Results

Figure 7 shows the convergence plots of sampling-based estimators for Experiment 1, highlighting the data efficiency of all estimators when the identifying assumptions hold. As expected, the blocking and complete case estimators are biased and do not converge to the true value of  $J$ . Among the unbiased IPW estimators, the semi-offline RL IPW estimator

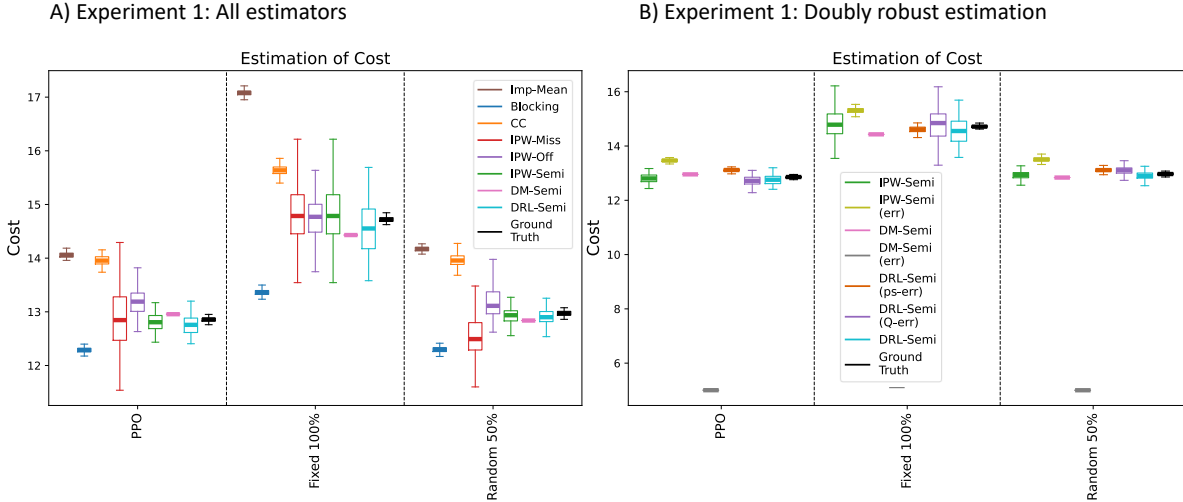


Figure 8: A) General estimation results for Experiment 1. The *Imp-Mean*, *Blocking*, and *CC* estimators show highly biased estimates, while the IPW, semi-offline DM, and DRL estimators align closely with the true target  $J$ . B) Estimation results highlighting the double robustness property of the DRL estimator. The *DRL-Semi* estimator continues to provide accurate estimates even when either the propensity score model  $\hat{\pi}_\beta$  (*DRL-Semi (ps-err)*) or the Q-model  $\hat{Q}_{Semi}$  (*DRL-Semi (Q-err)*) is misspecified.

is the most efficient, achieving the fastest convergence for the 'Random 50%' agent, which acquires each feature with a 50% probability. However, this data efficiency disappears when evaluating the 'Fixed 100%' agent, which acquires all features every time. As noted in Remark 3, all three IPW estimators perform identically in this scenario, as reflected by their equal convergence speeds.

Figure 8A) displays the overall performance of various estimators in Experiment 1. Confidence intervals were computed using the non-parametric bootstrap, excluding the retraining of nuisance functions due to high computational complexity. As a result, the confidence intervals are overly narrow, particularly for the semi-offline DM estimator. The experiment demonstrates that all semi-offline RL estimators accurately approximate the true target parameter  $J$ . However, the DM estimator may exhibit slight bias due to potential misspecification of the Q-function. In contrast, the biased mean imputation, blocking, and complete case analysis estimators fail to consistently estimate  $J$ .

Figure 8B) illustrates the double robustness property of the semi-offline RL version of the DRL estimator. Even when one of the nuisance functions is misspecified, the DRL estimator still provides estimates that closely approximate the true value of  $J$ .

Figures 9A) and B) underscore the importance of the positivity assumption. In Figure 9A), Experiment 2 shows the consequences of violating the missing data positivity assumption (Assumption 6.2)—the fraction of complete cases is only 0.007%. As a result,  $J$  for the 'Fixed 100%' agent cannot be identified from any of the views, and all IPW estimators fail

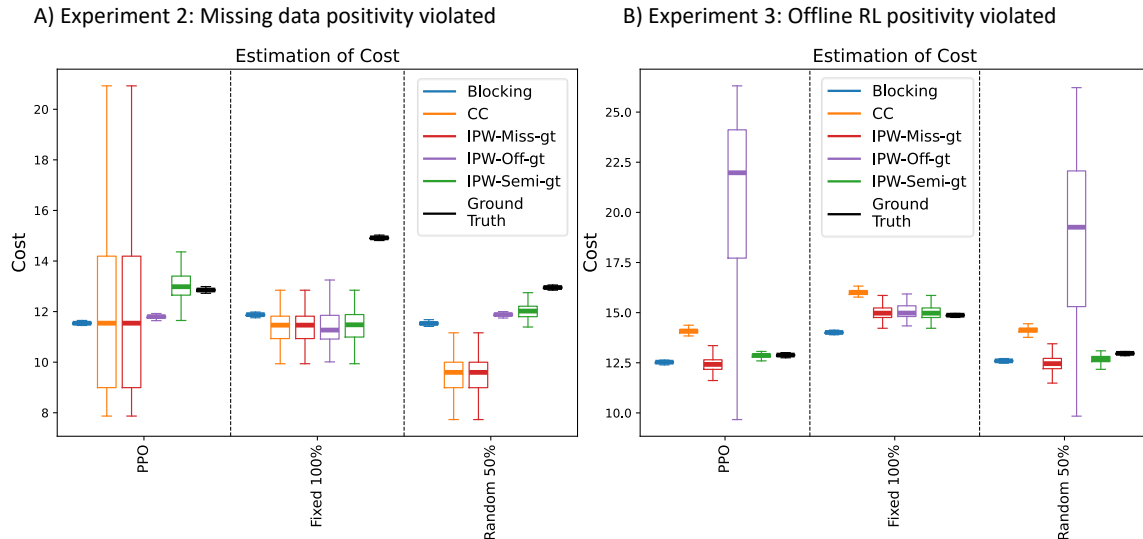


Figure 9: A) Experiment 2: AFA setting with an extremely low fraction of complete cases (0.007%), leading to violations of the missing data positivity assumption. For the 'Fixed 100%' agent, the target  $J$  is not identified from any view. However, the semi-offline IPW estimator is less impacted by the lack of complete cases when evaluating the other agents. B) Experiment 3: AFA setting with positivity violations specific to the offline RL view. While other estimators continue to produce accurate estimates, the offline RL IPW estimator fails entirely for the PPO and 'Random 50%' agents.

to provide accurate estimates. The positivity assumption is also violated for the 'Random 50%' agent, though the impact is less severe for the offline and semi-offline RL estimators.

In Figure 9B, Experiment 3 shows the failure of the positivity assumption required by the offline RL view (Assumption 6.1) for two agents. In this experiment,  $A_2 = \vec{1}$  for all data points, meaning trajectories where  $A_{(\pi_\alpha)} = 0$  have no support. While the other IPW estimators still produce accurate estimates for  $J$ , the estimates from the offline RL IPW estimators fail completely.

Finally, Figures 10A) and B) explore the AFA setting when either the NUC or NDE assumptions are violated. In Figure 10A, which depicts an MNAR scenario, all IPW estimators appear to perform well despite the violation of the NUC assumption. For the 'Fixed 100%' agent, this is expected, as all estimators reduce to the missing data IPW estimator, which is identified. However, for the 'Random 50%' agent, the AFAPE target  $J$  is not identified from the offline or semi-offline RL views, though the estimation from the semi-offline RL IPW estimator remains relatively robust. It is worth cautioning that violating the NUC assumption may have more severe consequences in real-world applications.

Figure 10B demonstrates the effect of violating the NDE assumption. In this case, only the offline RL view yields consistent estimators, as reflected in the experiment. All other estimators show significant deviations from the ground truth. However, the offline RL IPW estimator shows slight biases, too, possibly due to minor positivity violations.

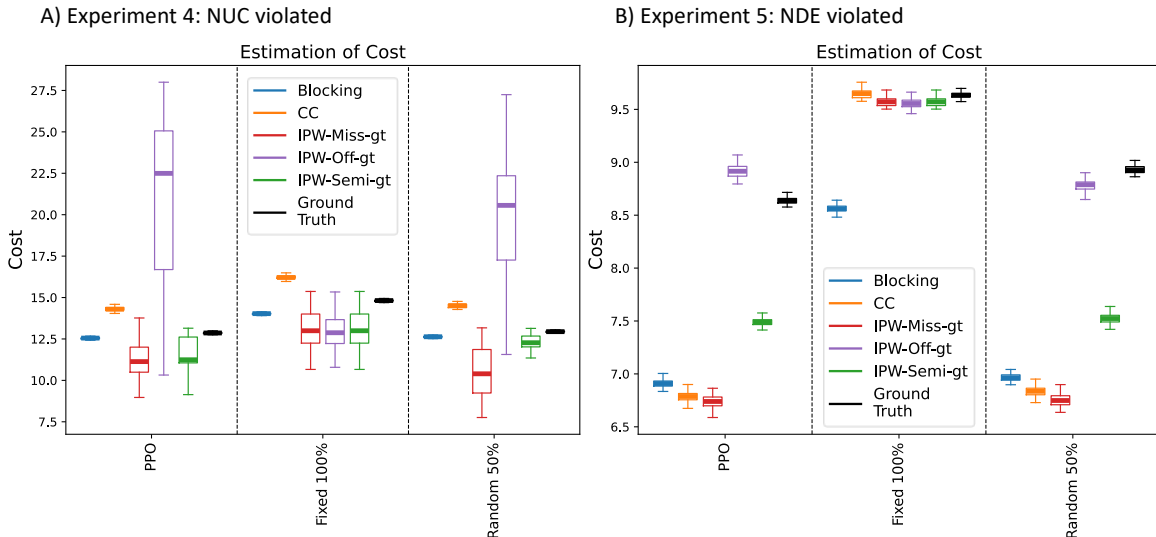


Figure 10: A) Experiment 4: AFA setting with an MNAR acquisition process, resulting in violations of the NUC assumption. Despite this, the semi-offline RL IPW estimator still provides accurate estimates of the target  $J$ . B) Experiment 5: AFA setting with a violation of the NDE assumption. All estimators, except the offline RL IPW estimator, produce highly biased estimates. The offline RL IPW estimator also shows some slight deviation from the true  $J$ , too, potentially caused by minor positivity violations.

## 9. Discussion and Future Work

In this study, we explored the various aspects of solving the AFAPE problem. We acknowledge that there is no one-size-fits-all solution, as the choice of assumptions can vary across AFA settings. To facilitate this discussion, we propose a set of questions that data scientists should ask themselves when tackling the AFAPE problem before choosing a viewpoint and estimator.

1) *What (conditional) independences hold in the data?* The choice of conditional independence assumptions directly influences identifiability and the selection of optimal viewpoints and estimators. When both the NDE and NUC assumptions are violated, the target parameter becomes unidentifiable, making estimation infeasible. If only the NDE assumption fails, the offline RL view is still applicable, while violations of only the NUC assumption allow for the use of the missing data (+online RL) view. In each scenario, leveraging the absence of an edge in the causal graph can effectively eliminate estimation bias that would otherwise persist. When both assumptions are satisfied, one can select from the offline RL, missing data (+online RL), or the novel semi-offline RL view. In such cases, leveraging both the NDE and NUC assumptions through the semi-offline RL view can be advantageous, as it allows for relaxed positivity assumptions and a reduction in estimation variance.

*Conclusion:* Under NUC, one can apply offline RL methods. Under NDE, one can apply missing data methods. Under both NUC and NDE, one can apply semi-offline RL methods.

2) *How much exploration was performed by the retrospective acquisition policy  $\pi_\beta$ ?* Positivity requirements, crucial for all viewpoints, demand certain action sequences to be present in the retrospective data. However, real-world data sets—especially in fields like medicine—often breach these assumptions due to the tendency of professionals to follow similar paths with minimal deviation. The semi-offline RL view, fortunately, imposes weaker positivity constraints than both the offline RL and missing data (+online RL) views. Nevertheless, for "data-hungry" AFA policies, the benefits may be less pronounced compared to the missing data view.

Additionally, the choice of the identifying policy  $\pi_{id}$  depends on which positivity assumptions hold in the data. We leave the adaptation of known positivity assessment methods (Petersen et al., 2012) to the semi-offline RL setting as future work.

*Conclusion:* The semi-offline RL view requires significantly weaker positivity assumptions than the offline RL and missing data (+ online RL) viewpoints.

3) *Can the nuisance models be correctly specified and trained?* The accuracy of different estimators hinges on the proper specification and training of nuisance functions. Despite the double robustness property of DRL estimators, achieving unbiased estimates still depends on how well these functions are modeled. Estimators like the multiple imputation (MI) method can outperform others in certain contexts, particularly when feature smoothness assumptions are reasonable and easily modeled over time. While machine learning techniques such as deep learning offer flexibility, they demand large data sets, which may not always be available.

*Conclusion:* No single viewpoint or estimator is superior across all settings. The choice between MI and semi-offline RL estimators depends on prior knowledge and the feasibility of training the nuisance models.

4) *Is the available data set size sufficient?* Efficient use of data is crucial for accurate estimation. Our experiments demonstrate that estimators based on the semi-offline RL view achieve greater data efficiency compared to both offline RL and missing data estimators.

While our semiparametric analysis shows that the efficiency of all estimators can still be improved, a closed-form efficient influence function does not exist. Some computationally intensive methods, though complex to implement, may still enhance efficiency if the respective strong positivity assumptions hold (Tsiatis, 2006; Liu et al., 2021).

*Conclusion:* Estimators derived from the semi-offline RL view demonstrate in experiments notably higher data efficiency compared to estimators from the offline RL and missing data (+ online RL) viewpoints.

When the answers to the above questions are uncertain, it is advisable to use multiple views and estimators in tandem as part of a broader sensitivity analysis. This approach enhances confidence in the reliability and safety of AFA agents before deployment.

Our study assumes that feature values change over time, making the timing of measurements critical. In our companion paper (von Kleist et al., 2023), we address how a static feature assumption can be incorporated into the AFAPE problem, and we explore how the semi-offline RL and missing data views can be combined when the NUC assumption is violated (i.e., in MNAR scenarios).



Looking forward, we aim to tackle the AFA optimization problem outlined in Section 3.4. Once the AFAPPE problem is resolved and the estimation of the target parameter  $J$  is successful, optimization can commence. This includes training new AFA agents and classifiers. A natural next step is adapting established DTR (Murphy, 2003; Robins, 2004) or offline RL methods—such as off-policy policy gradient methods, actor-critic techniques, and model-based RL approaches (Levine et al., 2020)—to the semi-offline RL framework for further development. These adapted methods could also integrate insights from the online RL literature, such as employing an adaptive exploration policy to enhance the sampling process.

## 10. Conclusion

We study the problem of active feature acquisition performance evaluation (AFAPPE), which involves estimating the acquisition and misclassification costs that an AFA agent would generate after being deployed, using retrospective data. We demonstrate that, depending on the assumptions, one can apply different existing viewpoints to solve AFAPPE. Under the no unobserved confounding (NUC) assumption, one can apply identification and estimation methods from the offline RL literature. Under the no direct effect (NDE) assumption, which assumes the underlying feature values are not affected by their measurement, one can instead apply missing data methods. For settings where both the NUC and the NDE assumptions hold, we propose a novel semi-offline RL viewpoint, which requires weaker positivity assumptions for identification. Within the semi-offline RL viewpoint, we developed several novel estimators that correspond to semi-offline RL versions of the direct method (DM), inverse probability weighting (IPW), and double reinforcement learning (DRL). Finally, we conducted synthetic data experiments to highlight the significance of utilizing proper unbiased estimators for AFAPPE to ensure the reliability and safety of AFA systems.

## Acknowledgments

The present contribution is supported by the Helmholtz Association under the joint research school “HIDSS-006 - Munich School for Data Science @ Helmholtz, TUM & LMU”. Henrik von Kleist received a Carl-Duisberg Fellowship by the Bayer Foundation.

## Appendix A. Literature Review for Active Feature Acquisition (AFA)

In this appendix, we explain in more detail the difference between AFA and related fields and introduce some common approaches to training AFA agents from the literature.

### A.1 Distinction between AFA and Related Fields

AFA is different from active learning (Settles, 2009). In active learning, one assumes a classification task with a training data set that contains many unlabeled data points. The active learning task is then to decide which label acquisitions will improve the training performance the most. Similar research also exists for the acquisition of features for optimal improvement of training. This task has been referred to as "active selection of classification features" (Kok et al., 2021), and unfortunately also as "active feature acquisition" (Huang et al., 2018; Beyer et al., 2020), but its objective differs fundamentally from ours. Huang et al. (2018) attempt to find out which missing values within the retrospective data set would improve training the most when retroactively acquired. In this paper, we are, however, interested which features for a new data point would improve the individual prediction for that data point the most.

### A.2 Approaches to Training AFA Agents

The AFA setting is most generally described as a sequential decision process, which motivates the use of RL-based solutions. One variant, model-based RL, focuses on learning a model for the state transitions. Under the NDE assumption, utilizing an imputation model to capture state transitions becomes feasible, exploiting the unique AFA structure for more straightforward learning (Yoon et al., 2018; Yin et al., 2020; Li and Oliva, 2021a,b; Ma et al., 2019). During deployment, this imputation model can simulate potential outcomes of feature acquisitions, facilitating the derivation of optimal acquisition strategies. Conversely, model-free RL methods do not require a state-transition function. One variant, Q-learning, involves estimating the expected cost of specific acquisition decisions (Chang et al., 2019; Janisch et al., 2020; Shim et al., 2018). For instance, Shim et al. (2018) illustrate the use of double Q-learning for the AFA agent, incorporating a deep neural network that shares network layers for the acquisition decision and classification tasks.

## Appendix B. Review of Semiparametric Theory

We give here a short review of the basic concepts of semiparametric theory and some results for missing data problems. The review is based on work by Tsiatis (2006), which we recommend for more in-depth explanations.

### B.1 General Semiparametric Theory

Semi-parametric theory aims at finding data-efficient estimators for a target parameter  $J = J(p)$  without imposing unnecessarily strict assumptions on  $p$ . In this review, we restrict ourselves to only scalar parameters  $J$ . We let  $p$  denote the distribution  $p(Z)$  over a set of random variables  $Z$  from which we have  $n$  independent and identically distributed samples  $(Z^1, \dots, Z^n)$ . It is possible in many cases to obtain estimators for  $J$  that are consistent at a

rate of  $\sqrt{n}$  without imposing many assumptions. The derivation of such estimators relies on influence functions which are discussed next.

### B.1.1 INFLUENCE FUNCTIONS AND ESTIMATORS

A central element of semi-parametric theory are influence functions as they characterize asymptotically linear estimators in the following sense. An estimator  $J_{est}$  is asymptotically linear and has an influence function  $\varphi(Z) \equiv \varphi$  if it allows the following equality (Tsiatis, 2006):

$$J_{est}(n) - J = \frac{1}{n} \sum_{i=1}^n \varphi(Z^i) + o_p\left(\frac{1}{\sqrt{n}}\right) \quad (37)$$

where  $\varphi \in \mathcal{H}$  and  $\mathcal{H}$  represents the space of all random functions of zero mean and finite variance. The central limit theorem implies that  $J_{est}$  is asymptotically normally distributed (Tsiatis, 2006):

$$\sqrt{n}(J_{est}(n) - J) \rightsquigarrow \mathcal{N}(0, \mathbb{E}[\varphi^2])$$

where  $\rightsquigarrow$  denotes convergence in distribution. The estimation error is thus asymptotically bounded by the variance of the influence function. The efficient influence function  $\varphi_{eff}$  is the one with the smallest asymptotic variance.

Many influence functions, such as the ones of the DRL estimators in this work, depend linearly on the target parameter  $J$ , such that  $\varphi = f(Z) + J$  for some function  $f$ . In these cases, one can very easily derive a corresponding, so-called "1-step" estimator by leveraging Eq. 37 to obtain:

$$J_{est} \equiv -J + \frac{1}{n} \sum_{i=1}^n \varphi(Z_i) = -J + \frac{1}{n} \sum_{i=1}^n f(Z_i) + J = \frac{1}{n} \sum_{i=1}^n f(Z_i).$$

### B.1.2 DERIVING INFLUENCE FUNCTIONS

Deriving the space of influence functions or the efficient influence function for a new target parameter or new model restrictions can be complex. There are, however, some known properties that influence functions in general, or the efficient influence function in particular, have to fulfill and these can be used for their derivation. We examine these now in more detail:

#### 1) An influence function must be in the orthocomp of the nuisance tangent space: $\varphi \in \Lambda_{nuis}^\perp$

To clarify this condition, we first separate the space of model parameters into  $J$ , the target parameter, and  $\eta$ , the nuisance parameters. We denote the nuisance tangent space as  $\Lambda_{nuis}$  and the space orthogonal to it, i.e., its orthocomplement (or orthocomp), as  $\Lambda_{nuis}^\perp$ . The nuisance tangent space can be seen as the collection of directions in which the nuisance parameters can vary without affecting the parameter of interest. It is defined as the mean

square closure of parametric submodel nuisance tangent spaces, where a parametric submodel nuisance tangent space is the linear subspace spanned by the nuisance scores, which are defined as:

$$S_\eta = \left. \frac{\partial \log p_Z(z, \eta, J)}{\partial \eta} \right|_{\eta=\eta_0}.$$

Here  $\eta_0$  denotes the true value of  $\eta$ .

