

Towards a better understanding of incremental learning

Sanjay Jain^{1*} and Steffen Lange² and Sandra Zilles³

¹ School of Computing, National University of Singapore, Singapore 117543,
sanjay@comp.nus.edu.sg

² FB Informatik, Hochschule Darmstadt, Haardtring 100, D-64295 Darmstadt,
slange@fbi.h-da.de

³ DFKI GmbH, Erwin-Schrödinger-Str. 57, D-67663 Kaiserslautern, zilles@dfki.de

Abstract. The present study aims at insights into the nature of incremental learning in the context of Gold’s model of identification in the limit. With a focus on natural requirements such as consistency and conservativeness, incremental learning is analysed both for learning from positive examples and for learning from positive and negative examples. The results obtained illustrate in which way different consistency and conservativeness demands can affect the capabilities of incremental learners. These results may serve as a first step towards characterising the structure of typical classes learnable incrementally and thus towards elaborating uniform incremental learning methods.

1 Introduction

Considering data mining tasks, where specific knowledge has to be induced from a huge amount of more or less unstructured data, several approaches have been studied empirically in machine learning and formally in the field of learning theory. These approaches differ in terms of the form of interaction between the learning machine and its environment. For instance, scenarios have been analysed, where the learner receives instances of some target concept to be identified or where the learner may pose queries concerning the target concept [6, 2, 11]. For learning from examples, one critical aspect is the limitation of a learning machine in terms of its memory capacity. In particular, if huge amounts of data have to be processed, it is conceivable that this capacity is too low to memorise all relevant information during the whole learning process. This has motivated the analysis of so-called *incremental learning*, cf. [4, 5, 7–9, 12], where in each step of the learning process, the learner has access only to a limited number of examples. Thus, in each step, its hypothesis can be built upon these examples and its former hypothesis, only. Other examples seen before have to be ‘forgotten’.

It has been analysed how such constraints affect the capabilities of learning machines, thus revealing models in which certain classes of target concepts are learnable, but not learnable in an incremental manner. However, some quite

* Supported in part by NUS grant number R252-000-127-112 and R252-000-212-112.

natural constraints for successful learning have mainly been neglected in the corresponding studies. These constraints are (a) the requirement for *consistent* learning, i. e., the demand that none of the intermediate hypotheses a learner explicates should contradict the data processed so far, and (b) the requirement for *conservative* learning, i. e., the demand that each intermediate hypothesis should be maintained as long as it is consistent with the data seen.

The fact that there is no comprehensive analysis of how these demands affect the capabilities of incremental learners can be traced back to a lack of knowledge about the nature of incremental learning. In particular, there is no formal basis explaining typical or uniform ways for solving learning tasks in an incremental way. In terms of learning theory, incremental learning is one of the very few models, for which no characterisation of the typical structure of learnable classes is known. For other models of learning from examples, characterisations and uniform learning methods have often been the outcome of analysing the impact of consistency or conservativeness, see, e. g., [13]. Thus, also in the context of incremental learning, it is conceivable that studying these natural requirements may yield insights into typical learning methods. In other words, analysing consistency and conservativeness may be the key for a better understanding of the nature of incremental learning and may thus, in the long term, provide characterisations of learnable classes and uniform incremental learning methods.

The present study aims at insights into the nature of incremental learning in the context of Gold’s model of learning in the limit from examples [6]. For that purpose, we analyse Wiehagen’s version of incremental learning, namely *iterative learning* [12] with a focus on consistent and conservative learners. In Gold’s approach, learning is considered as an infinite process, where in each step the learner is presented an example e_n for the target concept and is supposed to return an intermediate hypothesis. In the limit, the hypotheses must stabilise on a correct representation of the target concept. Here, in step $n + 1$ of the learning process, the learner has access to all examples e_0, \dots, e_n provided up to step n plus the current example e_{n+1} . In contrast, an iterative learner has no capacities for memorising any examples seen so far, i. e., its hypothesis h_{n+1} in step $n + 1$ is built only upon the example e_{n+1} and its previous hypothesis h_n .