An influence function must be an element in  $\Lambda_{nuis}^\perp$ , but must fulfill also the following normalization (Theorem 3.2 from Tsiatis (2006)):

$$\mathbb{E}[\varphi(Z)S_J(Z)] = 1$$

where  $S_J(Z)$  denotes the scores with respect to the target parameter.

A nuisance function can thus be obtained by taking any nonzero element  $h(Z) \in \mathcal{H}$ , projecting it onto the orthocomp of the nuisance tangent space  $\Lambda_{nuis}^\perp$ , and normalizing it. We denote the mentioned orthogonal projection by  $\Pi$  such that:

$$\varphi^*(Z) = \Pi(h(Z)|\Lambda_{nuis}^\perp) = h(Z) - \Pi(h(Z)|\Lambda_{nuis}).$$

where  $\varphi^*(Z)$  represents an element of  $\Lambda_{nuis}^\perp$  which if normalized would be an influence function.

**2) The efficient influence function must be in the tangent space:**  $\varphi \in \Lambda$

The condition to obtain efficiency for an influence function is that we choose the influence function that is in the tangent space  $\Lambda$ . It can thus be obtained by projecting any influence function onto the tangent space:

$$\varphi_{eff}(Z) = \Pi(\varphi(Z)|\Lambda) = \varphi(Z) - \Pi(\varphi(Z)|\Lambda^\perp).$$

**B.1.3 CONSTRUCTING TANGENT SPACES AND PROJECTING ON THEM**

Here, we provide some useful further details about the construction of tangent spaces with specific restrictions and the projection of random variables on them. We start by decomposing the space of all zero mean, finite variance functions  $\mathcal{H}$  into orthogonal subspaces for the case of a multivariate  $Z$ . We then examine how tangent space restrictions given by conditional independence assumptions can be incorporated.

**1) The decomposition of  $\mathcal{H}$  for a multivariate variable  $Z$**

Firstly, we look at a useful decomposition of the space of a multivariate variable  $Z$  of dimensions  $d$ . The space  $\mathcal{H}$  of all functions  $h(Z_1, Z_2, \dots, Z_d)$  separates into orthogonal subspaces (Theorem 4.5 from Tsiatis (2006)):

$$\mathcal{H} = \mathcal{H}_{Z_1} \oplus \mathcal{H}_{Z_2|Z_1} \oplus \dots \oplus \mathcal{H}_{Z_d|Z_{d-1}, \dots, Z_1} \tag{38}$$

where  $\mathcal{H}_{Z_i|Z_{i-1}, \dots, Z_1}$  denotes the space spanned by the conditional scores:

$$\mathcal{H}_{Z_i|Z_{i-1}, \dots, Z_1} = \left\{ h(Z_i, Z_{i-1}, \dots, Z_1) \in \mathcal{H} : \mathbb{E}[h(Z_i, Z_{i-1}, \dots, Z_1)|Z_{i-1}, \dots, Z_1] = 0 \right\}.$$

Fortunately, also the projection onto such a subspace is known and is given for an arbitrary element  $h^*(Z) \in \mathcal{H}$  as:

$$\Pi(h^*(Z)|\mathcal{H}_{Z_i|Z_{i-1}, \dots, Z_1}) = \mathbb{E}[h^*(Z)|Z_i, Z_{i-1}, \dots, Z_1] - \mathbb{E}[h^*(Z)|Z_{i-1}, \dots, Z_1].$$

## 2) Tangent space restrictions under a conditional independence assumption

It is also possible to define the tangent space under conditional independence restrictions and to project on it. In particular, let's revisit the decomposition of the tangent space from Eq. 38. Let's assume the following independence holds:  $Z_i \perp\!\!\!\perp Z_j | Z_{i-1}, \dots, Z_{j+1}, Z_{j-1}, \dots, Z_1$  for  $i > j$ . We want to find  $\Lambda_r$ , the tangent space restricted by the conditional independence, its orthocomp  $\Lambda_r^\perp$ , and projections on it.

The independence restriction only affects the space of the respective conditional scores such that:

$$\Lambda_r = \mathcal{H}_{Z_1} \oplus \mathcal{H}_{Z_2|Z_1} \oplus \dots \oplus \Lambda_{r, Z_i|Z_{i-1}, \dots, Z_{j+1}, Z_{j-1}, \dots, Z_1} \oplus \dots \oplus \mathcal{H}_{Z_d|Z_{d-1}, \dots, Z_1}.$$

and  $\Lambda_{r, Z_i|Z_{i-1}, \dots, Z_{j+1}, Z_{j-1}, \dots, Z_1} = \mathcal{H}_{Z_i|Z_{i-1}, \dots, Z_{j+1}, Z_{j-1}, \dots, Z_1}$ . The projections onto  $\Lambda_r$  are thus straightforward:

$$\begin{aligned} \Pi(h^*(Z)|\Lambda_{r, Z_i|Z_{i-1}, \dots, Z_{j+1}, Z_{j-1}, \dots, Z_1}) &= \Pi(h^*(Z)|\mathcal{H}_{Z_i|Z_{i-1}, \dots, Z_{j+1}, Z_{j-1}, \dots, Z_1}) \\ &= \mathbb{E}[h^*(Z)|Z_i, Z_{i-1}, \dots, Z_{j+1}, Z_{j-1}, \dots, Z_1] - \mathbb{E}[h^*(Z)|Z_{i-1}, \dots, Z_{j+1}, Z_{j-1}, \dots, Z_1] \end{aligned}$$

Correspondingly, we can obtain the orthocomp as

$$\begin{aligned} \Lambda_r^\perp &= \mathcal{H}_{Z_i|Z_{i-1}, \dots, Z_1} - \Pi(\mathcal{H}_{Z_i|Z_{i-1}, \dots, Z_1} | \Lambda_{r, Z_i|Z_{i-1}, \dots, Z_{j+1}, Z_{j-1}, \dots, Z_1}) \\ &= \left\{ h(Z_i, Z_{i-1}, \dots, Z_1) - \mathbb{E}[h(Z_i, Z_{i-1}, \dots, Z_1) | Z_i, Z_{i-1}, \dots, Z_{j+1}, Z_{j-1}, \dots, Z_1] \right. \\ &\quad \left. + \underbrace{\mathbb{E}[h(Z_i, Z_{i-1}, \dots, Z_1) | Z_{i-1}, \dots, Z_{j+1}, Z_{j-1}, \dots, Z_1]}_{=0} : \right. \\ &\quad \left. \mathbb{E}[h(Z_i, Z_{i-1}, \dots, Z_1) | Z_{i-1}, \dots, Z_1] = 0; h \in \mathcal{H} \right\} \\ &= \left\{ h(Z_i, Z_{i-1}, \dots, Z_1) - \mathbb{E}[h(Z_i, Z_{i-1}, \dots, Z_1) | Z_i, Z_{i-1}, \dots, Z_{j+1}, Z_{j-1}, \dots, Z_1] : \right. \\ &\quad \left. \mathbb{E}[h(Z_i, Z_{i-1}, \dots, Z_1) | Z_{i-1}, \dots, Z_1] = 0; h \in \mathcal{H} \right\}. \end{aligned}$$

## B.2 Semiparametric theory for missing data under the MAR assumption

Semiparametric methods have further been applied to missing data problems. As we take on a missing data view in this work, we now briefly introduce known results for such settings. We restrict our review to scenarios where the missingness process follows a missing-at-random (MAR) scenario. For more details, see also Tsiatis (2006).

We now distinguish between observed data (with missingness) and full data (without missingness). Let the full data be denoted by  $X_{(1)} \in \mathbb{R}^d$ , the missingness indicators by  $A \in \{0, 1\}^d$  and the observed data by  $X \equiv G_A(X_{(1)})$  such that  $X_i = X_{(1),i}$  if  $A_i = 1$  and  $X_i = "?"$ , otherwise.

Semiparametric theory methods for missing data aim at finding observed data (efficient) influence functions from full data influence functions. Throughout this review and the whole paper, we assume no restrictions on the full data  $X_{(1)}$ , meaning that there is only one full data influence function. There are, however, in general multiple corresponding observed data influence functions. In order to find these, the observed data nuisance tangent space needs to be constructed. This construction is simplified in MAR scenarios, as the likelihood factorizes into two separate terms related to the acquisition process and the observed data part of the likelihood (Tsiatis, 2006):

$$p_{A,G_A(X_{(1)})}(A, G_A(X_{(1)}); \psi, J, \eta) = \underbrace{p_{A|G_A(X_{(1)})}(A|G_A(X_{(1)}); \psi)}_{\text{acquisition process}} \sum_{G_A(X_{(1)})} \underbrace{p_{X_{(1)}}(X_{(1)}; J, \eta)}_{\text{process of full data}} \quad (39)$$

where we let  $\psi$  denote the nuisance parameter of the acquisition process, and  $\eta$  denote the nuisance parameter of the full data generating process.

This factorization allows the following decomposition of the nuisance tangent space into orthogonal subspaces (Theorem 8.2 from Tsiatis (2006)):

$$\Lambda_{\text{nuis}} = \Lambda_{\text{nuis},\psi} \oplus \Lambda_{\text{nuis},\eta} = \Lambda_{\psi} \oplus \Lambda_{\text{nuis},\eta}. \quad (40)$$

Here, we used  $\Lambda_{\psi} = \Lambda_{\text{nuis},\psi}$ , which holds since the target parameter is not part of this acquisition process.

We continue with the discussion for the derivation of observed data influence functions in the main text in Section 7.

## Appendix C. Glossary of Terms and Symbols

Term	Description
<i>AFAPE</i>	Active feature acquisition performance evaluation: The problem of estimating the counterfactual cost that would arise if an AFA agent was deployed.
<i>NDE assumption</i>	No direct effect assumption: States that the action of measuring a feature does not impact the values of any features or the label.
<i>NUC assumption</i>	No unobserved confounding assumption: States that acquisition decisions within the retrospective data set were only based on measured feature values.
<i>Semi-offline RL</i>	Novel framework that allows an agent to interact with the environment (the online part) but forbids the exploration of certain actions (the offline part).
<i>DTR</i>	Dynamic treatment regimes
<i>G-formula</i>	Identification formula from causal inference (Robins, 1986)

<i>Plug-in of the G-formula</i>	Estimation formula from causal inference that replaces unknown densities in the G-formula with estimated versions (Robins, 1986).
<i>IPW</i>	Inverse probability weighting: Estimator that is also known as importance sampling or the Horvitz-Thompson estimator.
<i>DM</i>	Direct method: Estimator based on a Q-function.
<i>DRL</i>	Double reinforcement learning: Double robust estimator that uses IPW weights and a Q-function.
<i>m-graph</i>	Missing data graph: Graph to visualize assumptions in missing data problems.
<i>MI</i>	Multiple imputation: Estimator for missing data problems that is a special case of the plug-in of the G-formula.
<i>influence function</i>	Function of mean zero and finite variance that is used to analyze the asymptotic properties of regular and asymptotically linear (RAL) estimators.
<i>MCAR assumption</i>	Missing-completely-at-random assumption: States that the reason for the missingness of certain features does not depend on any feature values.
<i>MAR assumption</i>	Missing-at-random assumption: States that the reason for missingness of certain features does only depend on observed feature values.
<i>MNAR assumption</i>	Missing-not-at-random assumption: States that the reason for the missingness of certain features may depend on feature values that are not observed.
<i>nuisance function</i>	A function that needs to be fitted from data in order to use a corresponding estimator but which is not of primary interest itself. Examples are the propensity score model and the Q-function.
<i>local positivity assumption</i>	Positivity assumption for semi-offline RL that ensures the simulation of a desired next action is possible from the retrospective data set.
<i>regional positivity assumption</i>	Positivity assumption for semi-offline RL that ensures the simulation of all future desired actions is possible from the retrospective data set.
<i>global positivity assumption</i>	Positivity assumption for semi-offline RL that ensures the simulation of all desired actions is possible from step 1 on.
<i>maximal regional positivity assumption</i>	A special, stronger version of the regional positivity assumption.

<i>maximal global positivity assumption</i>	A special, stronger version of the global positivity assumption.
<i>tangent space</i>	Space of scores (i.e., derivatives of the log-likelihood)
<i>nuisance tangent space</i>	Space of nuisance scores (i.e., the scores with respect to the nuisance parameters)

<b>Symbol</b>	<b>Description</b>
$t \in (0, \dots, T)$	Time
$U^t$	Unobserved state variables at time t
$d$	Number of features, i.e., dimension of $U^t$
$X^t = G_{A^t}(U^t)$	Observed feature values at time t (retrospective data set)
$A^t$	Acquisition action at time t (retrospective data set)
$\mathcal{A}$	Space of A
$Y$	Label
$Y^* = f_{cl}(\underline{X}^T, \underline{A}^T)$	Predicted label based on classifier $f_{cl}$
$C_a^t$	Acquisition cost for action $A^t$
$C_{mc} = f_C(Y^*, Y)$	Misclassification cost (if $Y$ and $Y^*$ differ)
$\pi_\beta$	Retrospective acquisition policy
$\pi_\alpha$	AFA policy
$C_{mc,(\pi_\alpha)}$	Counterfactual misclassification cost had $\pi_\alpha$ instead of $\pi_\beta$ been applied
$g(\cdot)$	known deterministic distribution
$J / J_{mc}$	Expected misclassification cost under the AFA policy and classifier
$J_a$	Expected acquisition cost under the AFA policy and classifier
$\phi_1^*, \phi_2^*$	Sets of parameters that parameterize the AFA policy and the classifier, respectively
$q(\cdot)$	Counterfactual distribution
$Q_{Off}^t$	State-action value function from offline RL (at time t)
$V_{Off}^t$	State value function from offline RL (at time t)
$\pi'$	Blocked policy
$\pi'_{sim}$	(Blocked) simulation policy
$p'(\cdot)$	Simulated distribution
$C', Y'^*, X', A'$	Simulated cost, predicted label, features, and actions



Symbol	Description
$\mathcal{D}$	Retrospective data set
$\mathcal{D}'$	Simulated data set
$\mathcal{A}_{adm}$	Local admissible set
$\tilde{\mathcal{A}}_{adm}$	Regional admissible set
$\pi_{id}$	Distribution for $A$ that allows identification of $J$ under the semi-offline RL view (subject to support restrictions)
$Q_{Semi}^t$	State-action value function from semi-offline RL (at time $t$ )
$\Xi \subseteq \{O, Y\}$	Arbitrary subset of the always observed features $O \subset X_{(1)}$ and the label $Y$
$q'(\cdot)$	Counterfactual simulated distribution
$\varphi$	Influence function
$\Lambda$	(Observed data) tangent space
$\Lambda^\perp$	Orthocomplement of the (observed data) tangent space
$\Pi([\cdot] \Lambda)$	(Orthogonal) projection onto the tangent space
$\Lambda_{nuis}$	(Observed data) nuisance tangent space
$\Lambda^F$	Full data tangent space
$\Lambda_{nuis}^F$	Full data nuisance tangent space
$\Lambda_{nuis,\psi} = \Lambda_\psi$	(Nuisance) tangent space of the acquisition process
$\Lambda_{nuis,\eta}$	Nuisance tangent space of the observed part of the full data process
$\Lambda_{IPW}$	IPW space
$\Lambda_2$	Augmentation space
$\Lambda_{2,Semi}(\Xi)$	Subspace of $\Lambda_2$ proposed for projection onto under the semi-offline RL view

## Appendix D. Identification of the Block-conditional Model

In this appendix, we demonstrate how identification of  $p(X_{(1)}, Y)$  can be achieved when the NUC assumption (Assumption 5) is violated. This corresponds to the block-conditional model (Zhou et al., 2010). We show that the propensity score model  $p(A = \bar{1}|X_{(1)}, Y)$  is identified, which in turn results in identification of  $p(X_{(1)}, Y)$ , as  $p(X_{(1)}, Y) = \frac{p(X_{(1)}, Y, A=\bar{1})}{p(A=\bar{1}|X_{(1)}, Y)}$ .

*Identification of  $p(A = \vec{1}|X_{(1)}, Y)$ :*

The propensity score is identified by

$$\begin{aligned} p(A = \vec{1}|X_{(1)}, Y) &= \prod_{t=1}^T p(A^t = \vec{1}|X_{(1)}, Y, \underline{A}^{t-1} = \vec{1}) \\ &\stackrel{*1}{=} \prod_{t=1}^T \pi_{\beta}^t(A^t = \vec{1}|\underline{X}_{(1)}^{t-1}, \underline{A}^{t-1} = \vec{1}) \\ &\stackrel{*2}{=} \prod_{t=1}^T \pi_{\beta}^t(A^t = \vec{1}|\underline{X}^{t-1}, \underline{A}^{t-1} = \vec{1}) \end{aligned}$$

where we used in \*1) the fact that future feature values do not affect current acquisition decisions:  $A^t \perp\!\!\!\perp \overline{X}_{(1)}^t, Y|\underline{X}_{(1)}^{t-1}, \underline{A}^{t-1}$ . We further use in \*2) that counterfactual feature values  $\underline{X}_{(1)}^{t-1}$  are equal to  $\underline{X}^{t-1}$  if  $\underline{A}^{t-1} = \vec{1}$ . The last expression is a function of only observed variables and is thus identified.

## Appendix E. Multiple Imputation (MI) for the AFAPe Problem

In this appendix, we aim to delve deeper into the multiple imputation (MI) estimator in the AFAPe context and highlight advantages as well as some common pitfalls associated with using MI approaches in AFA.

Let us begin by emphasizing a significant advantage of the MI estimator compared to other estimators discussed in this paper. It offers an elegant solution to the temporal coarsening problem. In time-series settings, where fixed time intervals are assumed ( $t \in \{0, 1, \dots, T\}$ ), employing a very fine resolution of time steps would inevitably result in a considerable increase in missingness, thereby making the AFAPe problem more challenging. The MI estimator can typically overcome this issue by assuming an often justifiable temporal smoothness of the feature distributions.

However, there are drawbacks to MI. MI requires modeling joint distributions, which is a complex task in practice, particularly in high-dimensional settings and when dealing with complex missingness patterns. For instance, the multiple imputation by chained equations (MICE) method (van Buuren, 2007) necessitates fitting  $d$  conditional densities for  $d$  partially observed features in static settings. In comparison, IPW only requires the specification of the propensity score, which is often more feasible. This effect is especially drastic for high-dimensional features such as images, which necessitate modeling for each pixel when using multiple imputation, but only the modeling of one joint missingness indicator when using IPW.

Furthermore, the MI estimator implies imputation of the missing features  $X_m$  by conditioning on the observed features  $X_o$  and the label  $Y$  (i.e., estimating  $\hat{p}(X_m|X_o, Y)$ ). This introduces the risk of data leakage, as the imputed features may carry predictive information not because of the true data generation mechanism but due to the imputation itself, resulting in a potentially over-optimistic estimation of prediction performance. A common alternative, frequently employed in machine learning, is to impute the data without conditioning on  $Y$ . However, this assumption implies that a missing feature  $X_{(1),i} \in X_m$  is

conditionally independent of the label given the observed features ( $X_{(1),i} \perp\!\!\!\perp Y|X_o$ ). Determining the marginal predictive value of a feature for predicting  $Y$  is, however, the whole task of AFA, which renders this approach impractical.

Conditional mean imputation represents a simplified imputation approach that reduces the complexity of modeling. It has been applied in AFA settings (An et al., 2022; Erion et al., 2021; Janisch et al., 2020). In this approach, missing values are imputed using a conditional mean model for  $\hat{\mathbb{E}}[X_m|X_o]$  (or  $\hat{\mathbb{E}}[X_m|X_o, Y]$ ). Therefore, conditional mean imputation assumes:

$$\begin{aligned} \hat{J}_{MI-Miss} &= \sum_{X_m, X_o, Y} \mathbb{E}[C_{(\pi_\alpha)}|X_m, X_o, Y] \hat{p}(X_m|X_o, Y) p(X_o, Y) \\ &\approx \sum_{X_o, Y} \mathbb{E}[C_{(\pi_\alpha)}|\hat{\mathbb{E}}[X_m|X_o], X_o, Y] p(X_o, Y) \end{aligned}$$

which does not hold in general and can lead to strongly biased results when  $\mathbb{E}[C_{(\pi_\alpha)}|X_m, X_o, Y]$  is nonlinear as is the case generally in AFA settings.

## Appendix F. Proof of Lemma 6

In this appendix, we prove Lemma 6, which we repeat here for readability:

**Lemma 6.** *(Sufficiency conditions for global positivity). The global positivity assumption for semi-offline RL (Assumption 6.3) holds if the positivity assumption from offline RL (Assumption 6.1) or from missing data (Assumption 6.2) holds.*

We split the proof into the following propositions:

**Proposition 18.** *If the positivity assumption for offline RL (Assumption 6.1) holds, then the global positivity assumption for semi-offline RL also holds.*

**Proposition 19.** *If the positivity assumption for missing data (Assumption 6.2) holds, then the global positivity assumption for semi-offline RL also holds.*

We begin by proving Proposition 18:

**Proof** Consider the initial time step  $t = 1$ , focusing on data points  $(a^1, x^0)$  such that  $p(x^0)\pi_\alpha(a^1|x^0) > 0$ . To establish the global positivity assumption, we need to ensure that the regional positivity assumption is satisfied.