The present paper addresses consistency and conservativeness in the context of iterative learning. Here several possible ways to formalise the demands for consistency and conservativeness become apparent. Assume an iterative learner has processed the examples e_0, \dots, e_{n+1} for some target concept and returns some hypothesis h_{n+1} in step $n + 1$. From a global perspective, one would define h_{n+1} consistent, if it agrees with the examples e_0, \dots, e_{n+1} . But since the learner has not memorised e_0, \dots, e_n , it might be considered natural to just demand that h_{n+1} agrees with the current example e_{n+1} . This is justified from a rather local perspective. Similarly, when defining conservativeness from a global point of view, one might demand that $h_{n+1} = h_n$ in case h_n does not contradict any of the examples e_0, \dots, e_{n+1} , whereas a local variant of conservativeness would mean to require that $h_{n+1} = h_n$ in case h_n does not contradict the current example e_{n+1} .

Note that local consistency is a weaker requirement than global consistency, whereas local conservativeness is stronger than global conservativeness.

In the present paper, we restrict our focus on recursive languages as target concepts [1, 13]. In particular, the target classes are required to be indexable, i. e., there exist algorithms deciding the membership problem uniformly for all possible target languages. This restriction is motivated by the fact that many classes of target concepts relevant for typical learning tasks are indexable.

The paper is structured as follows. In Section 2, we provide the definitions and notations necessary for our formal analysis. Then Section 3 is concerned with a case study of iterative learning of regular erasing pattern languages – a quite natural and simple to define indexable class which has shown to be suitable for representing target concepts in many application scenarios. This case study shows how consistency and conservativeness may affect the learnability of such pattern languages in case quite natural hypothesis spaces are chosen for learning. Section 4 focuses on consistency in iterative learning. It has turned out, that iterative learners can be normalised to work in a locally consistent way, whereas global consistency is a constraint reducing the capabilities of iterative learners. Both results hold for learning from positive examples as well as for learning from both positive and negative examples. Section 5 then is concerned with conservativeness. Here we show that, in the scenario of learning from only positive examples, the effects of global conservativeness demands and local conservativeness demands are equal, as far as the capabilities of iterative learners are concerned. In contrast to that there are classes which can be learned iteratively from positive and negative examples by a globally conservative learner, but not in a locally conservative manner. Concerning the effect of weak conservativeness demands (i. e., of global conservativeness), we can show that they strictly reduce the capabilities of iterative learners which are given both positive and negative examples as information. However, the corresponding comparison in the case of learning from only positive examples is still open. In our point of view, not only the mere results presented here, but in particular the proof constructions and separating classes give an impression of characteristic methods of iterative learning and characteristic properties of iteratively learnable classes, even though we cannot provide a formal characterisation yet. Section 6 contains our conclusions.

2 Preliminaries

Let Σ be a fixed finite alphabet, Σ^* the set of all finite strings over Σ , and Σ^+ its subset excluding the empty string. $|w|$ denotes the length of a string w . Any non-empty subset of Σ^* is called a *language*. For any language L , $co(L) = \Sigma^* \setminus L$. \mathbb{N} is the set of all natural numbers. If L is a language, then any infinite sequence $t = (w_j)_{j \in \mathbb{N}}$ with $\{w_j \mid j \in \mathbb{N}\} = L$ is called a *text* for L . Moreover, any infinite sequence $i = ((w_j, b_j))_{j \in \mathbb{N}}$ over $\Sigma^* \times \{+, -\}$ such that $\{w_j \mid j \in \mathbb{N}\} = \Sigma^*$, $\{w_j \mid j \in \mathbb{N}, b_j = +\} = L$, and $\{w_j \mid j \in \mathbb{N}, b_j = -\} = co(L)$ is referred to as an *informant* for L . Then, for any $n \in \mathbb{N}$, $t[n]$ and $i[n]$ denote the initial segment of t and i of length $n + 1$, while $t(n) = w_n$ and $i(n) = (w_n, b_n)$.