We will show that  $a^1 = a'^1$  belongs to the regional admissible set  $\tilde{\mathcal{A}}_{adm}^1(x^0, a^1)$ , which implies that the regional positivity condition is met. Two conditions must be fulfilled for  $a^1 = a'^1$  to be included in  $\tilde{\mathcal{A}}_{adm}^1(x^0, a^1)$ . First,  $a^1 = a'^1$  must belong to the local admissible set  $\mathcal{A}_{adm}^1(x^0, a^1)$ , which directly follows from the offline RL positivity assumption. Second, we must ensure that for  $a^1 = a'^1$ , the regional admissible set exists at the subsequent time step. Specifically, since  $a^1 = a'^1$  (and therefore  $x^1 = x'^1$ ), the regional admissible set at time step 2,  $\tilde{\mathcal{A}}_{adm}^2(\underline{a}^2, \underline{x}^1, \underline{a}^1)$ , must exist for all  $x^1$  and  $a'^2$  such that:

$$p(x^1|x^0, a^1)\pi_\alpha(a'^2|\underline{x}^1, \underline{a}^1) > 0.$$

For these conditions, the positivity assumption under the offline RL view again implies that  $a^2 = \underline{a}'^2$  is included in the local admissible set  $\mathcal{A}_{adm}^2(a'^2, \underline{x}^1, \underline{a}^1)$ . This reasoning can be extended iteratively through all time steps up to  $T$ , thus proving that regional positivity holds at every prior time point, which in turn establishes global positivity. ■

Next, we prove Proposition 19, following a similar strategy:

**Proof** Again, consider the initial time step  $t = 1$ , focusing on data points  $(a^1, x^0)$  such that  $p(x^0)\pi_\alpha(a^1|x^0) > 0$ . We show that  $a^1 = \vec{1}$  is included in the regional admissible set  $\tilde{\mathcal{A}}_{adm}^1(x^0, a^1)$ , irrespective of the value of  $a^1$ . First,  $a^1 = \vec{1}$  must be included in the local admissible set  $\mathcal{A}_{adm}^1(x^0, a^1)$ , a condition that directly follows from the missing data positivity assumption.

Second,  $a^1 = \vec{1}$  must permit the existence of a regional admissible set at the next time step. Specifically, given  $a^1 = \vec{1}$ , the regional admissible set at time step 2,  $\tilde{\mathcal{A}}_{adm}^2(\underline{a}'^2, \underline{x}^1, \underline{a}^1) = \tilde{\mathcal{A}}_{adm}^2(\underline{a}'^2, \underline{x}^1, \underline{a}^1 = \vec{1})$ , must exist for all  $x^1$  and  $a'^2$  such that:

$$p(x^1|x^0, A^1 = \vec{1})\pi_\alpha(a'^2|\underline{x}^1, \underline{a}^1) > 0.$$

Under these conditions, the missing data positivity assumption again directly implies that  $a^2 = \vec{1}$  belongs to the local admissible set at time step 2,  $\tilde{\mathcal{A}}_{adm}^2(\underline{a}'^2, \underline{x}^1, \underline{a}^1 = \vec{1})$ , regardless of the values of  $a'^2$  and  $x^1$ . This reasoning can be extended step-by-step until  $T$ , thereby ensuring that regional positivity holds at each prior time point and, consequently, that global positivity is satisfied as well. ■

## Appendix G. Proof of Theorems 7 and 8

In this Appendix, we prove Theorems 7 and 8. We also demonstrate how the positivity assumption arises. We restate the theorems here for clarity and ease of reference.

**Theorem 7.** *(Identification of  $J$  for the semi-offline RL view). The reformulated AFAPe problem of estimating  $J$  under the semi-offline RL view (Eq. 17) is under Assumption 1 (no measurement noise), Assumption 2 (consistency), Assumption 3 (no interference), Assumption 4 (NDE), Assumption 5 (NUC) and Assumption 6.3 (global positivity) identified by*

$$J = \mathbb{E}_{p'}[C'_{(\pi_\alpha)}] = \sum_{A', A, G_A(X_{(1)}), Y} f_C(A', X', Y)q'(A', A, X, Y) \quad (19)$$

with the distribution

$$q'(A', A, X, Y) = \prod_{t=1}^T \underbrace{\pi_{id}^t(A^t|\underline{A}^t, \underline{X}^{t-1}, \underline{A}^{t-1})}_{\text{distr. subject to constraints}} \underbrace{\pi_\alpha^t(A^t|\underline{X}^{t-1}, \underline{A}^{t-1})}_{\text{target policy}} \prod_{t=0}^T p(X^t|\underline{X}^{t-1}, \underline{A}^t, Y)p(Y) \quad (20)$$

where

$$\pi_{id}^t(A^t | \underline{A}^t, \underline{X}^{t-1}, \underline{A}^{t-1}) = \underbrace{\mathbb{I}(A^t \in \tilde{\mathcal{A}}_{adm}^t(\underline{X}^{t-1}, \underline{A}^{t-1}, \underline{A}^t))}_{\text{support restriction}} f_{id}^t(\underline{A}^t, \underline{X}^{t-1}, \underline{A}^{t-1}) \quad (21)$$

for any function  $f_{id}^t$  s.t.  $\pi_{id}^t$  is a valid density.

**Theorem 8.** (Bellman equation for semi-offline RL). *The semi-offline RL view admits under Assumption 1 (no measurement noise), Assumption 2 (consistency), Assumption 3 (no interference), Assumption 4 (NDE), Assumption 5 (NUC) and the local positivity assumption at datapoint  $\underline{x}^{t-1}, \underline{a}^{t-1}, a^t$  (from Definition 4), the following semi-offline RL version of the Bellman equation:*

$$Q_{Semi}(\underline{A}^t, \underline{X}^{t-1}, \underline{A}^{t-1}, \Xi) = \sum_{X^t} V_{Semi}(\underline{A}^t, \underline{X}^t, \underline{A}^{t-1}, A^t = a^t, \Xi) p(X^t | \underline{X}^{t-1}, \underline{A}^{t-1}, A^t = a^t, \Xi) \quad (22)$$

for any  $a^t \in \mathcal{A}_{adm}^t(\underline{X}^{t-1}, \underline{A}^{t-1}, \underline{A}^t)$

$$V_{Semi}(\underline{A}^t, \underline{X}^t, \underline{A}^t, \Xi) = \sum_{A^{t+1}} Q_{Semi}(\underline{A}^{t+1}, \underline{X}^t, \underline{A}^t, \Xi) \pi_{\alpha}^{t+1}(A^{t+1} | \underline{X}^t, \underline{A}^t) \quad (23)$$

with semi-offline RL versions of the state-action value function  $Q_{Semi}$  and state value function  $V_{Semi}$ :

$$\begin{aligned} Q_{Semi}^t &\equiv Q_{Semi}(\underline{A}^t, \underline{X}^{t-1}, \underline{A}^{t-1}, \Xi) \equiv \mathbb{E}_{p'}[C'_{(\bar{\pi}_{\alpha}^{t+1})} | \underline{A}^t, \underline{X}^{t-1}, \underline{A}^{t-1}, \Xi] \\ V_{Semi}^t &\equiv V_{Semi}(\underline{A}^t, \underline{X}^t, \underline{A}^t, \Xi) \equiv \mathbb{E}_{p'}[C'_{(\bar{\pi}_{\alpha}^{t+1})} | \underline{A}^t, \underline{X}^t, \underline{A}^t, \Xi] \end{aligned}$$

where  $C'_{(\bar{\pi}_{\alpha}^{t+1})}$  denotes the potential outcome of  $C'$  under interventions from time step  $t+1$  onwards.  $\Xi \subseteq \{Y, O\}$ , with  $O$  denoting all features that are always available, denotes an optional subset of additional variables that can be conditioned on. Furthermore,  $Q_{Semi}^t$  and  $V_{Semi}^t$  are identified if the regional positivity assumption (from Definition 5) holds at  $\underline{X}^{t-1}, \underline{A}^{t-1}, \underline{A}^t$  and  $a^t \in \tilde{\mathcal{A}}_{adm}^t(\underline{X}^{t-1}, \underline{A}^{t-1}, \underline{A}^t)$ .

**Proof** Firstly, we factorize the counterfactual distribution, denoted by  $q'$ , expressing it as a function of the observed (simulated) data. We factorize the graph step-by-step to show how the semi-offline RL version of the Bellman equation arises. We split identification in each step into two parts to emphasize the two parts of the Bellman equation. To help guide the identification, we duplicate Figure 6 of the causal graph describing the simulation process in Figure 11A). Alongside it, we show the counterfactual graph (for identification step  $t = 1$ ) in Figure 11B).

### Step 0

*Counterfactual factorization (step  $t = 0$ , part 1):*

$$p' \left( C'_{(\pi_{\alpha})} \right) \stackrel{*1}{\equiv} p' \left( C'_{(\bar{\pi}_{\alpha})} \right) = \sum_{X^0, \Xi} p' \left( C'_{(\bar{\pi}_{\alpha})} \middle| X^0, \Xi \right) p(X^0, \Xi)$$

where we denote in \*1)  $C'_{(\bar{\pi}_\alpha^1)}$  as the counterfactual  $C'$  under an intervention of  $\pi_\alpha$  from step  $t = 1$  onwards. The extension by  $X^0$  is needed for adjustment. The inclusion of  $\Xi \subseteq \{Y, O\}$ , where  $O$  denotes the subset of always observed features amongst  $X_{(1)}$ , is optional.

*Counterfactual factorization (step  $t = 0$ , part 2):*

$$\begin{aligned}
 p' \left( C'_{(\bar{\pi}_\alpha^1)} \middle| X^0, \Xi \right) &= \sum_{\mathbf{a}^1} p' \left( C'_{(\bar{\pi}_\alpha^2, \mathbf{a}^1)} \middle| X^0, \Xi \right) \pi_\alpha^1(\mathbf{a}^1 | \mathbf{X}^0) \\
 &\stackrel{*1}{=} \sum_{\mathbf{a}^1} p' \left( C'_{(\bar{\pi}_\alpha^2, \mathbf{a}^1, \pi_{id}^1)} \middle| X^0, \Xi \right) \pi_\alpha^1(\mathbf{a}^1 | X^0) \\
 &= \sum_{\mathbf{a}^1, \mathbf{a}^1} p' \left( C'_{(\bar{\pi}_\alpha^2, \mathbf{a}^1, \mathbf{a}^1)} \middle| X^0, \Xi \right) \pi_{id}^1(\mathbf{a}^1 | \mathbf{X}^0, \mathbf{a}^1) \pi_\alpha^1(\mathbf{a}^1 | X^0) \\
 &\stackrel{*2}{=} \sum_{\mathbf{a}^1, \mathbf{a}^1} p' \left( C'_{(\bar{\pi}_\alpha^2, \mathbf{a}^1, \mathbf{a}^1)} \middle| X^0, \mathbf{a}^1, \Xi \right) \pi_{id}^1(\mathbf{a}^1 | X^0, \mathbf{a}^1) \pi_\alpha^1(\mathbf{a}^1 | X^0) \\
 &\stackrel{*3}{=} \sum_{\mathbf{a}^1, \mathbf{a}^1} p' \left( C'_{(\bar{\pi}_\alpha^2, \mathbf{a}^1)} \middle| X^0, \mathbf{a}^1, \Xi \right) \pi_{id}^1(\mathbf{a}^1 | X^0, \mathbf{a}^1) \pi_\alpha^1(\mathbf{a}^1 | X^0) \\
 &\stackrel{*4}{=} \sum_{\mathbf{a}^1, \mathbf{a}^1} p' \left( C'_{(\bar{\pi}_\alpha^2, \mathbf{a}^1)} \middle| X^0, \mathbf{a}^1, \mathbf{a}^1, \Xi \right) \pi_{id}^1(\mathbf{a}^1 | X^0, \mathbf{a}^1) \pi_\alpha^1(\mathbf{a}^1 | X^0) \\
 &\stackrel{*5}{=} \sum_{\mathbf{a}^1, \mathbf{a}^1} p' \left( C'_{(\bar{\pi}_\alpha^2)} \middle| X^0, \mathbf{a}^1, \mathbf{a}^1, \Xi \right) \pi_{id}^1(\mathbf{a}^1 | X^0, \mathbf{a}^1) \pi_\alpha^1(\mathbf{a}^1 | X^0) \\
 &\stackrel{*6}{=} \sum_{\mathbf{a}^1, \mathbf{a}^1} p' \left( C'_{(\bar{\pi}_\alpha^2)} \middle| \mathbf{X}^0, \mathbf{a}^1, \Xi \right) \pi_{id}^1(\mathbf{a}^1 | X^0, \mathbf{a}^1) \pi_\alpha^1(\mathbf{a}^1 | X^0)
 \end{aligned}$$

with the following explanations:

- \*1): We notice that  $C'_{(\bar{\pi}_\alpha^2, \mathbf{a}^1)}$  is independent of any interventions  $\pi_{id}^1$  on  $A^1$ . This step prevents positivity problems in subsequent steps.
- \*2): We use the exchangeability  $C'_{(\bar{\pi}_\alpha^2, \mathbf{a}^1, \mathbf{a}^1)} \perp\!\!\!\perp A^1 | X^0, \Xi$ , which follows from the NUC assumption.
- \*3): We use the consistency assumption:  

$$p' \left( C'_{(\bar{\pi}_\alpha^2, \mathbf{a}^1, \mathbf{a}^1)} \middle| X^0, \mathbf{a}^1, \Xi \right) = p' \left( C'_{(\bar{\pi}_\alpha^2, \mathbf{a}^1)} \middle| X^0, \mathbf{a}^1, \Xi \right)$$
- \*4): We use the exchangeability:  $C'_{(\bar{\pi}_\alpha^2, \mathbf{a}^1)} \perp\!\!\!\perp A^1 | X^0, A^1, \Xi$
- \*5): We use the consistency assumption:  $p' \left( C'_{(\bar{\pi}_\alpha^2, \mathbf{a}^1)} \middle| X^0, \mathbf{a}^1, \mathbf{a}^1, \Xi \right) = p' \left( C'_{(\bar{\pi}_\alpha^2)} \middle| X^0, \mathbf{a}^1, \mathbf{a}^1, \Xi \right)$
- \*6): We use the conditional independence  $C'_{(\bar{\pi}_\alpha^2)} \perp\!\!\!\perp A^1 | X^0, A^1, \Xi$

We must also ensure that  $p' \left( C'_{(\bar{\pi}_\alpha^2, \mathbf{a}^1)} \middle| X^0, \mathbf{a}^1, \mathbf{a}^1, \Xi \right)$ , i.e. conditioning on  $X^0, A^1, A^1, \Xi$ , is well specified in \*4). To understand what positivity requirements are necessary, we first factorize the "observational" (i.e., simulated) distribution for step  $t = 0$ . By observational distribution for step  $t = 0$ , we refer to a distribution which only contains interventions from step  $t = 2$  onwards:




Figure 11: Causal graph for the distribution  $p'$ . A) Simulated ("observational") distribution. B) Counterfactual distribution under the intervention  $\pi_{\alpha}^2$ . Edges showing long-term dependencies are omitted from the graphs for visual clarity. These include:  $\underline{X}_{(1)}^{t-1} \leftrightarrow X_{(1)}^t$ ;  $\underline{X}_{(1)}^T \leftrightarrow Y$ ;  $\underline{X}^{t-1}, \underline{A}^{t-1} \rightarrow A^t$ ;  $\underline{X}^{t-1}/\underline{X}_{(\pi_{\alpha}^2)}^{t-1}, \underline{A}^{t-1}/\underline{A}_{(\pi_{\alpha}^2)}^{t-1} \rightarrow A^t/A_{(\pi_{\alpha}^2)}^t$ ; and  $\underline{X}^T/\underline{X}_{(\pi_{\alpha}^2)}^T, \underline{A}^T/\underline{A}_{(\pi_{\alpha}^2)}^T \rightarrow Y'^*/Y_{(\pi_{\alpha}^2)}'^*$ .

Observational factorization (step  $t = 0$ ):

$$p'(C'_{(\pi_{\alpha}^2)}) = \sum_{X^0, A^1, A'^1, \Xi} p'(C'_{(\pi_{\alpha}^2)} | X^0, A^1, \Xi) \underbrace{\pi'_{sim}(A'^1 | X^0, A^1)}_{\text{simulation policy}} \underbrace{\pi_{\beta}^1(A^1 | X^0)}_{\text{retro. acq. policy}} p(X^0, \Xi)$$

By comparing the observational and counterfactual factorizations, we see that the following positivity assumption is required:

$$\begin{aligned} \text{if} \quad & p(x^0)q'(a'^1, a^1 | x^0) = p(x^0)\pi_{\alpha}^1(a'^1 | x^0)\pi_{id}^1(a^1 | x^0, a'^1) > 0 \\ \text{then} \quad & p(x^0)p'(a'^1, a^1 | x^0) = p(x^0)\pi_{sim}^1(a'^1 | x^0, a^1)\pi_{\beta}^1(a^1 | x^0) \geq \mathcal{O} \\ & \forall x^0, a'^1, a^1, \text{ and some constant } \mathcal{O} > 0 \end{aligned} \quad (41)$$

Since none of the distributions  $\pi_{\alpha}$ ,  $\pi_{id}$ ,  $\pi'_{sim}$ , and  $\pi_{\beta}$  depend on  $\Xi$ , the choice for  $\Xi$  will not influence the positivity requirements. We can further simplify the positivity assumption by using knowledge about the known simulation policy  $\pi'_{sim}$ . By the construction of the blocking operation of the simulation policy  $\pi'_{sim}$  (Definition 2), one observes that

$$\text{if } \pi_{\alpha}^1(a'^1 | x^0) > 0, \quad \text{then } \pi_{sim}^1(a'^1 | x^0, a^1) \geq \mathcal{O}_1, \quad \text{if and only if } a'^1 \leq a^1$$

where,  $\mathcal{O}_1$  is some constant  $> 0$ , and, as before, we let  $a'^1 \leq a^1$  denote the element-wise comparison. The resulting positivity violation for the case  $a'^1 \not\leq a^1$  can be avoided by restricting  $\pi_{id}$  in the following way:

*Restriction 1 for  $\pi_{id}$  (step  $t = 0$ ):*

$$\text{if } a'^1 \not\subseteq a^1, \quad \text{then } \pi_{id}^1(a^1|x^0, a'^1) = 0 \quad \forall x^0, a'^1, a^1.$$

A second possible positivity violation arises if  $\pi_{\beta}^1(a^1|x^0) = 0$  for some values of  $a^1$ . This poses a second requirement for  $\pi_{id}$ :

*Restriction 2 for  $\pi_{id}$  (step  $t = 0$ ):*

$$\text{if } \pi_{\beta}^1(a^1|x^0) = 0, \quad \text{then } \pi_{id}^1(a^1|x^0, a'^1) = 0 \quad \forall x^0, a'^1, a^1.$$

Since  $\pi_{id}$  is required to be a valid probability distribution (it cannot be 0 for all  $a^1$ ), this imposes the following requirement for  $\pi_{\beta}$ :

$$\text{if } p(x^0)\pi_{\alpha}^1(a'^1|x^0) > 0, \quad \text{then } \pi_{\beta}^1(A^1 \geq a'^1|x^0) \geq \mathcal{O} \quad \forall x^0, a'^1, \text{ and some constant } \mathcal{O} > 0.$$

The positivity assumption implies that for any desired action  $a'^1$  by the target policy  $\pi_{\alpha}$ , there exists at least positive support for one set of acquisitions  $a^1$  that include equal or more acquisitions than what is contained in  $a'^1$ . This is equivalent to the local positivity assumption at  $x^0, a'^0$  (i.e. the existence of  $\mathcal{A}_{adm}^1$  from Definition 4). In the next steps, we show that these are only minimal requirements for  $\pi_{id}^1(A^1|X^0, A'^1)$ . To avoid running into positivity violations in later time steps, a further restriction can be necessary.

### Step 1

In the following, we continue the identification for step  $t = 1$ .

*Counterfactual factorization (step  $t = 1$ , part 1):*

$$\begin{aligned} p' \left( C'_{(\bar{\pi}_{\alpha}^2)} \middle| X^0, A'^1, \Xi \right) &= p' \left( C'_{(\bar{\pi}_{\alpha}^2)} \middle| X^0, A'^1, a^1, \Xi \right) = \\ &= \sum_{X^1} p' \left( C'_{(\bar{\pi}_{\alpha}^2)} \middle| \underline{A}^1, \underline{X}^1, a^1, \Xi \right) p(X^1|X^0, a^1, \Xi) \end{aligned}$$

which holds for any  $a^1 \in \mathcal{A}_{adm}^1(X^0, A'^1)$  (because local positivity must hold). Therefore, the term  $p' \left( C'_{(\bar{\pi}_{\alpha}^2)} \middle| \underline{A}^1, \underline{X}^1, a^1, \Xi \right)$  needs to be only identified for *at least one value*  $a^1 \in \mathcal{A}_{adm}^1(X^0, A'^1)$ .



Counterfactual factorization (step  $t = 1$ , part 2):

$$\begin{aligned}
 p' \left( C'_{(\bar{\pi}_\alpha^2)} \middle| \underline{A}^1, \underline{X}^1, \underline{A}^1, \Xi \right) &= \\
 &= \sum_{\mathbf{a}'^2} p' \left( C'_{(\bar{\pi}_\alpha^3, \mathbf{a}'^2)} \middle| \underline{A}^1, \underline{X}^1, \underline{A}^1, \Xi \right) \pi_\alpha^2(\mathbf{a}'^2 | \underline{X}'^1, \underline{A}'^1) \\
 &\stackrel{*1}{=} \sum_{\mathbf{a}'^2} p' \left( C'_{(\bar{\pi}_\alpha^3, \mathbf{a}'^2, \pi_{id}^2)} \middle| \underline{A}^1, \underline{X}^1, \underline{A}^1, \Xi \right) \pi_\alpha^2(\mathbf{a}'^2 | \underline{X}'^1, \underline{A}'^1) \\
 &= \sum_{\mathbf{a}'^2, \mathbf{a}^2} p' \left( C'_{(\bar{\pi}_\alpha^3, \mathbf{a}'^2, \mathbf{a}^2)} \middle| \underline{A}^1, \underline{X}^1, \underline{A}^1, \Xi \right) \pi_{id}^2(\mathbf{a}^2 | \underline{A}'^1, \mathbf{a}'^2, \underline{X}^1, \underline{A}^1) \pi_\alpha^2(\mathbf{a}'^2 | \underline{X}'^1, \underline{A}'^1) \\
 &\stackrel{*2}{=} \sum_{\mathbf{a}'^2, \mathbf{a}^2} p' \left( C'_{(\bar{\pi}_\alpha^3, \mathbf{a}'^2)} \middle| \underline{A}^1, \underline{X}^1, \underline{A}^1, \mathbf{a}^2, \Xi \right) \pi_{id}^2(\mathbf{a}^2 | \underline{A}'^1, \mathbf{a}'^2, \underline{X}^1, \underline{A}^1) \pi_\alpha^2(\mathbf{a}'^2 | \underline{X}'^1, \underline{A}'^1) \\
 &\stackrel{*3}{=} \sum_{\mathbf{a}'^2, \mathbf{a}^2} p' \left( C'_{(\bar{\pi}_\alpha^3)} \middle| \underline{A}'^1, \mathbf{a}'^2, \underline{X}^1, \underline{A}^1, \mathbf{a}^2, \Xi \right) \pi_{id}^2(\mathbf{a}^2 | \underline{A}'^1, \mathbf{a}'^2, \underline{X}^1, \underline{A}^1) \pi_\alpha^2(\mathbf{a}'^2 | \underline{X}'^1, \underline{A}'^1) \\
 &\stackrel{*4}{=} \sum_{\mathbf{a}'^2, \mathbf{a}^2} p' \left( C'_{(\bar{\pi}_\alpha^3)} \middle| \underline{A}'^1, \mathbf{a}'^2, \underline{X}^1, \underline{A}^1, \Xi \right) \pi_{id}^2(\mathbf{a}^2 | \underline{A}'^1, \mathbf{a}'^2, \underline{X}^1, \underline{A}^1) \pi_\alpha^2(\mathbf{a}'^2 | \underline{X}'^1, \underline{A}'^1)
 \end{aligned}$$

where we denote  $\underline{X}'^1 \equiv G_{\underline{A}'^1}(X_{(1)})$  which contains a subset of the features in  $\underline{X}^1 = G_{\underline{A}^1}(X_{(1)})$ . The derivation was based on the following arguments:

- \*1): We use that  $C'_{(\bar{\pi}_\alpha^3, \mathbf{a}'^2)}$  is independent of any interventions  $\pi_{id}^2$  on  $A^2$ .
- \*2): We use exchangeability :  $C'_{(\bar{\pi}_\alpha^3, \mathbf{a}'^2, \mathbf{a}^2)} \perp\!\!\!\perp A^2 | \underline{A}'^1, \underline{X}^1, \underline{A}^1, \Xi$  and consistency:  
 $p' \left( C'_{(\bar{\pi}_\alpha^3, \mathbf{a}'^2, \mathbf{a}^2)} \middle| \underline{A}'^1, \underline{X}^1, \underline{A}^1, \mathbf{a}^2, \Xi \right) = p' \left( C'_{(\bar{\pi}_\alpha^3, \mathbf{a}'^2)} \middle| \underline{A}'^1, \underline{X}^1, \underline{A}^1, \mathbf{a}^2, \Xi \right)$  for  $A^2$ .
- \*3): We use the exchangeability  $C'_{(\bar{\pi}_\alpha^3, \mathbf{a}'^2)} \perp\!\!\!\perp A^2 | \underline{A}'^1, \underline{X}^1, \underline{A}^1, \mathbf{a}^2, \Xi$  and consistency:  
 $p' \left( C'_{(\bar{\pi}_\alpha^3, \mathbf{a}'^2)} \middle| \underline{A}'^1, \mathbf{a}'^2, \underline{X}^1, \underline{A}^1, \mathbf{a}^2, \Xi \right) = p' \left( C'_{(\bar{\pi}_\alpha^3)} \middle| \underline{A}'^1, \mathbf{a}'^2, \underline{X}^1, \underline{A}^1, \mathbf{a}^2, \Xi \right)$  for  $A^2$ .
- \*4): We use the conditional independence  $C'_{(\bar{\pi}_\alpha^3)} \perp\!\!\!\perp A^2 | \underline{A}'^1, \mathbf{a}'^2, \underline{X}^1, \underline{A}^1, \Xi$

We must also ensure in \*3) that  $p' \left( C'_{(\bar{\pi}_\alpha^3, \mathbf{a}'^2)} \middle| \underline{A}'^1, \mathbf{a}'^2, \underline{X}^1, \underline{A}^1, \mathbf{a}^2, \Xi \right)$ , i.e., conditioning on  $\underline{A}'^1, \mathbf{a}'^2, \underline{X}^1, \underline{A}^1, \mathbf{a}^2, \Xi$ , is well specified. To understand what positivity requirements are necessary, we factorize the "observational" (i.e., simulated) distribution for step  $t = 1$ .

Observational factorization (step  $t = 1$ ):

$$\begin{aligned}
 p' \left( C'_{(\bar{\pi}_\alpha^3)} \middle| \underline{A}^1, X^0, \underline{A}^1, \Xi \right) &= \sum_{X^1, \mathbf{a}^2, \mathbf{a}'^2} p' \left( C'_{(\bar{\pi}_\alpha^3)} \middle| \underline{A}'^1, \mathbf{a}'^2, \underline{X}^1, \underline{A}^1, \Xi \right) \\
 &\quad \cdot \underbrace{\pi_{sim}^2(\mathbf{a}'^2 | \underline{X}'^1, \underline{A}'^1, \mathbf{a}^2)}_{\text{known simulation policy}} \underbrace{\pi_\beta^2(\mathbf{a}^2 | \underline{X}^1, \underline{A}^1)}_{\text{retro. acquisition policy}} p(X^1 | X^0, \underline{A}^1, \Xi)
 \end{aligned}$$

By comparing the observational and counterfactual factorizations, we see that the following positivity assumption is required:

$$\begin{aligned}
 \text{if} \quad & q'(\underline{a}^1, a'^2, \underline{x}^1, \underline{a}^1, a^2) = \\
 & = q'(\underline{a}^1, \underline{x}^1, \underline{a}^1) \pi_{\alpha}^2(a'^2 | \underline{x}^1, \underline{a}^1) \pi_{id}^2(a^2 | \underline{a}^1, a'^2, \underline{x}^1, \underline{a}^1) > 0 \\
 \text{then} \quad & p'(\underline{x}^1, a'^2, \underline{x}^1, \underline{a}^1, a^2) = \\
 & = p'(\underline{a}^1, \underline{x}^1, \underline{a}^1) \pi_{sim}^2(a'^2 | \underline{x}^1, \underline{a}^1, a^2) \pi_{\beta}^2(a^2 | \underline{x}^1, \underline{a}^1) \geq \mathcal{O} \\
 & \forall \underline{x}^1, \underline{a}^1, \underline{a}^1, a'^2, a^2, \text{ and some constant } \mathcal{O} > 0
 \end{aligned}$$

with the following factorizations:

$$\begin{aligned}
 q'(\underline{A}^1, \underline{X}^1, \underline{A}^1) &= q'(\underline{A}^1, \underline{X}^0, \underline{A}^1) p(X^1 | \underline{X}^0, \underline{A}^1) \\
 p'(\underline{A}^1, \underline{X}^1, \underline{A}^1) &= p'(\underline{A}^1, \underline{X}^0, \underline{A}^1) p(X^1 | \underline{X}^0, \underline{A}^1)
 \end{aligned}$$

Note that we can again ignore  $\Xi$  since it doesn't affect positivity requirements. The positivity condition can again be simplified through the two restrictions on  $\pi_{id}^2$ :

*Restrictions 1 and 2 for  $\pi_{id}$  (step  $t = 2$ ):*

$$\text{if } a'^2 \not\leq a^2 \text{ or } \pi_{\beta}^2(a^2 | \underline{x}^1, \underline{a}^1) = 0, \quad \text{then } \pi_{id}^2(a^2 | \underline{a}^1, a'^2, \underline{x}^1, \underline{a}^1) = 0 \quad \forall a^2, a'^2$$

This imposes the requirement for  $\pi_{\beta}^2$  that there exists at least one value  $a^2$  such that  $a'^2 \leq a^2$  and  $\pi_{\beta}^2(a^2 | \underline{x}^1, \underline{a}^1) \geq \mathcal{O}$  (i.e. local positivity at  $\underline{x}^1, \underline{a}^1, a'^1$ ). Notice, however, that this has to hold for *all values*  $a^1$  that were "allowed" in step  $t = 1$  (i.e. where  $\pi_{id}^1(a^1 | x^0, a^1) \geq \mathcal{O}$ ). As  $\pi_{id}^1$  only needs to have support for at least one  $a^1 \in \mathcal{A}_{adm}^1$ , we can restrict  $\pi_{id}^1$  at step  $t = 0$  further to reduce the positivity assumption for step  $t = 2$ . We do, however, only want to restrict  $\pi_{id}^1$  as much as necessary because if  $\pi_{id}^1$  has wider support, this means that more data points are used in the analysis. Therefore, we introduce the notion of regional positivity and the regional admissible set  $\tilde{\mathcal{A}}_{adm}$  (from Definition 5). In particular, Definition 5 defines  $\tilde{\mathcal{A}}_{adm}^1(x^0, a^1)$  as the subset of  $\mathcal{A}_{adm}^1(x^0, a^1)$  such that local positivity holds at step  $t = 1$  for all possible values of  $x^1$ , and  $a'^2$ . As this has to hold for future time-steps as well (as will be shown next), the definition for  $\tilde{\mathcal{A}}_{adm}$  even states regional positivity has to hold recursively, i.e., also at  $t = 1$ .

In summary, local positivity at step  $t = 2$  ensures that the available data allows the simulation of the currently desired action  $a'^2$ . Regional positivity at step  $t = 0$  ensures that only those simulations are used at step  $t = 0$  such that simulations of desired actions in the future (at step  $t = 1$ ) are possible with the data.

### Step t

Now, we generalize the factorization to step  $t$ .

*Counterfactual factorization (step  $t$ , part 1):*

$$\begin{aligned}
 p' \left( C'_{(\bar{\pi}_{\alpha}^{t+1})} \middle| \underline{A}^t, \underline{X}^{t-1}, \underline{A}^{t-1}, \Xi \right) &= p' \left( C'_{(\bar{\pi}_{\alpha}^{t+1})} \middle| \underline{A}^t, \underline{X}^{t-1}, \underline{A}^{t-1}, a^t, \Xi \right) = \\
 &= \sum_{X^t, \underline{X}^t} p' \left( C'_{(\bar{\pi}_{\alpha}^{t+1})} \middle| \underline{A}^t, \underline{X}^t, \underline{A}^{t-1}, a^t, \Xi \right) p(X^t | \underline{X}^{t-1}, \underline{A}^{t-1}, a^t, \Xi)
 \end{aligned} \tag{42}$$

which holds for any  $a^t \in \mathcal{A}_{adm}^t(\underline{X}^{t-1}, \underline{A}^{t-1}, A^t)$  (because local positivity must hold). Therefore, the term  $p' \left( C'_{(\bar{\pi}_\alpha^{t+1})} \middle| \underline{A}^t, \underline{X}^t, \underline{A}^{t-1}, a^t, \Xi \right)$  needs to be only identified for *at least one value*  $a^t \in \mathcal{A}_{adm}^t(\underline{X}^{t-1}, \underline{A}^{t-1}, A^t)$ .

*Counterfactual factorization (step t, part 2):*

$$\begin{aligned}
 & p' \left( C'_{(\bar{\pi}_\alpha^{t+1})} \middle| \underline{A}^t, \underline{X}^t, \underline{A}^t, \Xi \right) = \tag{43} \\
 &= \sum_{a^{t+1}} p' \left( C'_{(\bar{\pi}_\alpha^{t+2}, a^{t+1})} \middle| \underline{A}^t, \underline{X}^t, \underline{A}^t, \Xi \right) \pi_\alpha^{t+1}(a^{t+1} | \underline{X}^t, \underline{A}^t) \\
 &\stackrel{*1}{=} \sum_{a^{t+1}} p' \left( C'_{(\bar{\pi}_\alpha^{t+2}, a^{t+1}, \pi_{id}^{t+1})} \middle| \underline{A}^t, \underline{X}^t, \underline{A}^t, \Xi \right) \pi_\alpha^{t+1}(a^{t+1} | \underline{X}^t, \underline{A}^t) \\
 &= \sum_{a^{t+1}, a^{t+1}} p' \left( C'_{(\bar{\pi}_\alpha^{t+2}, a^{t+1}, a^{t+1})} \middle| \underline{A}^t, \underline{X}^t, \underline{A}^t, \Xi \right) \pi_{id}^{t+1}(a^{t+1} | \underline{A}^t, a^{t+1}, \underline{X}^t, \underline{A}^t) \pi_\alpha^{t+1}(a^{t+1} | \underline{X}^t, \underline{A}^t) \\
 &\stackrel{*2}{=} \sum_{a^{t+1}, a^{t+1}} p' \left( C'_{(\bar{\pi}_\alpha^{t+2}, a^{t+1})} \middle| \underline{A}^t, \underline{X}^t, \underline{A}^t, a^{t+1}, \Xi \right) \pi_{id}^{t+1}(a^{t+1} | \underline{A}^t, a^{t+1}, \underline{X}^t, \underline{A}^t) \pi_\alpha^{t+1}(a^{t+1} | \underline{X}^t, \underline{A}^t) \\
 &\stackrel{*3}{=} \sum_{a^{t+1}, a^{t+1}} p' \left( C'_{(\bar{\pi}_\alpha^{t+2})} \middle| \underline{A}^t, a^{t+1}, \underline{X}^t, \underline{A}^t, a^{t+1}, \Xi \right) \pi_{id}^{t+1}(a^{t+1} | \underline{A}^t, a^{t+1}, \underline{X}^t, \underline{A}^t) \pi_\alpha^{t+1}(a^{t+1} | \underline{X}^t, \underline{A}^t) \\
 &\stackrel{*4}{=} \sum_{a^{t+1}, a^{t+1}} p' \left( C'_{(\bar{\pi}_\alpha^{t+2})} \middle| \underline{A}^t, a^{t+1}, \underline{X}^t, \underline{A}^t, \Xi \right) \underbrace{\pi_{id}^{t+1}(a^{t+1} | \underline{A}^t, a^{t+1}, \underline{X}^t, \underline{A}^t)}_{\text{arbitrary dist. subject to constraints}} \underbrace{\pi_\alpha^{t+1}(a^{t+1} | \underline{X}^t, \underline{A}^t)}_{\text{target policy}}
 \end{aligned}$$

with the following explanations:

- \*1): We use that  $C'_{(\bar{\pi}_\alpha^{t+2}, a^{t+1})}$  is independent of any interventions  $\pi_{id}^{t+1}$  on  $A^{t+1}$ .
- \*2): We use exchangeability :  $C'_{(\bar{\pi}_\alpha^{t+2}, a^{t+1}, a^{t+1})} \perp\!\!\!\perp A^{t+1} | \underline{A}^t, \underline{X}^t, \underline{A}^t, \Xi$  and consistency:  $p' \left( C'_{(\bar{\pi}_\alpha^{t+2}, a^{t+1}, a^{t+1})} \middle| \underline{A}^t, \underline{X}^t, \underline{A}^t, a^{t+1}, \Xi \right) = p' \left( C'_{(\bar{\pi}_\alpha^{t+2}, a^{t+1})} \middle| \underline{A}^t, \underline{X}^t, \underline{A}^t, a^{t+1}, \Xi \right)$  for  $A^{t+1}$ .
- \*3): We use the exchangeability  $C'_{(\bar{\pi}_\alpha^{t+2}, a^{t+1})} \perp\!\!\!\perp A^{t+1} | \underline{A}^t, \underline{X}^t, \underline{A}^t, a^{t+1}, \Xi$  and consistency:  $p' \left( C'_{(\bar{\pi}_\alpha^{t+2}, a^{t+1})} \middle| \underline{A}^t, a^{t+1}, \underline{X}^t, \underline{A}^t, a^{t+1}, \Xi \right) = p' \left( C'_{(\bar{\pi}_\alpha^{t+2})} \middle| \underline{A}^t, a^{t+1}, \underline{X}^t, \underline{A}^t, a^{t+1}, \Xi \right)$  for  $A^{t+1}$ .
- \*4): We use the conditional independence  $C'_{(\bar{\pi}_\alpha^{t+2})} \perp\!\!\!\perp A^{t+1} | \underline{A}^t, a^{t+1}, \underline{X}^t, \underline{A}^t, \Xi$

As before, we have to make sure in \*3) that  $p' \left( C'_{(\bar{\pi}_\alpha^{t+2}, a^{t+1})} \middle| \underline{A}^t, a^{t+1}, \underline{X}^t, \underline{A}^t, a^{t+1}, \Xi \right)$ , i.e. conditioning on  $\underline{A}^t, a^{t+1}, \underline{X}^t, \underline{A}^t, a^{t+1}, \Xi$ , is well specified. To understand what positivity requirements are necessary, we factorize the "observational" (i.e., simulated) distribution for step  $t$ :

*Observational factorization (step  $t$ ):*

$$\begin{aligned}
 p' \left( C'_{(\bar{\pi}_\alpha^{t+2})} \middle| \underline{A}^t, \underline{X}^{t-1}, \underline{A}^t, \Xi \right) &= \sum_{X^t, a^{t+1}, a'^{t+1}} p' \left( C'_{(\bar{\pi}_\alpha^{t+2})} \middle| \underline{A}^t, a^{t+1}, \underline{X}^t, \underline{A}^t, \Xi \right) \\
 &\quad \cdot \underbrace{\pi_{sim}^{t+1}(a^{t+1} | \underline{X}^t, \underline{A}^t, a^{t+1})}_{\text{known simulation policy}} \underbrace{\pi_\beta^{t+1}(a^{t+1} | \underline{X}^t, \underline{A}^t)}_{\text{retro. acquisition policy}} p(X^t | \underline{X}^{t-1}, \underline{A}^t, \Xi)
 \end{aligned}$$

By comparing the observational and counterfactual factorizations, we see that the following positivity assumption is required:

$$\begin{aligned}
 \text{if} \quad & q'(\underline{a}^t, a^{t+1}, \underline{x}^t, \underline{a}^t, a^{t+1}) = \\
 & q'(\underline{a}^t, \underline{x}^t, \underline{a}^t) \pi_\alpha^{t+1}(a^{t+1} | \underline{x}^t, \underline{a}^t) \pi_{id}^{t+1}(a^{t+1} | \underline{a}^t, a^{t+1}, \underline{x}^t, \underline{a}^t) > 0 \\
 \text{then} \quad & p'(a^{t+1}, \underline{x}^t, \underline{a}^t, a^{t+1}) = \\
 & p'(\underline{a}^t, \underline{x}^t, \underline{a}^t) \pi_{sim}^{t+1}(a^{t+1} | \underline{x}^t, \underline{a}^t, a^{t+1}) \pi_\beta^{t+1}(a^{t+1} | \underline{x}^t, \underline{a}^t) \geq \mathcal{O} \\
 & \forall \underline{x}^t, \underline{a}^t, \underline{a}^t, a^{t+1}, a^{t+1}, \text{ and some constant } \mathcal{O} > 0
 \end{aligned}$$

with the following factorizations:

$$\begin{aligned}
 q'(\underline{A}^t, \underline{X}^t, \underline{A}^t) &= q'(\underline{A}^t, \underline{X}^{t-1}, \underline{A}^t) p(X^t | \underline{X}^{t-1}, \underline{A}^t) \\
 p'(\underline{A}^t, \underline{X}^t, \underline{A}^t) &= p'(\underline{A}^t, \underline{X}^{t-1}, \underline{A}^t) p(X^t | \underline{X}^{t-1}, \underline{A}^t)
 \end{aligned}$$

The positivity condition can again be simplified through the two restrictions on  $\pi_{id}$ :

*Restrictions 1 and 2 for  $\pi_{id}$  (step  $t$ ):*

$$\begin{aligned}
 \text{if} \quad & a^{t+1} \not\leq a^{t+1} \text{ or } \pi_\beta^{t+1}(a^{t+1} | \underline{x}^t, \underline{a}^t) = 0, \\
 \text{then} \quad & \pi_{id}^{t+1}(a^{t+1} | \underline{a}^t, a^{t+1}, \underline{x}^t, \underline{a}^t) = 0 \\
 & \forall a^{t+1}, a'^{t+1}
 \end{aligned}$$

This imposes the requirement for  $\pi_\beta$  that there exists at least one value  $a^{t+1}$  such that  $a^{t+1} \leq a^{t+1}$  and  $\pi_\beta^{t+1}(a^{t+1} | \underline{x}^t, \underline{a}^t) \geq \mathcal{O}$  (i.e. local positivity at  $\underline{x}^t, \underline{a}^t, a^t$ ). This has to hold for *all values*  $\underline{a}^t$  that were "allowed" in all previous steps (i.e. all  $a^t$ , for all  $\tau \leq t$ , s.t.  $\pi_{id}^t(a^\tau | \underline{a}^{\tau-1}, a'^{\tau-1}, \underline{x}^{\tau-1}, \underline{a}^{\tau-1}) > 0$  and which could later on have led to the current state). As  $\pi_{id}$  only needs to have support for at least one  $a^t \in \mathcal{A}_{adm}^t$  per step, we can restrict  $\pi_{id}$  at *all* previous steps to reduce the positivity assumption for step  $t$ . Again, we do not want to restrict  $\pi_{id}$  too much because if  $\pi_{id}$  has wider support, this means that more data points are used in the analysis. The regional positivity assumption (from Definition 5) ensures in this case that only those simulations are used (and exist) at all previous steps such that simulations of the desired actions can be performed at step  $t$  (and for future steps).