Furthermore, $\text{content}(t[n]) = \{w_j \mid j \leq n\}$. Let $\text{content}(i[n])$, $\text{content}^+(i[n])$, and $\text{content}^-(i[n])$ denote the sets $\{(w_j, b_j) \mid j \leq n\}$, $\{w_j \mid j \leq n, b_j = +\}$, and $\{w_j \mid j \leq n, b_j = -\}$.

A family $(L_j)_{j \in \mathbb{N}}$ of languages is called an *indexing* for a class \mathcal{C} of recursive languages, if $\mathcal{C} = \{L_j \mid j \in \mathbb{N}\}$ and there is a recursive function f such that $L_j = \{w \in \Sigma^* \mid f(j, w) = 1\}$ for all $j \in \mathbb{N}$. \mathcal{C} is called an *indexable class* (of recursive languages), if \mathcal{C} possesses an indexing.

In our proofs, we will use a fixed Gödel numbering $(\varphi_j)_{j \in \mathbb{N}}$ of all (and only all) partial recursive functions over \mathbb{N} as well as an associated complexity measure $(\Phi_j)_{j \in \mathbb{N}}$, see [3]. Then, for $k, x \in \mathbb{N}$, φ_k is the partial recursive function computed by program k and we write $\varphi_k(x) \downarrow$ ($\varphi_k(x) \uparrow$), if $\varphi_k(x)$ is defined (undefined).

2.1 Learning from text

Let \mathcal{C} be an indexable class, $\mathcal{H} = (L_j)_{j \in \mathbb{N}}$ any indexing of some $\mathcal{C}' \supseteq \mathcal{C}$ (called *hypothesis space*), and $L \in \mathcal{C}$. An *inductive inference machine* (IIM for short) M is an algorithmic device that reads longer and longer initial segments σ of a text and outputs numbers $M(\sigma)$ as its hypotheses. An IIM M returning some j is construed to hypothesise the language L_j . The following definition of learning from positive data is based on Gold [6]. Given a text t for L , M *learns L from t with respect to \mathcal{H}* , if (a) the sequence of hypotheses output by M , when fed t , stabilises on a number j (* i. e., past some point M always outputs the hypothesis j *) and (b) this number j fulfils $L_j = L$.

An *iterative inductive inference machine* is only allowed to use its previous hypothesis and the current string in a text for computing its current hypothesis. More formally, an *iterative IIM* M is an algorithmic device that maps elements from $\mathbb{N} \cup \{\text{init}\} \times \Sigma^*$ into \mathbb{N} , where *init* denotes a fixed initial ‘hypothesis’ which the IIM may never output. Let $t = (w_n)_{n \in \mathbb{N}}$ be any text for some language $L \subseteq \Sigma^*$. Then we denote by $(M[\text{init}, t[n]])_{n \in \mathbb{N}}$ the sequence of hypotheses generated by M when processing t , i. e., $M[\text{init}, w_0] = M(\text{init}, w_0)$ and, for all $n \in \mathbb{N}$, $M[\text{init}, t[n+1]] = M(M[\text{init}, t[n]], w_{n+1})$.

Definition 1. [12] *Let \mathcal{C} be an indexable class, $\mathcal{H} = (L_j)_{j \in \mathbb{N}}$ a hypothesis space, and $L \in \mathcal{C}$. An iterative IIM M learns L from text with respect to \mathcal{H} iff, for any text $t = (w_n)_{n \in \mathbb{N}}$ for L , the sequence $(M[\text{init}, t[n]])_{n \in \mathbb{N}}$ stabilises on a number j with $L_j = L$. Moreover, M learns \mathcal{C} from text with respect to \mathcal{H} , if it identifies every $L' \in \mathcal{C}$ from text with respect to \mathcal{H} . Finally, ItTxt denotes the collection of all indexable classes \mathcal{C}' for which there is a hypothesis space \mathcal{H}' and an iterative IIM learning \mathcal{C}' from text with respect to \mathcal{H}' .*

In the definition of *consistent learning*, a hypothesis of a learner is said to be consistent, if it reflects the data it was built upon correctly. Since an iterative IIM M , when processing some text t , is only allowed to use its previous hypothesis, say $L_{j'}$, and the current string v in t for computing its current hypothesis L_j , it is quite natural to distinguish two variants of consistent learning. In the first case, it is demanded that L_j contains all elements of t seen so far, while, in the second case, it is only required that L_j contains the string v .