## Full factorization

Bringing all time-steps  $t = 0, \dots, T$  together and including  $Y$  (as a part of  $\Xi$ ) one obtains the full factorization of the identifying distribution  $q'$ :

$$\begin{aligned}
 q'(A', A, X, Y) &= \prod_{t=1}^T \underbrace{\pi_{id}^t(A^t | \underline{A}^t, \underline{X}^{t-1}, \underline{A}^{t-1})}_{\text{arb. distr. subject to constraints}} \underbrace{\pi_{\alpha}^t(A^t | \underline{X}^{t-1}, \underline{A}^{t-1})}_{\text{target policy}} \prod_{t=0}^T p(X^t | \underline{X}^{t-1}, \underline{A}^t, Y) p(Y) \\
 &= \prod_{t=1}^T \pi_{id}^t(A^t | \underline{A}^t, G_{\underline{A}^{t-1}}(X_{(1)}), \underline{A}^{t-1}) \pi_{\alpha}^t(A^t | G_{\underline{A}^{t-1}}(X_{(1)}), \underline{A}^{t-1}) \\
 &\quad \cdot \prod_{t=0}^T p(G_{A^t}(X_{(1)}) | G_{\underline{A}^{t-1}}(X_{(1)}), \underline{A}^t, Y) p(Y) \\
 &= q'(A', A, G_A(X_{(1)}), Y)
 \end{aligned}$$

In order for this expression to hold,  $\pi_{id}^t$  must be restricted to have support only on  $\tilde{\mathcal{A}}_{adm}$ . This leads to the following restriction:

$$\pi_{id}^t(A^t | \underline{A}^t, \underline{X}^{t-1}, \underline{A}^{t-1}) = \mathbb{I}(A^t \in \underbrace{\tilde{\mathcal{A}}_{adm}(\underline{A}^t, \underline{X}^{t-1}, \underline{A}^{t-1})}_{\text{support restriction}}) f_{id}^t(\underline{A}^t, \underline{X}^{t-1}, \underline{A}^{t-1}).$$

where  $f_{id}^t$  is an arbitrary function that ensures that  $\pi_{id}^t$  is a valid density. This concludes the proof of Theorem 7.

### Bellman equation

Equations 42 and 43 correspond to the two parts of the semi-offline RL version of the Bellman equation:

$$\begin{aligned}
 \mathbb{E} \left[ C'_{(\bar{\pi}_{\alpha}^{t+1})} | \underline{A}^t, \underline{X}^{t-1}, \underline{A}^{t-1}, \Xi \right] &= \sum_{X^t} \mathbb{E} \left[ C'_{(\bar{\pi}_{\alpha}^{t+1})} | \underline{A}^t, \underline{X}^t, \underline{A}^{t-1}, a^t, \Xi \right] p(X^t | \underline{X}^{t-1}, \underline{A}^{t-1}, a^t, \Xi) \\
 &= \sum_{G_{A^t}(X_{(1)})} \mathbb{E} \left[ C'_{(\bar{\pi}_{\alpha}^{t+1})} | \underline{A}^t, G_{\underline{A}^t}(X_{(1)}), \underline{A}^{t-1}, a^t, \Xi \right] p(G_{A^t}(X_{(1)}) | G_{\underline{A}^{t-1}}(X_{(1)}), \underline{A}^{t-1}, a^t, \Xi) \\
 &\text{for any } a^t \in \mathcal{A}_{adm}^t(\underline{X}^{t-1}, \underline{A}^{t-1}, \underline{A}^t)
 \end{aligned}$$

$$\begin{aligned}
 \mathbb{E} \left[ C'_{(\bar{\pi}_{\alpha}^{t+1})} | \underline{A}^t, \underline{X}^t, \underline{A}^t, \Xi \right] &= \sum_{A^{t+1}} \mathbb{E} \left[ C'_{(\bar{\pi}_{\alpha}^{t+2})} | \underline{A}^{t+1}, \underline{X}^t, \underline{A}^t, \Xi \right] \pi_{\alpha}^{t+1}(A^{t+1} | \underline{X}^t, \underline{A}^t) \\
 &= \sum_{A^{t+1}} \mathbb{E} \left[ C'_{(\bar{\pi}_{\alpha}^{t+2})} | \underline{A}^{t+1}, G_{\underline{A}^t}(X_{(1)}), \underline{A}^t, \Xi \right] \pi_{\alpha}^{t+1}(A^{t+1} | G_{\underline{A}^{t-1}}(X_{(1)}), \underline{A}^t)
 \end{aligned}$$

The factorization holds under local positivity (if  $\mathcal{A}_{adm}^t \neq \emptyset$  exists). Furthermore, the individual terms are identified if regional positivity holds, which concludes the proof of Theorem 8.  $\blacksquare$

## Appendix H. Proof of Corollary 10

In this appendix, we prove Corollary 10, stating identification under the maximal global positivity assumption. We repeat it here for ease of reference:

**Corollary 10.** *(Identification of  $J$  for the semi-offline RL view under maximal global positivity). The reformulated AFAPPE problem of estimating  $J$  under the semi-offline RL view (Eq. 17) is under Assumption 1 (no measurement noise), Assumption 2 (consistency), Assumption 3 (no interference), Assumption 4 (NDE), Assumption 5 (NUC) and Assumption 6.4 (maximum global positivity) identified by Eqs. 19 and 20 where*

$$\begin{aligned} \pi_{id}^t(A^t | \underline{A}^{t-1}, A^t = a^t, \underline{X}^{t-1}, \underline{A}^{t-1}) &= \\ &= \mathbb{I}(A^t \geq a^t) \pi_{\beta}^t(A^t | \underline{X}^{t-1}, \underline{A}^{t-1}) f_{id}^t(\underline{A}^{t-1}, A^t = a^t, \underline{X}^{t-1}, \underline{A}^{t-1}) \end{aligned}$$

for any function  $f_{id}^t$  s.t.  $\pi_{id}^t$  is a valid density. This holds in particular for the choice of a truncated  $\pi_{\beta}$ :

$$\begin{aligned} \pi_{id}^t(A^t | \underline{A}^{t-1}, A^t = a^t, \underline{X}^{t-1}, \underline{A}^{t-1}) &= \pi_{\beta}^t(A^t | A^t \geq a^t, \underline{X}^{t-1}, \underline{A}^{t-1}) = \\ &= \frac{\mathbb{I}(A^t \geq a^t) \pi_{\beta}^t(A^t | \underline{X}^{t-1}, \underline{A}^{t-1})}{\pi_{\beta}^t(A^t \geq a^t | \underline{X}^{t-1}, \underline{A}^{t-1})}. \end{aligned}$$

**Proof** Under the maximal global positivity assumption, we have  $\tilde{\mathcal{A}}_{adm}^t(\underline{a}^t, \underline{x}^{t-1}, \underline{a}^{t-1}) = \mathcal{A}_{adm}^t(\underline{x}^{t-1}, \underline{a}^{t-1}, a^t)$ . We can now insert this assumption into Eq. 21 which states the identification of  $J$  under global positivity:

$$\begin{aligned} \pi_{id}^t(A^t | \underline{A}^t, \underline{X}^{t-1}, \underline{A}^{t-1}) &= \mathbb{I}(A^t \in \tilde{\mathcal{A}}_{adm}^t(\underline{A}^t, \underline{X}^{t-1}, \underline{A}^{t-1})) f_{id}^t(\underline{A}^t, \underline{X}^{t-1}, \underline{A}^{t-1}) \\ &= \mathbb{I}(A^t \in \mathcal{A}_{adm}^t(\underline{X}^{t-1}, \underline{A}^{t-1}, A^t)) f_{id}^t(\underline{A}^t, \underline{X}^{t-1}, \underline{A}^{t-1}) \\ &= \mathbb{I}(A^t \geq a^t) \pi_{\beta}^t(A^t | \underline{X}^{t-1}, \underline{A}^{t-1}) f_{id}^{*t}(\underline{X}^{t-1}, \underline{A}^t, \underline{X}^{t-1}, \underline{A}^{t-1}) \end{aligned}$$

where we let  $f_{id}^{*t}$  denote another arbitrary function that ensures that  $\pi_{id}^t$  is a valid density. This concludes the proof for Corollary 10. ■

## Appendix I. Comparison of the Semi-offline RL IPW Estimator with Related Methods

In this appendix, we demonstrate that our proposed IPW estimator  $J_{IPW-Semi}$  is a more general version of an adapted version of the IPW estimator introduced by Caniglia et al. (2019). We refer to this estimator as  $\hat{J}_{IPW-Cen}$  since it is derived from a censoring viewpoint. While  $\hat{J}_{IPW-Cen}$  was developed for a scenario where both feature acquisition decisions and treatment decisions are made by the agent, it can be adapted to the AFA setting. However,  $\hat{J}_{IPW-Cen}$  is only applicable to simpler settings with one acquisition option per time-point

( $A^t \in \{0, 1\}$ ).  $\hat{J}_{IPW-Cen}$  is also only consistent if the maximal global positivity assumption holds, as will be shown.

Caniglia et al. (2019) derived  $J_{IPW-Cen}$  under the NDE and NUC assumptions. Instead of using a semi-offline sampling policy that avoids the acquisition of non-available features as proposed in this paper, Caniglia et al. (2019) simply sample from  $\pi_\alpha$ , even without knowledge about  $X_{(1)}$ . As the feature revelation is not possible if a non-available feature is acquired, they treat the resulting trajectory as censored. Known missing data methods are then applied to adjust for this censoring. Hence, in the wording of this paper, we would describe this viewpoint as an *online RL + censoring viewpoint*.

Adapted to the AFA setting under the consideration of deterministic AFA policies  $\pi_\alpha$ ,  $\hat{J}_{IPW-Cen}$  becomes:

$$\hat{J}_{IPW-Cen} = \hat{\mathbb{E}}_{n,uncen}[\hat{\rho}_{Cen}^T C(\pi_\alpha)] \text{ where } \hat{\rho}_{Cen}^T = \prod_{t=1}^T \left( \frac{\mathbb{I}(A^t = 1)}{\hat{\pi}_\beta^t(A^t = 1 | \underline{X}^{t-1}, \underline{A}^{t-1})} \right)^{A^t(\pi_\alpha)} \quad (44)$$

where  $\hat{\mathbb{E}}_{n',uncen}[\cdot]$  denotes the empirical average over the uncensored data points which have the known deterministic counterfactuals  $A^t_{(\pi_\alpha)}, X^t_{(\pi_\alpha)}$  and  $C(\pi_\alpha)$ .

Since  $A^t \in \{0, 1\}$ , it can be observed that the propensity score for a specific time-point  $t$  only appears in the factorization if the corresponding action is  $A^t_{(\pi_\alpha)} = 1$ . An "acquire nothing" AFA policy (where  $\pi_\alpha^t(A^t = 0 | \underline{X}^{t-1}, \underline{A}^{t-1}) = 1 \forall t$ ) would thus require no adjustment ( $\rho_{Cen}^T = 1$ ). In their example, this estimator achieved a 50-fold increase in data efficiency compared to the standard offline RL IPW estimator (Caniglia et al., 2019).

We establish the equivalence of our estimator  $\hat{J}_{IPW-Semi}$  and  $\hat{J}_{IPW-Cen}$  in the following proposition:

**Proposition 20.** (*Equivalence of  $\hat{J}_{IPW-Cen}$  and  $\hat{J}_{IPW-Semi}$* ). *The estimators  $\hat{J}_{IPW-Cen}$  and  $\hat{J}_{IPW-Semi}$  are equivalent for AFA settings with one action option per time-point, deterministic AFA policies  $\pi_\alpha$ , the maximal global positivity assumption (Assumption 6.4), and a simulation policy  $\pi_{sim} = \pi_\alpha$ .*

**Proof** Firstly, we clarify the blocking operation (from Definition 2) for this setting:

$$\pi'_{sim}(A^t = a^t | \underline{X}^{t-1}, \underline{A}^{t-1}, A^t = a^t) = \begin{cases} 1, & \text{if } a^t = 0 \ \& \ a^t = 0 \\ 0, & \text{if } a^t = 1 \ \& \ a^t = 0 \\ \pi_\alpha^t(A^t = a^t | \underline{X}^{t-1}, \underline{A}^{t-1}), & \text{if } a^t = 1. \end{cases}$$

The inverse probability weights of  $\hat{J}_{IPW-Semi}$  become:

$$\begin{aligned} \rho_{Semi}^T &= \prod_{t=1}^T \frac{\pi_\alpha^t(A^t = a^t | \underline{X}^{t-1}, \underline{A}^{t-1})}{\pi'_{sim}(A^t = a^t | \underline{X}^{t-1}, \underline{A}^{t-1}, A^t = a^t)} \frac{\mathbb{I}(A^t \geq a^t)}{\pi_\beta^t(A^t \geq a^t | \underline{X}^{t-1}, \underline{A}^{t-1})} \\ &\stackrel{*1}{=} \prod_{t=1}^T (\pi_\alpha^t(A^t = a^t | \underline{X}^{t-1}, \underline{A}^{t-1}))^{1-a^t} \left( \frac{\mathbb{I}(A^t = 1)}{\pi_\beta^t(A^t = 1 | \underline{X}^{t-1}, \underline{A}^{t-1})} \right)^{a^t} \\ &\stackrel{*2}{=} \prod_{t=1}^T \mathbb{I}(A^t = A^t_{(\pi_\alpha)}) \left( \frac{\mathbb{I}(A^t = 1)}{\pi_\beta^t(A^t = 1 | \underline{X}^{t-1}, \underline{A}^{t-1})} \right)^{A^t(\pi_\alpha)} \end{aligned}$$

where we used in \*1) the above definition of  $\pi'_{sim}$  and that  $\mathbb{I}(A^t \geq 0) = 1 = \pi_\beta(A^t \geq 0 | \underline{X}^{t-1}, \underline{A}^{t-1})$ . In \*2), we see that the first term corresponds to whether  $\pi_\alpha$  could be applied without running into censoring. It thus gives 0 weights to all datapoints where blocking occurred (i.e., which are censored under the online RL + censoring viewpoint). The second term then corresponds to the same weights as  $\rho_{Cen}^T$ , which concludes the proof for Proposition 20.  $\blacksquare$

We have demonstrated that, although the two estimators are derived from different concepts (online RL with censoring vs semi-offline RL), they are equal in this specific AFA setting of one action option per time-step, deterministic policies and under the maximal global positivity assumption. However, the key distinction lies in the generality of our estimator. Unlike  $\hat{J}_{IPW-Cen}$ , which is limited to the described setting, we developed an IPW estimator that can be applied for multiple acquisition options (i.e., higher dimensional  $A^t$ ) under the weaker global positivity assumption and in a modified version for static features settings as we show in our companion paper (von Kleist et al., 2023). It can further be combined with a Q-model to build the DRL estimator.

## Appendix J. Proof of Theorem 13

In this appendix, we proof Theorem 13, which we repeat here:

**Theorem 13.** *(Double robustness of  $\hat{J}_{DRL-Semi}$ ). The estimator  $\hat{J}_{DRL-Semi}$  is doubly robust, in the sense that it is consistent if either the Q-function  $\hat{Q}_{Semi}$  or the propensity score model  $\hat{\pi}_\beta$  is correctly specified.*

**Proof** To prove the double robustness property of the semi-offline RL version of the DRL estimator (i.e., Theorem 13), we decompose  $\hat{J}_{DRL-Semi}$  in two different ways:

*Scenario 1:* If  $\hat{\pi}_\beta$  is correctly specified, we find

$$\begin{aligned} \hat{J}_{DRL-Semi} &= \hat{\mathbb{E}}_n \left[ \hat{\mathbb{E}}_{n'} \left[ \rho_{Semi}^T C' + \sum_{t=1}^T \left( -\hat{\rho}_{Semi}^t \hat{Q}_{Semi}^t + \rho_{Semi}^{t-1} \hat{V}_{Semi}^{t-1} \right) \middle| A, X, Y \right] \right] \\ &= \underbrace{\mathbb{E}[\rho_{Semi}^T C']}_{=J} + \sum_{t=1}^T \underbrace{\mathbb{E} \left[ -\rho_{Semi}^t \hat{Q}_{Semi}^t + \rho_{Semi}^{t-1} \hat{V}_{Semi}^{t-1} \right]}_{=0}, \end{aligned}$$

where the first term is just the IPW estimator. As  $\hat{\pi}_\beta = \pi_\beta$  is correctly specified, the IPW estimator consistently estimates  $J$ . The fact that the second term equals 0 is shown in the



following:

$$\begin{aligned}
 & \mathbb{E} \left[ -\rho_{Semi}^t \hat{Q}_{Semi}^t + \rho_{Semi}^{t-1} \hat{V}_{Semi}^{t-1} \right] = \\
 & \stackrel{*1}{=} \mathbb{E} \left[ \rho_{Semi}^{t-1} \left( -\frac{\pi_{\alpha}^t(A^{t-1} | \underline{X}^{t-1}, \underline{A}^{t-1})}{\pi_{sim}^{t-1}(A^{t-1} | \underline{X}^{t-1}, \underline{A}^{t-1}, A^t)} \frac{\pi_{id}^t(A^t | \underline{A}^t, \underline{X}^{t-1}, \underline{A}^{t-1})}{\pi_{\beta}^t(A^t | \underline{X}^{t-1}, \underline{A}^{t-1})} \hat{Q}_{Semi}^t \right. \right. \\
 & \quad \left. \left. + \sum_{A^t} \pi_{\alpha}^t(A^t | \underline{X}^{t-1}, \underline{A}^{t-1}) \hat{Q}_{Semi}^t \right) \right] \\
 & \stackrel{*2}{=} \mathbb{E} \left[ \rho_{Semi}^{t-1} \left( -\sum_{A^t, A^t} \pi_{sim}^{t-1}(A^{t-1} | \underline{X}^{t-1}, \underline{A}^{t-1}, A^t) \pi_{\beta}^t(A^t | \underline{X}^{t-1}, \underline{A}^{t-1}) \frac{\pi_{\alpha}^t(A^t | \underline{X}^{t-1}, \underline{A}^{t-1})}{\pi_{sim}^{t-1}(A^t | \underline{X}^{t-1}, \underline{A}^{t-1}, A^t)} \right. \right. \\
 & \quad \left. \left. \cdot \frac{\pi_{id}^t(A^t | \underline{A}^t, \underline{X}^{t-1}, \underline{A}^{t-1})}{\pi_{\beta}^t(A^t | \underline{X}^{t-1}, \underline{A}^{t-1})} \hat{Q}_{Semi}^t + \sum_{A^t} \pi_{\alpha}^t(A^t | \underline{X}^{t-1}, \underline{A}^{t-1}) \hat{Q}_{Semi}^t \right) \right] \\
 & \stackrel{*3}{=} \mathbb{E} \left[ \rho_{Semi}^{t-1} \left( -\sum_{A^t} \pi_{\alpha}^t(A^t | \underline{X}^{t-1}, \underline{A}^{t-1}) \hat{Q}_{Semi}^t + \sum_{A^t} \pi_{\alpha}^t(A^t | \underline{X}^{t-1}, \underline{A}^{t-1}) \hat{Q}_{Semi}^t \right) \right] = 0
 \end{aligned}$$

with the following explanations:

- \*1): We use the relationship  $\hat{V}_{Semi}^{t-1} = \mathbb{E}_{\pi_{\alpha}}[\hat{Q}_{Semi}^t]$  and the decomposition of  $\rho_{Semi}$ .
- \*2): We use the fact that one can pull the expected value with respect to  $\pi_{sim}^{t-1}(A^{t-1} | \underline{X}^{t-1}, \underline{A}^{t-1}, A^t) \pi_{\beta}^t(A^t | \underline{X}^{t-1}, \underline{A}^{t-1})$  inside.
- \*3): We use the fact that  $\hat{Q}_{Semi}^t$  is independent of  $\pi_{id}^t(A^t | \underline{A}^t, \underline{X}^{t-1}, \underline{A}^{t-1})$  as long as it fulfills the positivity assumption.