Definition 2. Let \mathcal{C} be an indexable class, $\mathcal{H} = (L_j)_{j \in \mathbb{N}}$ a hypothesis space, and M an iterative IIM. M is globally (locally) consistent for \mathcal{C} iff $\text{content}(t[n]) \subseteq L_{M[\text{init}, t[n]]}$ ($t(n) \in L_{M[\text{init}, t[n]]}$) for every text segment $t[n]$ for some $L \in \mathcal{C}$. Finally, ItGConsTxt (ItLConsTxt) denotes the collection of all indexable classes \mathcal{C}' for which there is a hypothesis space \mathcal{H}' and an iterative IIM which is globally (locally) consistent for \mathcal{C}' and learns \mathcal{C}' from text with respect to \mathcal{H}' .

Finally we consider *conservative learning*. Informally speaking, a conservative learner maintains its current hypothesis as long as the latter does not contradict any data seen. Hence, whenever a conservative IIM changes its recent hypothesis, this must be justified by data having occurred which prove an inconsistency of its recent hypothesis. Similarly to the case of consistent iterative learning, it is quite natural to distinguish two variants of conservativeness.

Definition 3. Let \mathcal{C} be an indexable class, $\mathcal{H} = (L_j)_{j \in \mathbb{N}}$ be a hypothesis space, and M be an iterative IIM. M is globally (locally) conservative for \mathcal{C} iff, for every text segment $t[n+1]$ for some $L \in \mathcal{C}$, $M[\text{init}, t[n+1]] \neq M[\text{init}, t[n]]$ implies $\text{content}(t[n+1]) \not\subseteq L_{M[\text{init}, t[n]]}$ (implies $t(n+1) \notin L_{M[\text{init}, t[n]]}$). Finally, ItGConvTxt (ItLConvTxt) denotes the collection of all indexable classes \mathcal{C}' for which there is a hypothesis space \mathcal{H}' and an iterative IIM which is globally (locally) conservative for \mathcal{C}' and learns \mathcal{C}' from text with respect to \mathcal{H}' .

Note that we allow a mind change from *init* after the first input data is received.

2.2 Learning from informant

For all variants of *ItTxt* considered so far we define corresponding models capturing the case of learning from informant. Now an iterative IIM M maps $\mathbb{N} \times (\Sigma^* \times \{+, -\})$ into \mathbb{N} . Let $i = (w_n, b_n)_{n \in \mathbb{N}}$ be any informant for some language L , and let *init* be a fixed initial hypothesis. Then $(M[\text{init}, i[n]])_{n \in \mathbb{N}}$ is the sequence of hypotheses by M processing i , i. e., $M[\text{init}, (w_0, b_0)] = M(\text{init}, (w_0, b_0))$ and, for all $n \in \mathbb{N}$, $M[\text{init}, i[n+1]] = M(M[\text{init}, i[n]], (w_{n+1}, b_{n+1}))$.

Definition 4. [12] Let \mathcal{C} be an indexable class, $\mathcal{H} = (L_j)_{j \in \mathbb{N}}$ a hypothesis space, and $L \in \mathcal{C}$. An iterative IIM M learns L from informant with respect to \mathcal{H} , iff for every informant i for L , the sequence $(M[\text{init}, i[n]])_{n \in \mathbb{N}}$ stabilises on a number j with $L_j = L$. Moreover, M learns \mathcal{C} from informant with respect to \mathcal{H} , if M learns every $L' \in \mathcal{C}$ from informant with respect to \mathcal{H} .

The notion *ItInf* is defined similarly to the text case. Now also the consistency and conservativeness demands can be formalised. For instance, for consistency, let \mathcal{C} be an indexable class, $\mathcal{H} = (L_j)_{j \in \mathbb{N}}$ a hypothesis space, and M an iterative IIM. M is globally (locally) consistent for \mathcal{C} iff $\text{content}^+(i[n]) \subseteq L_{M[\text{init}, i[n]]}$ and $\text{content}^-(i[n]) \subseteq \text{co}(L_{M[\text{init}, i[n]]})$ ($b = +$ for $w \in L_{M[\text{init}, i[n]]}$ and $b = -$ for $w \notin L_{M[\text{init}, i[n]]}$) for every informant segment $i[n]$ for some $L \in \mathcal{C}$, where $i(n) = (w, b)$. Finally, the definitions of *ItGConsInf*, *ItLConsInf*, *ItGConvInf*, *ItLConvInf* can be adapted from the text case to the informant case.

3 A case study: The regular erasing pattern languages

Let Σ be any fixed finite alphabet. Let $X = \{x_1, x_2, \dots\}$ be an infinite set of variables, disjoint with Σ . A *regular pattern* α is a string from $(\Sigma \cup X)^+$ which contains every variable at most once. Let α be a regular pattern. Then $L_\varepsilon(\alpha)$, the *regular erasing pattern language* generated by α , contains all strings in Σ^* that can be obtained by replacing the variables in α by strings from Σ^* , see, e. g., [10]. Note that $L_\varepsilon(\alpha)$ constitutes a regular language. Subsequently, let \mathcal{C}_{rp} denote the collection of all regular erasing pattern languages.

Our first result can be achieved by adapting a standard idea, see, e. g., [4].

Theorem 1. *There is a learner witnessing both $\mathcal{C}_{\text{rp}} \in \text{ItGConsTxt}$ and $\mathcal{C}_{\text{rp}} \in \text{ItLConvTxt}$.*

Sketch of the proof. Let $(D_j)_{j \in \mathbb{N}}$ be the canonical enumeration of all finite subsets of \mathbb{N} and $(L_\varepsilon(\alpha_j))_{j \in \mathbb{N}}$ be an effective, repetition-free indexing of \mathcal{C}_{rp} . Moreover let $L'_j = \bigcap_{z \in D_j} L_\varepsilon(\alpha_z)$. Hence $(L'_j)_{j \in \mathbb{N}}$ is an indexing comprising the class \mathcal{C}_{rp} . The proof is essentially based on the following fact.

Fact 1. *There is an algorithm A which, given any string $w \in \Sigma^+$ as input, outputs an index j such that $D_j = \{z \in \mathbb{N} \mid w \in L_\varepsilon(\alpha_z)\}$.*

A learner M witnessing $\mathcal{C}_{\text{rp}} \in \text{ItGConsTxt}$ and $\mathcal{C}_{\text{rp}} \in \text{ItLConvTxt}$ with respect to $(L'_j)_{j \in \mathbb{N}}$ may simply work as follows:

Initially, if the first string w appears, M starts its subroutine A , determines $j = A(w)$, and guesses the language L'_j , i. e., $M(\text{init}, w) = j$. Next M , when receiving a new string v , refines its recent hypothesis, say j' , as follows. M determines the canonical index j of the set $\{z \mid z \in D_{j'}, v \in L_\varepsilon(\alpha_z)\} \subseteq D_{j'}$ and guesses the languages L'_j , i. e., $M(j', v) = j$.

It is not hard to see that M learns as required. \square

Although the iterative learner M used in this proof is locally conservative and globally consistent, M has the disadvantage of guessing languages not contained in the class of all regular erasing pattern languages. At first glance, it might seem that this weakness can easily be compensated, since the final guess returned by M is always a regular erasing pattern language and, moreover, one can effectively determine whether or not the recent guess of M equals a regular erasing pattern language. Surprisingly, even under this quite ‘perfect’ circumstances, it is impossible to replace M by an iterative, locally conservative, and globally consistent learner for \mathcal{C}_{rp} that hypothesises languages in \mathcal{C}_{rp} , exclusively.