*Scenario 2:* If  $\hat{Q}_{Semi}$  is correctly specified, we find

$$\begin{aligned}
 \hat{J}_{DRL-Semi} &= \mathbb{E}[V_{Semi}^0] + \mathbb{E}[\hat{\rho}_{Semi}^T (C' - Q_{Semi}^T)] + \mathbb{E} \left[ \sum_{t=1}^{T-1} \hat{\rho}_{Semi}^t (-Q_{Semi}^t + V_{Semi}^t) \right] \\
 &= \underbrace{\mathbb{E}[V_{Semi}^0]}_{=J} + \mathbb{E} \left[ \hat{\rho}_{Semi}^T \underbrace{(f_C(A', X', Y) - Q_{Semi}^T)}_{=0} \right] \\
 &+ \mathbb{E} \left[ \underbrace{\sum_{t=1}^{T-1} \hat{\rho}_{Semi}^t \left( -Q_{Semi}^t + \sum_{X^t} V_{Semi}^t p(X^t | \underline{X}^{t-1}, \underline{A}^t, \Xi) \right)}_{=0} \right]
 \end{aligned}$$

where  $\mathbb{E}[V_{Semi}^0]$  corresponds to the DM estimator which is consistent if  $\hat{Q}_{Semi} = Q_{Semi}$  is correctly specified. For the last term, we used that we can pull in the expected value with respect to the conditional distributions of  $X^t$ . The resulting term equals the first part of the semi-offline RL version of Bellman's equation. This concludes the proof of Theorem 13.  $\blacksquare$

## Appendix K. Estimation of Other Target Parameters from the Semi-offline RL View

In this appendix, we extend the target parameter to include time-dependent costs  $C^t$  (s.t.  $C = \bar{C}^1$ ). The newly defined target parameter becomes  $J = \mathbb{E} \left[ \sum_{t=1}^T C_{(\pi_\alpha)}^t \right]$ . In particular, these costs may include acquisition costs or misclassification costs for predictions at each time step. The acquisition costs are given by the known deterministic  $f_{C_a}^t(A^t)$ . When considering misclassification costs, we assume a per-step label  $Y^t$  to be available at each time step (s.t.  $Y = \bar{Y}^1$ ). The per-step misclassification costs can be computed by:

$$f_{C_{mc}}^t(Y^{*t}, Y^t) = f_{C_{mc}}^t(f_{cl}(A^t, \underline{X}^t), Y^t).$$

We combine both costs such that the target parameter is redefined to be:

$$\begin{aligned} J &= \mathbb{E} \left[ \sum_{t=1}^T \left( C_{a,(\pi_\alpha)}^t + C_{mc,(\pi_\alpha)}^t \right) \right] = \mathbb{E} \left[ \sum_{t=1}^T \left( f_{C_a}^t(A_{(\pi_\alpha)}^t) + f_{C_{mc}}^t(\underline{A}_{(\pi_\alpha)}^t, \underline{X}_{(\pi_\alpha)}^t, Y^t) \right) \right] \\ &\equiv \mathbb{E} \left[ \sum_{t=1}^T f_C^t(\underline{A}_{(\pi_\alpha)}^t, \underline{X}_{(\pi_\alpha)}^t, Y^t) \right] = \mathbb{E} \left[ \sum_{t=1}^T C_{(\pi_\alpha)}^t \right]. \end{aligned}$$

The reformulation, identification, and estimation steps from the semi-offline RL view can be extended to per-step costs. For this setting, we provide corollaries of the identification and estimation theorems from the main body. We do not provide additional proofs, as the extensions are straightforward.

### K.1 Identification

We start with a corollary that extends Theorem 7 for the per-step costs setting.

**Corollary 21.** *(Identification of  $J$  (for per-step costs) for the semi-offline RL view). The reformulated AFAPPE problem of estimating  $J$  (for per-step costs) under the semi-offline RL view (Eq. 17) is under Assumption 1 (no measurement noise), Assumption 2 (consistency), Assumption 3 (no interference), Assumption 4 (NDE), Assumption 5 (NUC) and Assumption 6.3 (global positivity) identified by*

$$J = \mathbb{E}_{p'} \left[ \sum_{t=1}^T C'_{(\pi_\alpha)} \right] = \sum_{A', A, G_A(X_{(1)}), Y} \sum_{t=1}^T f_C^t(\underline{A}^t, \underline{X}^t, Y^t) q'(A', A, X, Y) \quad (45)$$

where  $q'$  is given by Eq. 20.

Next, we continue with a corollary that extends Theorem 8 for the per-step costs setting.

**Corollary 22.** *(Bellman equation for semi-offline RL (for per-step costs)). The semi-offline RL view admits under Assumption 1 (no measurement noise), Assumption 2 (consistency), Assumption 3 (no interference), Assumption 4 (NDE), Assumption 5 (NUC) and*

the local positivity assumption at datapoint  $\underline{x}^{t-1}, \underline{a}^{t-1}, a^t$  (from Definition 4), the following semi-offline RL version of the Bellman equation for per-step costs:

$$Q_{Semi}(\underline{A}^t, \underline{X}^{t-1}, \underline{A}^{t-1}, \Xi) = \sum_{X^t} V_{Semi}(\underline{A}^t, \underline{X}^t, \underline{A}^{t-1}, A^t = a^t, \Xi) p(X^t | \underline{X}^{t-1}, \underline{A}^{t-1}, A^t = a^t, \Xi) \quad (46)$$

for any  $a^t \in \mathcal{A}_{adm}^t(\underline{X}^{t-1}, \underline{A}^{t-1}, A^t)$

$$\begin{aligned} V_{Semi}(\underline{A}^t, \underline{X}^t, \underline{A}^t, \Xi) &= \sum_{Y^t} f_C^t(\underline{X}^t, \underline{A}^t, Y^t) p(Y^t | \underline{X}^t, \underline{A}^t, \Xi) \\ &+ \sum_{A^{t+1}} Q_{Semi}(\underline{A}^{t+1}, \underline{X}^t, \underline{A}^t, \Xi) \pi_\alpha^{t+1}(A^{t+1} | \underline{X}^t, \underline{A}^t) \end{aligned} \quad (47)$$

with semi-offline RL versions of the state-action value function  $Q_{Semi}$  and state value function  $V_{Semi}$ :

$$\begin{aligned} Q_{Semi}^t &\equiv Q_{Semi}(\underline{A}^t, \underline{X}^{t-1}, \underline{A}^{t-1}, \Xi) \equiv \mathbb{E}_{p'} \left[ \sum_{\tau=t}^T C'_{(\bar{\pi}_\alpha^{t+1})} \middle| \underline{A}^t, \underline{X}^{t-1}, \underline{A}^{t-1}, \Xi \right] \\ V_{Semi}^t &\equiv V_{Semi}(\underline{A}^t, \underline{X}^t, \underline{A}^t, \Xi) \equiv \mathbb{E}_{p'} \left[ \sum_{\tau=t}^T C'_{(\bar{\pi}_\alpha^{t+1})} \middle| \underline{A}^t, \underline{X}^t, \underline{A}^t, \Xi \right]. \end{aligned}$$

Furthermore,  $Q_{Semi}^t$  and  $V_{Semi}^t$  are identified if the regional positivity assumption holds at  $\underline{A}^t, \underline{X}^{t-1}, \underline{A}^{t-1}$  and  $a^t \in \tilde{\mathcal{A}}_{adm}^t(\underline{A}^t, \underline{X}^{t-1}, \underline{A}^{t-1})$ .

## K.2 Estimation

The estimation formulas can be extended to the per-step setting as follows:

1) *Inverse probability weighting (IPW)*:

The target cost (for per-step costs) that is estimated by the semi-offline IPW estimator is

$$\hat{J}_{IPW-Semi} = \hat{\mathbb{E}}_n \left[ \hat{\mathbb{E}}_{n'} \left[ \sum_{t=1}^T \hat{\rho}_{Semi}^t C'^t | A, X, Y \right] \right],$$

with the same options for  $\rho_{Semi}^t$  as in the setting described in the main body.

2) *Direct method (DM)*:

The target cost (for per-step costs) that is estimated by the semi-offline DM estimator is

$$\hat{J}_{DM-Semi} = \hat{\mathbb{E}}_{n'} [\hat{V}_{Semi}^0]$$

with the adapted per-step cost version of  $V_{Semi}$  from Corollary 22.

3) *Double reinforcement learning (DRL)*:

The target cost (for per-step costs) that is estimated by the semi-offline DRL estimator is

$$\hat{J}_{DRL-Semi} = \hat{\mathbb{E}}_n \left[ \hat{\mathbb{E}}_{n'} \left[ \sum_{t=1}^T \left( \hat{\rho}_{Semi}^t C'^t - \hat{\rho}_{Semi}^t \hat{Q}_{Semi}^t + \hat{\rho}_{Semi}^{t-1} \hat{V}_{Semi}^{t-1} \right) | A, X, Y \right] \right].$$

with the adapted per-step cost version of  $V_{Semi}$  and  $Q_{Semi}$ .

## Appendix L. Derivation of the Missing Data Semiparametric Theory Approach

In this appendix, we show why Eq. 35 and Eq. 36 constitute valid choices for an element of the IPW space  $\Lambda_{IPW}$  and the augmentation space  $\Lambda_2$  but rely on the positivity assumption for missing data (Assumption 6.2). This is a simplified derivation from Tsiatis (2006).

### IPW space:

Clearly,  $h_{Miss} = \rho_{Miss} \mathbb{E}[C_{(\pi_\alpha)} | X_{(1)}, Y] - J$  is a function of the observed data, since  $\rho_{Miss}$  is non-zero only for complete cases. It thus remains to be shown that

$$\mathbb{E} [h_{Semi}(A, G_A(X_{(1)}), Y) | X_{(1)}, Y] = \varphi^F(X_{(1)}, Y).$$

This is shown in the following:

$$\begin{aligned} \mathbb{E}[h_{Semi}(A, G_A(X_{(1)}), Y) | X_{(1)}, Y] &= \mathbb{E} \left[ \rho_{Miss} \mathbb{E}[C_{(\pi_\alpha)} | X_{(1)}, Y] \Big| X_{(1)}, Y \right] - J \\ &= \mathbb{E} \left[ \prod_{t=1}^T \frac{\mathbb{I}(A^t = \vec{1})}{\pi_\beta^t(A^t = \vec{1} | \underline{X}_{(1)}^{t-1}, \underline{A}^{t-1} = \vec{1})} \mathbb{E}[C_{(\pi_\alpha)} | X_{(1)}, Y] \Big| X_{(1)}, Y \right] - J \\ &= \mathbb{E} \left[ \prod_{t=1}^T \frac{\mathbb{E}[\mathbb{I}(A^t = \vec{1}) | \underline{X}_{(1)}^{t-1}, \underline{A}^{t-1} = \vec{1}]}{\pi_\beta^t(A^t = \vec{1} | \underline{X}_{(1)}^{t-1}, \underline{A}^{t-1} = \vec{1})} \mathbb{E}[C_{(\pi_\alpha)} | X_{(1)}, Y] \Big| X_{(1)}, Y \right] - J \\ &= \mathbb{E} \left[ 1 \cdot \mathbb{E} \left[ C_{(\pi_\alpha)} \Big| X_{(1)}, Y \right] \Big| X_{(1)}, Y \right] - J \\ &= \mathbb{E} \left[ C_{(\pi_\alpha)} \Big| X_{(1)}, Y \right] - J = \varphi^F(X_{(1)}, Y). \end{aligned}$$

### Augmentation space $\Lambda_2$ :

To derive  $\Lambda_2$  under the missing data view, one redefines any function  $b(A, G_A(X_{(1)}), Y)$  using the fact that  $A$  is a categorical variable:

$$b(A, G_A(X_{(1)}), Y) = \sum_{a \in \mathcal{A}} \mathbb{I}(A = a) b_a(G_a(X_{(1)}), Y) \quad (48)$$

where  $b_a(G_a(X_{(1)}), Y)$  is any mean zero, finite variance function of  $G_a(X_{(1)}), Y$ . This allows the enforcement of the zero conditional mean condition that defines  $\Lambda_2$ :

$$\mathbb{E}[b(A, G_A(X_{(1)}), Y) | X_{(1)}, Y] = \sum_{a \in \mathcal{A}} p(A = a | G_a(X_{(1)}), Y) b_a(G_a(X_{(1)}), Y) = 0.$$

Under the missing data positivity assumption, one can now solve for  $b_{\vec{1}}$ :

$$b_{\vec{1}}(G_{\vec{1}}(X_{(1)}), Y) = - \frac{1}{p(A = \vec{1} | G_{\vec{1}}(X_{(1)}), Y)} \sum_{a \in \mathcal{A} \setminus \vec{1}} p(A = a | G_a(X_{(1)}), Y) b_a(G_a(X_{(1)}), Y)$$

Substituting  $b_{\vec{1}}(G_{\vec{1}}(X_{(1)}), Y)$  into Eq. 48 and applying the known factorization of the propensity score model for our AFA setting gives the desired space  $\Lambda_2$  consisting of all

$$\begin{aligned} b(A, G_A(X_{(1)}), Y) &= \sum_{a \in \mathcal{A} \setminus \vec{1}} \left[ \mathbb{I}(A = a) - \frac{\mathbb{I}(A = \vec{1})p(A = a|G_a(X_{(1)}), Y)}{p(A = \vec{1}|G_{\vec{1}}(X_{(1)}), Y)} \right] b_a(G_a(X_{(1)}), Y) \\ &= \sum_{a \in \mathcal{A} \setminus \vec{1}} \left[ \mathbb{I}(A = a) - \prod_{t=1}^T \frac{\mathbb{I}(A^t = \vec{1})\pi_\beta^t(A^t = a^t|G_{\underline{a}^{t-1}}(X_{(1)}), \underline{a}^{t-1})}{\pi_\beta^t(A^t = \vec{1}|G_{\underline{a}^{t-1}}(X_{(1)}), \underline{a}^{t-1})} \right] b_a(G_a(X_{(1)}), Y). \end{aligned}$$

## Appendix M. Proof of Lemma 15 and Lemma 16

In this appendix, we prove Lemma 15 and Lemma 16 which establish the equivalence of the missing data and offline RL semiparametric theory approaches to AFAPE under the NUC and NDE assumptions. We start with Lemma 15.

For ease of reference, we repeat it here:

**Lemma 15.** *(Relating the offline RL IPW estimator to the IPW space). The functional  $h_{\text{off}} \equiv \rho_{\text{off}}^T C - J$ , based on the IPW estimator from the offline RL view, is a valid element of the IPW space:  $h_{\text{off}} \in \Lambda_{\text{IPW}}$ .*

**Proof** We need to show that  $h_{\text{off}} = h_{\text{off}}(A, G_A(X_{(1)}), Y) = \rho_{\text{off}}^T C - J \in \Lambda_{\text{IPW}}$ . Clearly,  $h_{\text{off}}$  is a function of the observed data. It thus remains to be shown that

$$\mathbb{E} [h_{\text{off}}(A, G_A(X_{(1)}), Y) | X_{(1)}, Y] = \varphi^F(X_{(1)}, Y).$$

This is shown in the following:

$$\begin{aligned} \mathbb{E} [h_{\text{off}}(A, G_A(X_{(1)}), Y) | X_{(1)}, Y] &= \mathbb{E} \left[ \rho_{\text{off}}^T C - J \mid X_{(1)}, Y \right] \\ &= \mathbb{E} \left[ \prod_{t=1}^T \frac{\pi_\alpha^t(A^t | G_{\underline{A}^{t-1}}(X_{(1)}), \underline{A}^{t-1})}{\pi_\beta^t(A^t | G_{\underline{A}^{t-1}}(X_{(1)}), \underline{A}^{t-1})} f_C(Y, G_A(X_{(1)}), A) \mid X_{(1)}, Y \right] - J \\ &= \sum_{a \in \mathcal{A}} \prod_{t=1}^T \frac{\pi_\alpha^t(a^t | G_{\underline{a}^{t-1}}(X_{(1)}), \underline{a}^{t-1})}{\pi_\beta^t(a^t | G_{\underline{a}^{t-1}}(X_{(1)}), \underline{a}^{t-1})} f_C(Y, G_a(X_{(1)}), a) \pi_\beta^t(a^t | G_{\underline{a}^{t-1}}(X_{(1)}), \underline{a}^{t-1}) - J \\ &= \sum_{a \in \mathcal{A}} \prod_{t=1}^T \pi_\alpha^t(a^t | G_{\underline{a}^{t-1}}(X_{(1)}), \underline{a}^{t-1}) f_C(Y, G_a(X_{(1)}), a) - J = \varphi^F(X_{(1)}, Y) \end{aligned}$$

which completes the proof. ■

Next, we prove Lemma 16. We also repeat it here:

**Lemma 16.** *( $\Lambda_*$  is equal to the augmentation space). The augmentation space  $\Lambda_2$  is equal to  $\Lambda_*$ .*

We reuse properties about  $\Lambda_*$  shown by Liu et al. (2021) and do not repeat the corresponding proofs for these properties as they involve cumbersome notation.

**Proof** To demonstrate that  $\Lambda_* = \Lambda_2$ , we first introduce a new space, denoted as  $\Lambda_*^{AF}$ , containing all functions of  $b(A, X_{(1)}, Y)$  for which  $\mathbb{E}[b(A, X_{(1)}, Y)|X_{(1)}, Y] = 0$  holds:

$$\Lambda_*^{AF} \equiv \left\{ b^t(\underline{A}^{t-1}, X_{(1)}, Y) \left( \frac{A^t}{\pi_\beta^t} - 1 \right) : b^t(\underline{A}^{t-1}, X_{(1)}, Y) \in \mathcal{H}; t \in \{1, \dots, T\} \right\}.$$

The space  $\Lambda_*^{AF}$  indeed includes all functions  $b(A, X_{(1)}, Y)$  with mean zero and finite variance for which  $\mathbb{E}[b(A, X_{(1)}, Y)|X_{(1)}, Y] = 0$  holds. Specifically, the elements  $b^t(\underline{A}^{t-1}, X_{(1)}, Y)$  can represent any function of  $\underline{A}^{t-1}, X_{(1)}, Y$ . Since  $A^t$  is binary (taking values 0 or 1), adding a term  $\left( \frac{A^t}{c(\underline{A}^{t-1}, X_{(1)}, Y)} - 1 \right)$  (for some  $c(A, X_{(1)}, Y)$ ) generalizes the space to contain all functions of  $A^t, X_{(1)}, Y$ . The specific choice  $c(\underline{A}^{t-1}, X_{(1)}, Y) = \pi_\beta^t$  is enforced by the condition  $\mathbb{E}[b(A, X_{(1)}, Y)|X_{(1)}, Y] = 0$ :

$$\begin{aligned} \mathbb{E} \left[ b^t(\underline{A}^{t-1}, X_{(1)}, Y) \left( \frac{A^t}{\pi_\beta^t} - 1 \right) | X_{(1)}, Y \right] &= \\ &= \mathbb{E} \left[ b^t(\underline{A}^{t-1}, X_{(1)}, Y) \mathbb{E} \left[ \left( \frac{A^t}{\pi_\beta^t} - 1 \right) | X_{(1)}, Y, \underline{A}^{t-1} \right] | X_{(1)}, Y \right] \\ &= \mathbb{E} \left[ b^t(\underline{A}^{t-1}, X_{(1)}, Y) \underbrace{\left( \frac{\pi_\beta^t}{\pi_\beta^t} - 1 \right)}_0 | X_{(1)}, Y \right] = 0. \end{aligned}$$

Now, in order to find  $\Lambda_2$ , we must find the subspace of  $\Lambda_*^{AF}$ , which contains only the functions of the observed data. This is exactly  $\Lambda_*$  as was shown by Liu et al. (2021) (in the proof of Remark 13). In fact, they showed that there aren't observed data elements that are in  $\Lambda_*^{AF}$ , but not in  $\Lambda_*$ . They also showed that  $\Lambda_*$  only contains observed data elements. This concludes the proof.  $\blacksquare$

## Appendix N. Proof of Lemma 17

In this appendix, we prove Lemma 17. The proof follows a similar approach as for Lemma 15 for the offline RL IPW estimator. We repeat the lemma here for ease of reference.

**Lemma 17.** *(Relating the semi-offline RL IPW estimator to the IPW space).  $h_{Semi} \equiv h_{Semi}(A, G_A(X_{(1)}), Y) = \hat{\mathbb{E}}_{n'}[\rho_{Semi}^T C' | A, G_A(X_{(1)}), Y] - J$  is an element of the IPW space  $\Lambda_{IPW}$ .*

**Proof** We need to show that  $h_{Semi} = h_{Semi}(A, G_A(X_{(1)}), Y) = \hat{\mathbb{E}}_{n'}[\rho_{Semi}^T C' | A, G_A(X_{(1)}), Y] - J \in \Lambda_{IPW}$ . Clearly,  $h_{Semi}$  is a function of the observed data. It thus remains to be shown that

$$\mathbb{E} [h_{Semi}(A, G_A(X_{(1)}), Y) | X_{(1)}, Y] = \varphi^F(X_{(1)}, Y).$$

This is shown in the following:

$$\begin{aligned}
 \mathbb{E} \left[ h_{Semi}(A, G_A(X_{(1)}), Y) \middle| X_{(1)}, Y \right] &= \mathbb{E} \left[ \mathbb{E} \left[ \rho_{Semi}^T f_C(A', G_{A'}(X_{(1)}), Y) \middle| A, G_A(X_{(1)}), Y \right] - J \middle| X_{(1)}, Y \right] \\
 &= \mathbb{E} \left[ \prod_{t=1}^T \frac{\pi_\alpha^t(A^t | \underline{X}^{t-1}, \underline{A}^{t-1})}{\pi_{sim}^t(A^t | \underline{X}^{t-1}, \underline{A}^{t-1}, A^t)} \frac{\pi_{id}^t(A^t | \underline{A}^t, \underline{X}^{t-1}, \underline{A}^{t-1})}{\pi_\beta^t(A^t | \underline{X}^{t-1}, \underline{A}^{t-1})} f_C(A', G_{A'}(X_{(1)}), Y) \middle| X_{(1)}, Y \right] - J \\
 &= \sum_{a \in \mathcal{A}} \sum_{a' \in \mathcal{A}} \prod_{t=1}^T \frac{\pi_\alpha^t(a^t | \underline{X}^{t-1}, \underline{a}^{t-1})}{\pi_{sim}^t(a^t | \underline{X}^{t-1}, \underline{a}^{t-1}, a^t)} \frac{\pi_{id}^t(a^t | \underline{a}^t, \underline{X}^{t-1}, \underline{a}^{t-1})}{\pi_\beta^t(a^t | \underline{X}^{t-1}, \underline{a}^{t-1})} f_C(a', G_{a'}(X_{(1)}), Y) \\
 &\quad \cdot \pi_{sim}^t(a^t | \underline{X}^{t-1}, \underline{a}^{t-1}, a^t) \pi_\beta^t(a^t | \underline{X}^{t-1}, \underline{a}^{t-1}) - J \\
 &= \sum_{a \in \mathcal{A}} \sum_{a' \in \mathcal{A}} \prod_{t=1}^T \pi_\alpha^t(a^t | \underline{X}^{t-1}, \underline{a}^{t-1}) \pi_{id}^t(a^t | \underline{a}^t, \underline{X}^{t-1}, \underline{a}^{t-1}) f_C(a', G_{a'}(X_{(1)}), Y) - J \\
 &= \sum_{a' \in \mathcal{A}} \prod_{t=1}^T \pi_\alpha^t(a^t | \underline{X}^{t-1}, \underline{a}^{t-1}) f_C(a', G_{a'}(X_{(1)}), Y) - J \\
 &= \varphi^F(X_{(1)}, Y)
 \end{aligned}$$

which completes the proof. ■

## Appendix O. Derivation of the Influence Function for Semi-offline RL

In this Appendix, we provide the complete derivation of the projection of  $h_{Semi}$  onto  $\Lambda_{2, Semi}(\Xi)$ , thereby completing the proof for the class of influence functions under the semi-offline RL view, as proposed in Theorem 14.