Theorem 2. *Let $\text{card}(\Sigma) \geq 2$. Let $(L_j)_{j \in \mathbb{N}}$ be any indexing of \mathcal{C}_{rp} . Then there is no learner M witnessing both $\mathcal{C}_{\text{rp}} \in \text{ItGConsTxt}$ and $\mathcal{C}_{\text{rp}} \in \text{ItLConvTxt}$ with respect to $(L_j)_{j \in \mathbb{N}}$.*

Proof. Let $\{a, b\} \subseteq \Sigma$. Assume to the contrary that there is an iterative learner M which learns \mathcal{C}_{rp} locally conservatively and globally consistently, hypothesising only regular erasing pattern languages. Consider M for any text of some $L \in \mathcal{C}_{\text{rp}}$ with the initial segment $\sigma = aba, aab$. Since M must avoid overgeneralised

hypotheses, there are only two possible semantically different hypotheses which are globally consistent with σ , namely x_1abx_2 and ax_1ax_2 . Distinguish two cases:

Case (a). $L_{M[init,\sigma]} = L_\varepsilon(x_1abx_2)$.

Consider M processing $\sigma_1 = \sigma ab, aa$ and $\sigma_2 = \sigma aa$. Since $ab \in L_\varepsilon(x_1abx_2)$ and M is locally conservative for \mathcal{C}_{rp} , we obtain $M[init, \sigma ab] = M[init, \sigma]$. For reasons of global consistency, $L_{M[init,\sigma_1]} = L_\varepsilon(ax_1)$. Now, since $M[init, \sigma ab] = M[init, \sigma]$, this yields $L_{M[init,\sigma_2]} = L_\varepsilon(ax_1)$. However, σ_2 can be extended to a text for $L_\varepsilon(ax_1ax_2)$, on which M will fail to learn locally conservatively, since $M[init, \sigma_2]$ overgeneralises the target. This contradicts the assumptions on M .

Case (b). $L_{M[init,\sigma]} = L_\varepsilon(ax_1ax_2)$.

Here a similar contradiction can be obtained for M processing $\sigma_1 = \sigma aa, ab$ and $\sigma_2 = \sigma ab$.

Both cases yield a contradiction and thus the theorem is verified. \square

However, as Theorems 3 and 4 show, each of our natural requirements, in its stronger formulation, can be achieved separately, if an appropriate indexing of the regular erasing pattern languages is used as a hypothesis space. We provide the proof only for the first result; a similar idea can be used also for Theorem 4.

Theorem 3. *There is an indexing $(L_j^*)_{j \in \mathbb{N}}$ of \mathcal{C}_{rp} and a learner M witnessing $\mathcal{C}_{rp} \in ItLConvTxt$ with respect to $(L_j^*)_{j \in \mathbb{N}}$.*

Proof. As in the proof of Theorem 1, let $(D_j)_{j \in \mathbb{N}}$ be the canonical enumeration of all finite subsets of \mathbb{N} and $(L_\varepsilon(\alpha_j))_{j \in \mathbb{N}}$ an effective, repetition-free indexing of \mathcal{C}_{rp} . Moreover let $L'_j = \bigcap_{z \in D_j} L_\varepsilon(\alpha_z)$ for all $j \in \mathbb{N}$. Hence $(L'_j)_{j \in \mathbb{N}}$ is an indexing comprising the class \mathcal{C}_{rp} . The proof is based on the following fact.

Fact 2. *There is an algorithm A' which, given any index j as input, outputs an index k with $L_\varepsilon(\alpha_k) = L'_j$, if such an index exists, and 'no', otherwise.*

(* Since every regular erasing pattern language is a regular language and both the inclusion problem as well as the equivalence problem for regular languages are decidable, such an algorithm A' exists. *)

The required iterative learner uses the algorithm A' and the iterative learner M from the demonstration of Theorem 1 as its subroutines. Let $(L_{\langle k,j \rangle}^*)_{k,j \in \mathbb{N}}$ be an indexing of \mathcal{C}_{rp} with $L_{\langle k,j \rangle}^* = L_\varepsilon(\alpha_k)$ for all $k, j \in \mathbb{N}$. We define an iterative learner M' for \mathcal{C}_{rp} that uses the hypothesis space $(L_{\langle k,j \rangle}^*)_{k,j \in \mathbb{N}}$.

Initially, if the first string w appears, M' determines the canonical index k of the regular erasing pattern language $L_\varepsilon(w)$ as well as $j = M(init, w)$, and outputs the hypothesis $\langle k, j \rangle$, i. e., $M'(init, w) = \langle k, j \rangle$. Next M' , when receiving a string v , refines its recent hypothesis, say $\langle k', j' \rangle$, as follows. First, if $v \in L_{\langle k', j' \rangle}^*$, M' repeats its recent hypothesis, i. e., $M'(\langle k', j' \rangle, v) = \langle k', j' \rangle$. (* Note that $j' = M(j', v)$, too. *) Second, if $v \notin L_{\langle k', j' \rangle}^*$, M' determines $j = M(j', v)$ and runs A' on input j . If A' returns some $k \in \mathbb{N}$, M' returns $\langle k, j \rangle$, i. e., $M'(\langle k', j' \rangle, v) = \langle k, j \rangle$. If A' returns 'no', M' determines the canonical index k of the regular erasing pattern language $L_\varepsilon(v)$ and returns $\langle k, j \rangle$, i. e., $M'(\langle k', j' \rangle, v) = \langle k, j \rangle$.

By definition, M' is an iterative and locally conservative learner. Let t be any text for any $L \in \mathcal{C}_{\text{rp}}$. Since M learns L , there is some n such that $M[\text{init}, t[n]] = j$ with $L'_j = L$. By definition, for $\langle k, j \rangle = M'[\text{init}, t[n]]$, we have $L_\varepsilon(\alpha_k) = L'_j$. Thus $L_{\langle k, j \rangle}^* = L_\varepsilon(\alpha_k)$. Since M' is a locally conservative learner, M' learns L , too. \square

Theorem 4. *There is an indexing $(L_j)_{j \in \mathbb{N}}$ of \mathcal{C}_{rp} and a learner M witnessing $\mathcal{C}_{\text{rp}} \in \text{ItGConsTxt}$ with respect to $(L_j)_{j \in \mathbb{N}}$.*

This case study shows that the necessity of auxiliary hypotheses representing languages outside the target class may depend on whether both global consistency and local conservativeness or only one of these properties is required. In what follows, we analyse the impact of consistency and conservativeness separately in a more general context, assuming that auxiliary hypotheses are allowed.

4 Incremental learning and consistency

This section is concerned with the impact of consistency demands in iterative learning. In the case of learning from text, the weaker consistency demand, namely local consistency, does not restrict the capabilities of iterative learners.

Theorem 5. *$\text{ItLConsTxt} = \text{ItTxt}$.*

Proof. By definition, $\text{ItLConsTxt} \subseteq \text{ItTxt}$. To prove $\text{ItTxt} \subseteq \text{ItLConsTxt}$, fix an indexable class $\mathcal{C} \in \text{ItTxt}$. Let $(L_j)_{j \in \mathbb{N}}$ be an indexing comprising \mathcal{C} and M an iterative learner for \mathcal{C} with respect to $(L_j)_{j \in \mathbb{N}}$.

The required learner M' uses the indexing $(L'_{\langle j, w \rangle})_{j \in \mathbb{N}, w \in \Sigma^*}$, where $L'_{\langle j, w \rangle} = L_j \cup \{w\}$ for all $j \in \mathbb{N}$, $w \in \Sigma^*$. Initially, $M'(\text{init}, w) = \langle j, w \rangle$ for $j = M(\text{init}, w)$. Next M' , upon a string v , refines its recent hypothesis, say $\langle j', w' \rangle$, as follows. First, M' determines $j = M(j', v)$. Second, if $v \in L_j$, M returns $\langle j, w' \rangle$; otherwise, it returns $\langle j, v \rangle$. Obviously, M' witnesses $\mathcal{C} \in \text{ItLConsTxt}$. \square

In contrast to that, requiring local consistency results in a loss of learning potential, as the following theorem shows.