We now show all in-between steps of the following equalities shown in Section 7:

$$\begin{aligned}
 \varphi_{Semi}(A, G_A(X_{(1)}), Y; \Xi) &= h_{Semi} - \Pi(h_{Semi} | \Lambda_{2, Semi}(\Xi)) \\
 &= h_{Semi} - \sum_{t=1}^T \mathbb{E} \left[ h_{Semi} \middle| A^t, \underline{A}^{t-1}, G_{\underline{A}^{t-1}}(X_{(1)}), \Xi \right] + \sum_{t=1}^T \mathbb{E} \left[ h_{Semi} \middle| \underline{A}^{t-1}, G_{\underline{A}^{t-1}}(X_{(1)}), \Xi \right] \\
 &\stackrel{*1}{=} \mathbb{E} \left[ \rho_{Semi}^T f_C(A', G_{A'}(X_{(1)}), Y) \middle| Y, G_A(X_{(1)}), A \right] \\
 &\quad - \sum_{t=1}^T \mathbb{E} \left[ \rho_{Semi}^t Q_{Semi}^t \middle| A^t, \underline{A}^{t-1}, G_{\underline{A}^{t-1}}(X_{(1)}), \Xi \right] \\
 &\quad + \sum_{t=1}^T \mathbb{E} \left[ \rho_{Semi}^{t-1} V_{Semi}^{t-1} \middle| \underline{A}^{t-1}, G_{\underline{A}^{t-1}}(X_{(1)}), \Xi \right] - J \\
 &= \mathbb{E} \left[ \rho_{Semi}^T f_C(A', G_{A'}(X_{(1)}), Y) - \sum_{t=1}^T \rho_{Semi}^t Q_{Semi}^t + \sum_{t=1}^T \rho_{Semi}^{t-1} V_{Semi}^{t-1} \middle| A, G_A(X_{(1)}), Y \right] - J.
 \end{aligned}$$

We now go into more detail about why \*1) holds. We begin with the term including  $Q_{Semi}$ :

$$\begin{aligned}
 & \mathbb{E} \left[ h_{Semi} \left| A^t, \underline{A}^{t-1}, G_{\underline{A}^{t-1}}(X_{(1)}), \Xi \right] = \right. \\
 &= \mathbb{E} \left[ \mathbb{E} \left[ \rho_{Semi}^T f_C(A', G_{A'}(X_{(1)}), Y) | Y, G_A(X_{(1)}), A \right] \left| A^t, \underline{A}^{t-1}, G_{\underline{A}^{t-1}}(X_{(1)}), \Xi \right] - J \right. \\
 &= \mathbb{E} \left[ \rho_{Semi}^T f_C(A', G_{A'}(X_{(1)}), Y) \left| A^t, \underline{A}^{t-1}, G_{\underline{A}^{t-1}}(X_{(1)}), \Xi \right] - J \right. \\
 &= \mathbb{E} \left[ \prod_{\tau=1}^T \frac{\pi_{\alpha}^{\tau}(A'^{\tau} | \underline{X}'^{\tau-1}, \underline{A}'^{\tau-1})}{\pi_{sim}^{\tau}(A'^{\tau} | \underline{X}'^{\tau-1}, \underline{A}'^{\tau-1}, A^{\tau})} \frac{\pi_{id}^{\tau}(A^{\tau} | \underline{A}'^{\tau}, \underline{X}^{\tau-1}, \underline{A}^{\tau-1})}{\pi_{\beta}^{\tau}(A^{\tau} | \underline{X}^{\tau-1}, \underline{A}^{\tau-1})} \right. \\
 &\quad \cdot f_C(A', G_{A'}(X_{(1)}), Y) \left| A^t, \underline{A}^{t-1}, G_{\underline{A}^{t-1}}(X_{(1)}), \Xi \right] - J \\
 &= \mathbb{E} \left[ \prod_{\tau=1}^t \frac{\pi_{\alpha}^{\tau}(A'^1 | \underline{X}'^{\tau-1}, \underline{A}'^{\tau-1})}{\pi_{sim}^{\tau}(A'^{\tau} | \underline{X}'^{\tau-1}, \underline{A}'^{\tau-1}, A^{\tau})} \frac{\pi_{id}^{\tau}(A^{\tau} | \underline{A}'^{\tau}, \underline{X}^{\tau-1}, \underline{A}^{\tau-1})}{\pi_{\beta}^{\tau}(A^{\tau} | \underline{X}^{\tau-1}, \underline{A}^{\tau-1})} \right. \\
 &\quad \cdot \mathbb{E} \left[ \prod_{\tau=t+1}^T \frac{\pi_{\alpha}^{\tau}(A'^{\tau} | \underline{X}'^{\tau-1}, \underline{A}'^{\tau-1})}{\pi_{sim}^{\tau}(A'^{\tau} | \underline{X}'^{\tau-1}, \underline{A}'^{\tau-1}, A^{\tau})} \frac{\pi_{id}^{\tau}(A^{\tau} | \underline{A}'^{\tau}, \underline{X}^{\tau-1}, \underline{A}^{\tau-1})}{\pi_{\beta}^{\tau}(A^{\tau} | \underline{X}^{\tau-1}, \underline{A}^{\tau-1})} \right. \\
 &\quad \cdot f_C(A', G_{A'}(X_{(1)}), Y) \left| A^t, \underline{A}^{t-1}, A^t, \underline{A}^{t-1}, G_{\underline{A}^{t-1}}(X_{(1)}), \Xi \right] \left| A^t, \underline{A}^{t-1}, G_{\underline{A}^{t-1}}(X_{(1)}), \Xi \right] - J \\
 &= \mathbb{E} \left[ \rho_{Semi}^t Q_{Semi}(\underline{A}^t, \underline{A}^{t-1}, G_{\underline{A}^{t-1}}(X_{(1)}), \Xi) \left| A^t, \underline{A}^{t-1}, G_{\underline{A}^{t-1}}(X_{(1)}), \Xi \right] - J \right. \\
 &= \mathbb{E} \left[ \rho_{Semi}^t Q_{Semi}^t \left| A, G_A(X_{(1)}), Y \right] - J \right.
 \end{aligned}$$

Similarly, we can show for the term including  $V_{Semi}$ :

$$\begin{aligned}
 & \mathbb{E} \left[ h_{Semi} \left| \underline{A}^{t-1}, G_{\underline{A}^{t-1}}(X_{(1)}), \Xi \right] = \right. \\
 &= \mathbb{E} \left[ \mathbb{E} \left[ \rho_{Semi}^T f_C(A', G_{A'}(X_{(1)}), Y) | Y, G_A(X_{(1)}), A \right] \left| \underline{A}^{t-1}, G_{\underline{A}^{t-1}}(X_{(1)}), \Xi \right] - J \right. \\
 &= \mathbb{E} \left[ \rho_{Semi}^T f_C(A', G_{A'}(X_{(1)}), Y) \left| \underline{A}^{t-1}, G_{\underline{A}^{t-1}}(X_{(1)}), \Xi \right] - J \right. \\
 &= \mathbb{E} \left[ \prod_{\tau=1}^T \frac{\pi_{\alpha}^{\tau}(A'^{\tau} | \underline{X}'^{\tau-1}, \underline{A}'^{\tau-1})}{\pi_{sim}^{\tau}(A'^{\tau} | \underline{X}'^{\tau-1}, \underline{A}'^{\tau-1}, A^{\tau})} \frac{\pi_{id}^{\tau}(A^{\tau} | \underline{A}'^{\tau}, \underline{X}^{\tau-1}, \underline{A}^{\tau-1})}{\pi_{\beta}^{\tau}(A^{\tau} | \underline{X}^{\tau-1}, \underline{A}^{\tau-1})} \right. \\
 &\quad \cdot f_C(A', G_{A'}(X_{(1)}), Y) \left| \underline{A}^{t-1}, G_{\underline{A}^{t-1}}(X_{(1)}), \Xi \right] - J \\
 &= \mathbb{E} \left[ \prod_{\tau=1}^{t-1} \frac{\pi_{\alpha}^{\tau}(A'^1 | \underline{X}'^{\tau-1}, \underline{A}'^{\tau-1})}{\pi_{sim}^{\tau}(A'^{\tau} | \underline{X}'^{\tau-1}, \underline{A}'^{\tau-1}, A^{\tau})} \frac{\pi_{id}^{\tau}(A^{\tau} | \underline{A}'^{\tau}, \underline{X}^{\tau-1}, \underline{A}^{\tau-1})}{\pi_{\beta}^{\tau}(A^{\tau} | \underline{X}^{\tau-1}, \underline{A}^{\tau-1})} \right. \\
 &\quad \cdot \mathbb{E} \left[ \prod_{\tau=t}^T \frac{\pi_{\alpha}^{\tau}(A'^{\tau} | \underline{X}'^{\tau-1}, \underline{A}'^{\tau-1})}{\pi_{sim}^{\tau}(A'^{\tau} | \underline{X}'^{\tau-1}, \underline{A}'^{\tau-1}, A^{\tau})} \frac{\pi_{id}^{\tau}(A^{\tau} | \underline{A}'^{\tau}, \underline{X}^{\tau-1}, \underline{A}^{\tau-1})}{\pi_{\beta}^{\tau}(A^{\tau} | \underline{X}^{\tau-1}, \underline{A}^{\tau-1})} \right. \\
 &\quad \cdot f_C(A', G_{A'}(X_{(1)}), Y) \left| \underline{A}^{t-1}, \underline{A}^{t-1}, G_{\underline{A}^{t-1}}(X_{(1)}), \Xi \right] \left| \underline{A}^{t-1}, G_{\underline{A}^{t-1}}(X_{(1)}), \Xi \right] - J \\
 &= \mathbb{E} \left[ \rho_{Semi}^{t-1} V_{Semi}(\underline{A}^{t-1}, \underline{A}^{t-1}, G_{\underline{A}^{t-1}}(X_{(1)}), \Xi) \left| \underline{A}^{t-1}, G_{\underline{A}^{t-1}}(X_{(1)}), \Xi \right] - J \right. \\
 &= \mathbb{E} \left[ \rho_{Semi}^{t-1} V_{Semi}^{t-1} \left| A, G_A(X_{(1)}), Y \right] - J \right.
 \end{aligned}$$



## Appendix P. Experiment Details

In this section, we describe the experiment setup in more detail. We also provide a detailed list of the parameters and configurations for each experiment in Tables 3 and 4.

### P.1 Data, Costs and Missingness Mechanisms

For the experiments, we defined a "superfeature" as a feature that comprises multiple subfeatures, which are acquired jointly and which have a single cost. Furthermore, we assumed a subset of features is available at no cost (free features) and set fixed acquisition costs  $c_{acq}$  for the remaining features. A prediction was to be performed at each time step, which corresponds to the setting described in Appendix K. We chose misclassification costs such that good policies must find a balance between the feature acquisition cost and the predictive value of the features.

We evaluated and compared the described methods on synthetic data sets with and without violation of either the NDE or NUC assumption. In experiments where the NDE assumption holds, the features are distributed according to:

$$X_{(1),i}^t = \begin{cases} \gamma_i X_{(1),i}^{t-1} + (1 - \gamma_i)\epsilon_i, & \text{if } t > 0 \\ \epsilon_i, & \text{if } t = 0. \end{cases}$$

where  $\epsilon_i \sim \mathcal{N}(0, \sigma)$ . In experiments with a violation of the NDE assumption, the unobserved variables  $U$  were distributed according to:

$$U_i^t = \begin{cases} \gamma_i U_i^{t-1} + (1 - \gamma_i)\epsilon_i + 0.5 \sum_i A_i^{t-1} & \text{if } t > 1 \\ \gamma_i U_i^{t-1} + (1 - \gamma_i)\epsilon_i, & \text{if } t = 1 \\ \epsilon_i, & \text{if } t = 0. \end{cases}$$

The labels are distributed according to

$$p(Y^t = 1) = \begin{cases} 1, & \text{if } \zeta_1 \sum_i W_i X_{(1),i}^t + \zeta_2 \sum_i W_i X_{(1),i}^{t-1} > 0 \\ 0.3, & \text{otherwise.} \end{cases}$$

This choice for  $Y$  simulates a scenario where not all data points are equally easy to classify.

The retrospective policy  $\pi_\beta$  follows different logistic models depending on whether a MAR assumption (NUC holds) or MNAR assumption (NUC is violated) is assumed, as specified in Table 3. To evaluate the convergence of different estimators when the NDE assumption holds, we consider the average cost of running the AFA agent on the data set over all data points in the ground truth test set (without missingness) as the true expected cost  $J$ . When NDE is violated, we sample the ground truth data generating process while running the agent and do so the same number of times as there are data points in the test set.

We performed five different experiments:

- **Experiment 1:** A standard experiment where NUC, NDE, and all three positivity assumptions (Assumptions 6.1, 6.2, and 6.4) hold.

- **Experiment 2:** The NUC and NDE assumptions hold, but the missing data positivity assumption (Assumption 6.2) is violated. This is achieved by reducing the number of complete cases to 0.007%.
- **Experiment 3:** The NUC and NDE assumptions hold, but the offline RL (Assumption 6.1) positivity assumption is violated. This is achieved by letting  $A_2$  be always 1 under  $\pi_\beta$ .
- **Experiment 4:** The NDE assumption holds, but the NUC assumption is violated. This corresponds to an MNAR missing data scenario.
- **Experiment 5:** The NUC assumption holds, but the NDE assumption is violated.

For full experiment configurations for the acquisition processes, please see Table 4.

## P.2 Training

We used an impute-then-regress classifier (Le Morvan et al., 2021) with unconditional mean imputation and a logistic regression classifier for the classification task and trained it on the available and further randomly subsampled data (where  $p(A_i^t = 1) = 0.5$ ). We tested random and fixed acquisition policies that acquire each costly feature with a 50% or 100% probability. Furthermore, we evaluated a proximal policy optimization (PPO) RL agent (Schulman et al., 2017), which was trained on the semi-offline sampling distribution  $p'$  using  $\pi_\alpha$  as the semi-offline sampling policy, but without adjustment for the blocking of actions. The data sets were split into training set (for the training of the agent and the classifier), nuisance function training set, and test set, where the estimators were evaluated. The splitting of the data set in a nuisance function training set and a test set is necessary due to the complexity of the used nuisance model functions classes (Kennedy, 2022). The resulting loss of efficiency may, however, be avoided using a cross-fitting approach (Kennedy, 2022).

<b>Data and environment</b>	
Sample size $n_D$	100'000 divided into 30% training set (for agent and classifier), 30% nuisance function training set, and 40% test set.
Superfeatures	super $X_0$ : $[X_0]$ , super $X_1$ : $[X_1]$ , super $X_2$ : $[X_2, X_3]$
Label	$Y^t \in \{0, 1\}$ and for $t \leq T = 3$ .
Data generation parameters	$\gamma_i = 0.2 \forall i$ , $\sigma = 1$ , $\zeta_1 = 1$ , $\zeta_2 = 0.3$ , $W = [1, 1, 2, 2]/6$
Feature acquisition cost	$c_{acq} = [0, 1, 1]$
Misclassification cost	$c_{mc} = 12$
<b>Models</b>	
Classifier	Logistic regression
Agents	Random 50%, Fixed 100%, PPO (learning rate: 0.0001, number of layers: 2, hidden layer neurons per layer: 64, hidden layer activation function: tanh)
Nuisance functions	$\hat{\pi}_\beta$ (logistic regression), $\hat{Q}_{Semi}$ ( $\Xi = \emptyset$ , learning rate: 0.001, number of layers: 2, hidden layer neurons per layer: 16, hidden layer activation function: ReLU)

Table 3: Full experiment details except for the acquisition process

<b>Missingness mechanisms</b>	
Exp 1	$p(A_0^t = 1) = 1.0,$ $p(A_1^t = 1) = \sigma(0.8 - 3.0X_0^{t-1} + 0.02X_1^{t-1} - 0.02X_2^{t-1}),$ $p(A_2^t = 1) = \sigma(0.8 - 3.0X_0^{t-1} + 0.02X_1^{t-1} - 0.02X_2^{t-1})$ Complete cases ratio: $p(A = \vec{1}) = 11.71\%$
Exp 2	$p(A_0^t = 1) = 1.0,$ $p(A_1^t = 1) = 0.2,$ $p(A_2^t = 1) = 0.2$ Complete cases ratio: $p(A = \vec{1}) = 0.007\%$
Exp 3	$p(A_0^t = 1) = 1.0,$ $p(A_1^t = 1) = 1.0,$ $p(A_2^t = 1) = \sigma(-0.5 - 2.0X_0^{t-1} - 0.1X_1^{t-1} - 0.1X_2^{t-1})$ Complete cases ratio: $p(A = \vec{1}) = 7.09\%$
Exp 4	$p(A_0^t = 1) = 1.0,$ $p(A_1^t = 1) = 1.0,$ $p(A_2^t = 1) = \sigma(-0.6 - 1.5X_{(1),2}^{t-1} - 1.5X_{(1),3}^{t-1})$ Complete cases ratio: $p(A = \vec{1}) = 9.63\%$
Exp 5	$p(A_0^t = 1) = 1.0,$ $p(A_1^t = 1) = \sigma(0.8 - 0.2X_0^{t-1} - 0.1X_1^{t-1} + 0.5X_2^{t-1}),$ $p(A_2^t = 1) = \sigma(0.8 - 0.2X_0^{t-1} - 0.1X_1^{t-1} + 0.5X_2^{t-1})$ Complete cases ratio: $p(A = \vec{1}) = 11.71\%$

Table 4: Acquisition process details for all five experiments

## References

- Chaojie An, Qifeng Zhou, and Shen Yang. A reinforcement learning guided adaptive cost-sensitive feature acquisition method. *Applied Soft Computing*, 117:108437, March 2022. ISSN 1568-4946. doi: 10.1016/j.asoc.2022.108437. URL <https://www.sciencedirect.com/science/article/pii/S1568494622000163>.
- Qi An, Hui Li, Xuejun Liao, and Lawrence Carin. Active feature acquisition with POMDP models. *Submitted to Pattern Recognition Letters*, 2006.
- Dimitri Bertsekas. *Dynamic programming and optimal control: Volume I*, volume 1. Athena scientific, 2012.
- Christian Beyer, Maik Büttner, Vishnu Unnikrishnan, Miro Schleicher, Eirini Ntoutsi, and Myra Spiliopoulou. Active feature acquisition on data streams under feature drift. *Annals of Telecommunications*, 75(9-10):597–611, October 2020. ISSN 0003-4347, 1958-9395. doi: 10.1007/s12243-020-00775-2. URL <https://link.springer.com/10.1007/s12243-020-00775-2>.
- Rohit Bhattacharya, Razieh Nabi, Ilya Shpitser, and James M. Robins. Identification In Missing Data Models Represented By Directed Acyclic Graphs. In *Proceedings of The 35th Uncertainty in Artificial Intelligence Conference*, pages 1149–1158. PMLR, August 2020. URL <https://proceedings.mlr.press/v115/bhattacharya20b.html>.
- Marc F. Botteman, Chris L. Pashos, Alberto Redaelli, Benjamin Laskin, and Robert Hauser. The health economics of bladder cancer. *PharmacoEconomics*, 21(18):1315–1330, December 2003. ISSN 1179-2027. doi: 10.1007/BF03262330. URL <https://doi.org/10.1007/BF03262330>.
- Ellen C. Caniglia, James M. Robins, Lauren E. Cain, Caroline Sabin, Roger Logan, Sophie Abgrall, Michael J. Mugavero, Sonia Hernández-Díaz, Laurence Meyer, Remonie Seng, Daniel R. Drozd, George R. Seage III, Fabrice Bonnet, Fabien Le Marec, Richard D. Moore, Peter Reiss, Ard van Sighem, William C. Mathews, Inma Jarrín, Belén Alejos, Steven G. Deeks, Roberto Muga, Stephen L. Boswell, Elena Ferrer, Joseph J. Eron, John Gill, Antonio Pacheco, Beatriz Grinsztejn, Sonia Napravnik, Sophie Jose, Andrew Phillips, Amy Justice, Janet Tate, Heiner C. Bucher, Matthias Egger, Hansjakob Furrer, Jose M. Miro, Jordi Casabona, Kholoud Porter, Giota Touloumi, Heidi Crane, Dominique Costagliola, Michael Saag, and Miguel A. Hernán. Emulating a trial of joint dynamic strategies: An application to monitoring and treatment of HIV-positive individuals. *Statistics in Medicine*, 38(13):2428–2446, 2019. ISSN 1097-0258. doi: 10.1002/sim.8120. URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/sim.8120>. eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/sim.8120>.
- Chun-Hao Chang, Mingjie Mai, and Anna Goldenberg. Dynamic Measurement Scheduling for Event Forecasting using Deep RL. In *Proceedings of the 36th International Conference on Machine Learning*, pages 951–960. PMLR, May 2019. URL <https://proceedings.mlr.press/v97/chang19a.html>.

- Li-Fang Cheng, Niranjani Prasad, and Barbara E. Engelhardt. An Optimal Policy for Patient Laboratory Tests in Intensive Care Units. In *Biocomputing 2019*, pages 320–331. WORLD SCIENTIFIC, October 2018. ISBN 978-981-327-981-0. doi: 10.1142/9789813279827\_0029. URL [https://www.worldscientific.com/doi/abs/10.1142/9789813279827\\_0029](https://www.worldscientific.com/doi/abs/10.1142/9789813279827_0029).
- Victor Chernozhukov, Denis Chetverikov, Mert Demirer, Esther Duflo, Christian Hansen, Whitney Newey, and James Robins. Double/debiased machine learning for treatment and structural parameters. *The Econometrics Journal*, 21(1):C1–C68, February 2018. ISSN 1368-4221. doi: 10.1111/ectj.12097. URL <https://doi.org/10.1111/ectj.12097>.
- Srijita Das, Rishabh Iyer, and Sriraam Natarajan. A Clustering based Selection Framework for Cost Aware and Test-time Feature Elicitation. In *8th ACM IKDD CODS and 26th COMAD*, pages 20–28, Bangalore India, January 2021. ACM. ISBN 978-1-4503-8817-7. doi: 10.1145/3430984.3431008. URL <https://dl.acm.org/doi/10.1145/3430984.3431008>.
- Miroslav Dudik, John Langford, and Lihong Li. Doubly Robust Policy Evaluation and Learning, May 2011. URL <http://arxiv.org/abs/1103.4601>. arXiv:1103.4601 [cs, stat].
- Gabriel Erion, Joseph D. Janizek, Carly Hudelson, Richard B. Utarnachitt, Andrew M. McCoy, Michael R. Sayre, Nathan J. White, and Su-In Lee. CoAI: Cost-Aware Artificial Intelligence for Health Care. Technical report, medRxiv, January 2021. URL <https://www.medrxiv.org/content/10.1101/2021.01.19.21249356v1>.
- US Preventive Services Task Force\*. Screening for breast cancer: US Preventive Services Task Force recommendation statement. *Annals of internal medicine*, 151(10):716–726, 2009. Publisher: American College of Physicians.
- Wenbo Gong, Sebastian Tschitschek, Sebastian Nowozin, Richard E Turner, José Miguel Hernández-Lobato, and Cheng Zhang. Icebreaker: Element-wise Efficient Information Acquisition with a Bayesian Deep Latent Gaussian Model. In *Advances in Neural Information Processing Systems*, 2019. URL <https://proceedings.neurips.cc/paper/2019/hash/c055dcc749c2632fd4dd806301f05ba6-Abstract.html>.
- John P Gould. Risk, stochastic preference, and the value of information. *Journal of Economic Theory*, 8(1):64–84, May 1974. ISSN 0022-0531. doi: 10.1016/0022-0531(74)90006-4. URL <https://www.sciencedirect.com/science/article/pii/0022053174900064>.
- James Hess. Risk and the Gain from Information. *Journal of Economic Theory*, 27(1): 231–238, 1982.
- Ronald W. Hilton. The Determinants of Cost Information Value: An Illustrative Analysis. *Journal of Accounting Research*, 17(2):411–435, 1979. ISSN 0021-8456. doi: 10.2307/2490511. URL <https://www.jstor.org/stable/2490511>. Publisher: [Accounting Research Center, Booth School of Business, University of Chicago, Wiley].

- Daniel G. Horvitz and Donovan J. Thompson. A generalization of sampling without replacement from a finite universe. *Journal of the American statistical Association*, 47(260): 663–685, 1952. Publisher: Taylor & Francis.
- Sheng-Jun Huang, Miao Xu, Ming-Kun Xie, Masashi Sugiyama, Gang Niu, and Songcan Chen. Active Feature Acquisition with Supervised Matrix Completion. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 1571–1579, July 2018. ISBN 978-1-4503-5552-0. doi: 10.1145/3219819.3220084. URL <https://dl.acm.org/doi/10.1145/3219819.3220084>.
- Jaromír Janisch, Tomáš Pevný, and Viliam Lisý. Classification with costly features as a sequential decision-making problem. *Machine Learning*, 109(8):1587–1615, August 2020. ISSN 0885-6125, 1573-0565. doi: 10.1007/s10994-020-05874-8. URL <http://link.springer.com/10.1007/s10994-020-05874-8>.
- Daniel Jarrett and Mihaela van der Schaar. Inverse Active Sensing: Modeling and Understanding Timely Decision-Making. *arXiv:2006.14141 [cs, stat]*, June 2020. URL <http://arxiv.org/abs/2006.14141>.
- Nathan Kallus and Masatoshi Uehara. Double Reinforcement Learning for Efficient Off-Policy Evaluation in Markov Decision Processes. *Journal of Machine Learning Research*, 21(167):1–63, 2020. ISSN 1533-7928. URL <http://jmlr.org/papers/v21/19-827.html>.
- Jeffrey M. Keisler, Zachary A. Collier, Eric Chu, Nina Sinatra, and Igor Linkov. Value of information analysis: the state of application. *Environment Systems and Decisions*, 34(1):3–23, March 2014. ISSN 2194-5411. doi: 10.1007/s10669-013-9439-4. URL <https://doi.org/10.1007/s10669-013-9439-4>.
- Edward H. Kennedy. Semiparametric doubly robust targeted double machine learning: a review. *arXiv:2203.06469 [stat]*, March 2022. URL <http://arxiv.org/abs/2203.06469>. arXiv: 2203.06469.
- Thomas T. Kok, Rachel M. Brouwer, Rene M. Mandl, Hugo G. Schnack, and Georg Kreml. Active Selection of Classification Features. *arXiv:2102.13636 [cs]*, February 2021. URL <http://arxiv.org/abs/2102.13636>. arXiv: 2102.13636.
- Murray D. Krahn, John E. Mahoney, Mark H. Eckman, John Trachtenberg, Stephen G. Pauker, and Allan S. Detsky. Screening for Prostate Cancer: A Decision Analytic View. *JAMA*, 272(10):773–780, September 1994. ISSN 0098-7484. doi: 10.1001/jama.1994.03520100035030. URL <https://doi.org/10.1001/jama.1994.03520100035030>.
- Noémi Kreif, Oleg Sofrygin, Julie A. Schmittiel, Alyce S. Adams, Richard W. Grant, Zheng Zhu, Mark J. van der Laan, and Romain Neugebauer. Exploiting nonsystematic covariate monitoring to broaden the scope of evidence about the causal effects of adaptive treatment strategies. *Biometrics*, 77(1):329–342, 2021. ISSN 1541-0420. doi: 10.1111/biom.13271. URL <https://onlinelibrary.wiley.com/doi/abs/10.1111/biom.13271>. eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1111/biom.13271>.

- Irving H. LaValle. On cash equivalents and information evaluation in decisions under uncertainty part I: Basic theory. *Journal of the American Statistical Association*, 63(321): 252–276, 1968a. Publisher: Taylor & Francis.
- Irving H. LaValle. On cash equivalents and information evaluation in decisions under uncertainty Part II: Incremental information decisions. *Journal of the American Statistical Association*, 63(321):277–284, 1968b. Publisher: Taylor & Francis.
- Marine Le Morvan, Julie Josse, Erwan Scornet, and Gael Varoquaux. What’s a good imputation to predict with missing values? In *Advances in Neural Information Processing Systems*, volume 34, pages 11530–11540, 2021. URL <https://proceedings.neurips.cc/paper/2021/hash/5fe8fdc79ce292c39c5f209d734b7206-Abstract.html>.
- Sergey Levine, Aviral Kumar, George Tucker, and Justin Fu. Offline Reinforcement Learning: Tutorial, Review, and Perspectives on Open Problems. *arXiv:2005.01643 [cs, stat]*, November 2020. URL <http://arxiv.org/abs/2005.01643>. arXiv: 2005.01643.
- Yang Li and Junier Oliva. Active Feature Acquisition with Generative Surrogate Models. In *Proceedings of the 38th International Conference on Machine Learning*, pages 6450–6459. PMLR, July 2021a. URL <https://proceedings.mlr.press/v139/li21p.html>. ISSN: 2640-3498.
- Yang Li and Junier B. Oliva. Dynamic Feature Acquisition with Arbitrary Conditional Flows. *arXiv:2006.07701 [cs, stat]*, March 2021b. URL <http://arxiv.org/abs/2006.07701>.
- Yang Li, Siyuan Shan, Qin Liu, and Junier B. Oliva. Towards Robust Active Feature Acquisition. *arXiv:2107.04163 [cs]*, July 2021. URL <http://arxiv.org/abs/2107.04163>.
- Lin Liu, Zach Shahn, James M. Robins, and Andrea Rotnitzky. Efficient Estimation of Optimal Regimes Under a No Direct Effect Assumption. *Journal of the American Statistical Association*, 116(533):224–239, January 2021. ISSN 0162-1459, 1537-274X. doi: 10.1080/01621459.2020.1856117. URL <https://www.tandfonline.com/doi/full/10.1080/01621459.2020.1856117>.
- Chao Ma, Sebastian Tschitschek, Konstantina Palla, Jose Miguel Hernandez-Lobato, Sebastian Nowozin, and Cheng Zhang. EDDI: Efficient Dynamic Discovery of High-Value Information with Partial VAE. In *Proceedings of the 36th International Conference on Machine Learning*, pages 4234–4243. PMLR, May 2019. URL <https://proceedings.mlr.press/v97/ma19c.html>.
- Karthika Mohan and Judea Pearl. Graphical Models for Processing Missing Data. *Journal of the American Statistical Association*, 116(534):1023–1037, April 2021. ISSN 0162-1459. doi: 10.1080/01621459.2021.1874961. URL <https://doi.org/10.1080/01621459.2021.1874961>. Publisher: Taylor & Francis eprint: <https://doi.org/10.1080/01621459.2021.1874961>.



- Karthika Mohan, Judea Pearl, and Jin Tian. Graphical Models for Inference with Missing Data. In *Advances in Neural Information Processing Systems*, volume 26, 2013. URL <https://proceedings.neurips.cc/paper/2013/hash/0ff8033cf9437c213ee13937b1c4c455-Abstract.html>.
- Susan A. Murphy. Optimal dynamic treatment regimes. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 65(2):331–355, 2003. URL <https://academic.oup.com/jrsssb/article-abstract/65/2/331/7092852>. Publisher: Oxford University Press.
- Alvin I. Mushlin and Lou Fintor. Is screening for breast cancer cost-effective? *Cancer*, 69(S7):1957–1962, 1992. ISSN 1097-0142. doi: 10.1002/1097-0142(19920401)69:7+⟨1957::AID-CNCR2820691716⟩3.0.CO;2-T. URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/1097-0142%2819920401%2969%3A7%20%3C1957%3A%3AAID-CNCR2820691716%3E3.0.CO%3B2-T>. \_eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/1097-0142%2819920401%2969%3A7%2B%3C1957%3A%3AAID-CNCR2820691716%3E3.0.CO%3B2-T>.
- Razieh Nabi, Rohit Bhattacharya, and Ilya Shpitser. Full law identification in graphical models of missing data: Completeness results. In *International Conference on Machine Learning*, pages 7153–7163. PMLR, 2020. URL <http://proceedings.mlr.press/v119/nabi20a.html>.
- Sriraam Natarajan, Srijita Das, Nandini Ramanan, Gautam Kunapuli, and Predrag Radivojac. On Whom Should I Perform this Lab Test Next? An Active Feature Elicitation Approach. In *IJCAI*, pages 3498–3505, 2018.
- Romain Neugebauer, Julie A. Schmittdiel, Alyce S. Adams, Richard W. Grant, and Mark J. van der Laan. Identification of the Joint Effect of a Dynamic Treatment Intervention and a Stochastic Monitoring Intervention Under the No Direct Effect Assumption. *Journal of Causal Inference*, 5(1):20160015, September 2017. ISSN 2193-3685, 2193-3677. doi: 10.1515/jci-2016-0015. URL <https://www.degruyter.com/document/doi/10.1515/jci-2016-0015/html>.
- Maya L Petersen, Kristin E Porter, Susan Gruber, Yue Wang, and Mark J van der Laan. Diagnosing and responding to violations in the positivity assumption. *Statistical methods in medical research*, 21(1):31–54, February 2012. ISSN 0962-2802. doi: 10.1177/0962280210386207. URL <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC4107929/>.
- Sophia M. Rein, Jing Li, Miguel Hernan, and Andrew Beam. Deep Learning Methods for the Noniterative Conditional Expectation G-Formula for Causal Inference from Complex Observational Data, October 2024. URL <http://arxiv.org/abs/2410.21531>. arXiv:2410.21531.
- James Robins. A new approach to causal inference in mortality studies with a sustained exposure period—application to control of the healthy worker survivor effect. *Mathematical Modelling*, 7(9):1393–1512, January 1986. ISSN 0270-0255.

doi: 10.1016/0270-0255(86)90088-6. URL <https://www.sciencedirect.com/science/article/pii/0270025586900886>.

James Robins, Liliana Orellana, and Andrea Rotnitzky. Estimation and extrapolation of optimal treatment and testing strategies. *Statistics in Medicine*, 27(23):4678–4721, 2008. ISSN 1097-0258. doi: 10.1002/sim.3301. URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/sim.3301>. [\\_eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/sim.3301](https://onlinelibrary.wiley.com/doi/pdf/10.1002/sim.3301).

James M. Robins. Optimal Structural Nested Models for Optimal Sequential Decisions. In P. Bickel, P. Diggle, S. Fienberg, U. Gather, I. Olkin, S. Zeger, D. Y. Lin, and P. J. Heagerty, editors, *Proceedings of the Second Seattle Symposium in Biostatistics*, volume 179, pages 189–326. Springer New York, New York, NY, 2004. ISBN 978-0-387-20862-6 978-1-4419-9076-1. doi: 10.1007/978-1-4419-9076-1\_11. URL [http://link.springer.com/10.1007/978-1-4419-9076-1\\_11](http://link.springer.com/10.1007/978-1-4419-9076-1_11). Series Title: Lecture Notes in Statistics.

Patrick Rockenschaub, Ela Marie Akay, Benjamin Gregory Carlisle, Adam Hilbert, Falk Meyer-Eschenbach, Anatol-Fiete Näher, Dietmar Frey, and Vince Istvan Madai. Generalisability of AI-based scoring systems in the ICU: a systematic review and meta-analysis, October 2023a. URL <https://www.medrxiv.org/content/10.1101/2023.10.11.23296733v1>. Pages: 2023.10.11.23296733.

Patrick Rockenschaub, Adam Hilbert, Tabea Kossen, Falk von Dincklage, Vince Istvan Madai, and Dietmar Frey. From Single-Hospital to Multi-Centre Applications: Enhancing the Generalisability of Deep Learning Models for Adverse Event Prediction in the ICU, April 2023b. URL <http://arxiv.org/abs/2303.15354>. arXiv:2303.15354 [cs].

Andrea Rotnitzky, James Robins, and Lucia Babino. On the multiply robust estimation of the mean of the g-functional, May 2017. URL <http://arxiv.org/abs/1705.08582>. arXiv:1705.08582 [stat].

Daniel O. Scharfstein, Andrea Rotnitzky, and James M. Robins. Adjusting for Nonignorable Drop-Out Using Semiparametric Nonresponse Models. *Journal of the American Statistical Association*, 94(448):1096–1120, 1999. ISSN 0162-1459. doi: 10.2307/2669923. URL <https://www.jstor.org/stable/2669923>. Publisher: [American Statistical Association, Taylor & Francis, Ltd.].

John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal Policy Optimization Algorithms, August 2017. URL <http://arxiv.org/abs/1707.06347>. arXiv:1707.06347 [cs].

Shaun R. Seaman and Ian R. White. Review of inverse probability weighting for dealing with missing data. *Statistical methods in medical research*, 22(3):278–295, 2013. Publisher: Sage Publications Sage UK: London, England.

Burr Settles. Active learning literature survey. 2009.

- Hajin Shim, Sung Ju Hwang, and Eunho Yang. Joint Active Feature Acquisition and Classification with Variable-Size Set Encoding. *Advances in Neural Information Processing Systems*, 31, 2018. URL <https://proceedings.neurips.cc/paper/2018/hash/e5841df2166dd424a57127423d276bbe-Abstract.html>.
- Ilya Shpitser, Karthika Mohan, and Judea Pearl. Missing data as a causal and probabilistic problem. Technical report, CALIFORNIA UNIV LOS ANGELES DEPT OF COMPUTER SCIENCE, 2015.
- Jonathan A. C. Sterne, Ian R. White, John B. Carlin, Michael Spratt, Patrick Royston, Michael G. Kenward, Angela M. Wood, and James R. Carpenter. Multiple imputation for missing data in epidemiological and clinical research: potential and pitfalls. *BMJ*, 338: b2393, June 2009. ISSN 0959-8138, 1468-5833. doi: 10.1136/bmj.b2393. URL <https://www.bmj.com/content/338/bmj.b2393>. Publisher: British Medical Journal Publishing Group Section: Research Methods & Reporting.
- Fengyi Tang, Lifan Zeng, Fei Wang, and Jiayu Zhou. Adversarial Precision Sensing with Healthcare Applications. In *2020 IEEE International Conference on Data Mining (ICDM)*, pages 521–530, November 2020. doi: 10.1109/ICDM50108.2020.00061.
- Morteza Tavakol, Salman Ashraf, and Sorin J. Brener. Risks and Complications of Coronary Angiography: A Comprehensive Review. *Global Journal of Health Science*, 4(1):65–93, January 2012. ISSN 1916-9736. doi: 10.5539/gjhs.v4n1p65. URL <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC4777042/>.
- Philip Thomas and Emma Brunskill. Data-Efficient Off-Policy Policy Evaluation for Reinforcement Learning. In *Proceedings of The 33rd International Conference on Machine Learning*, pages 2139–2148. PMLR, June 2016. URL <https://proceedings.mlr.press/v48/thomasa16.html>. ISSN: 1938-7228.
- Anastasios A. Tsiatis. *Semiparametric theory and missing data*. Springer series in statistics. Springer, New York, 2006. ISBN 978-0-387-32448-7.
- Stef van Buuren. Multiple imputation of discrete and continuous data by fully conditional specification. *Statistical Methods in Medical Research*, 16(3):219–242, June 2007. ISSN 0962-2802. doi: 10.1177/0962280206074463. URL <https://doi.org/10.1177/0962280206074463>.
- Thomas S. Verma and Judea Pearl. Equivalence and Synthesis of Causal Models. Technical Report R-150, Department of Computer Science, University of California, Los Angeles, 1990.
- Henrik von Kleist, Alireza Zamanian, Ilya Shpitser, and Narges Ahmidi. Evaluation of Active Feature Acquisition Methods for Time-varying Feature Settings, December 2023. URL <http://arxiv.org/abs/2312.01530>. arXiv:2312.01530 [cs, stat].
- Lan Wen, Jessica G. Young, James M. Robins, and Miguel A. Hernán. Parametric G-Formula Implementations for Causal Survival Analyses. *Biometrics*, 77(2):740–753, June

2021. ISSN 0006-341X. doi: 10.1111/biom.13321. URL <https://doi.org/10.1111/biom.13321>.
- Xiaoyong Chai, Lin Deng, Qiang Yang, and C. X. Ling. Test-cost sensitive naive Bayes classification. In *Fourth IEEE International Conference on Data Mining (ICDM'04)*, pages 51–58, November 2004. doi: 10.1109/ICDM.2004.10092.
- Haiyan Yin, Yingzhen Li, Sinno Jialin Pan, Cheng Zhang, and Sebastian Tschiatschek. Reinforcement Learning with Efficient Active Feature Acquisition. *arXiv:2011.00825 [cs]*, November 2020. URL <http://arxiv.org/abs/2011.00825>.
- Jinsung Yoon, William R. Zame, and Mihaela van der Schaar. Deep sensing: Active sensing using multi-directional recurrent neural networks. In *International Conference on Learning Representations*, 2018.
- Jinsung Yoon, James Jordon, and Mihaela van der Schaar. ASAC: Active Sensing using Actor-Critic models. In *Machine Learning for Healthcare Conference*, pages 451–473. PMLR, October 2019. URL <http://proceedings.mlr.press/v106/yon19a.html>. ISSN: 2640-3498.
- Pin Zhang. A novel feature selection method based on global sensitivity analysis with application in machine learning-based prediction model. *Applied Soft Computing*, 85: 105859, December 2019. ISSN 1568-4946. doi: 10.1016/j.asoc.2019.105859. URL <http://www.sciencedirect.com/science/article/pii/S1568494619306404>.
- Yan Zhou, Roderick J. A. Little, and John D. Kalbfleisch. Block-Conditional Missing at Random Models for Missing Data. *Statistical Science*, 25(4), November 2010. ISSN 0883-4237. doi: 10.1214/10-STS344. URL <https://projecteuclid.org/journals/statistical-science/volume-25/issue-4/Block-Conditional-Missing-at-Random-Models-for-Missing-Data/10.1214/10-STS344.full>.