Theorem 6. *$\text{ItGConsTxt} \subset \text{ItTxt}$.*

Proof. By definition, $\text{ItGConsTxt} \subseteq \text{ItTxt}$. It remains to provide a separating class \mathcal{C} that witnesses $\text{ItTxt} \setminus \text{ItGConsTxt} \neq \emptyset$.

Let $\Sigma = \{a, b\}$ and let $(A_j)_{j \in \mathbb{N}}$ be the canonical enumeration of all finite subsets of $\{a\}^+$. Now \mathcal{C} contains the language $L = \{a\}^+$ and, for all $j \in \mathbb{N}$, the finite language $L_j = A_j \cup \{b^z \mid z \leq j\}$.

Claim 1. $\mathcal{C} \in \text{ItTxt}$.

The required iterative learner M may work as follows. As long as exclusively strings from $\{a\}^+$ appear, M just guesses L . If a string of form b^j appears for the first time, M guesses L_j . Past that point, M , when receiving a string v , refines its recent guess, say L_k , as follows. If $v \in L$ or $v = b^z$ for some $z \leq k$, M repeats its guess L_k . If $v = b^z$ for some $z > k$, M guesses L_z .

It is not hard to verify that M is an iterative learner that learns \mathcal{C} as required.
Claim 2. $\mathcal{C} \notin ItGConsTxt$.

Suppose to the contrary that there is an indexing $(L'_j)_{j \in \mathbb{N}}$ comprising \mathcal{C} and a learner M witnessing $\mathcal{C} \in ItGConsTxt$ with respect to $(L'_j)_{j \in \mathbb{N}}$.

Consider M when processing the text $t = a^1, a^2, \dots$ for L . Since M is a learner for \mathcal{C} , there has to be some n such that $M[init, t[n]] = M[init, t[n+m]]$ for all $m \geq 1$. (* Note that $M[init, t[n]] = M[init, t[n]a^z]$ for all $z > n+1$.*)

Now let j be fixed such that $A_j = content(t[n]) = \{a^1, \dots, a^{n+1}\}$. Consider M when processing any text \hat{t} for L_j with $\hat{t}[n] = t[n]$. Since M is a learner for \mathcal{C} , there is some $n' > n$ such that $content(\hat{t}[n']) = L_j$ as well as $L'_k = L_j$ for $k = M[init, \hat{t}[n']]$. (* Note that there is some finite sequence σ with $\hat{t}[n'] = t[n]\sigma$.*)

Next let $j' > j$ be fixed such that $A_j \subset A_{j'}$. Moreover fix any string a^z in $A_{j'} \setminus A_j$. (* Note that $z > n+1$ and $a^z \notin L_j$.*) Consider M when processing any text \tilde{t} for the language $L_{j'}$ having the initial segment $\tilde{t}[n'+1] = t[n]a^z\sigma$. Since $M[init, t[n]] = M[init, t[n]a^z]$, one obtains $M[init, \tilde{t}[n+1]] = M[init, \hat{t}[n]]$. Finally since M is an iterative learner, $\hat{t}[n'] = \hat{t}[n]\sigma$, and $\tilde{t}[n'+1] = \tilde{t}[n+1]\sigma$, one may conclude that $M[init, \tilde{t}[n'+1]] = M[init, \hat{t}[n']] = k$. But $L'_k = L_j$, and therefore $a^z \notin L'_k$. The latter implies $content(\tilde{t}[n'+1]) \not\subseteq L'_k$, contradicting the assumption that M is an iterative and globally consistent learner for \mathcal{C} . \square

In the case of learning from informant, the results obtained are parallel to those in the text case. Theorem 7 can be verified similarly to Theorem 5.

Theorem 7. $ItLConsInf = ItConsInf$.

Considering the stronger consistency requirement, there are even classes learnable iteratively from text, but not globally consistently from informant.

Theorem 8. $ItTxt \setminus ItGConsInf \neq \emptyset$.

Proof. It suffices to provide a class $\mathcal{C} \in ItTxt \setminus ItGConsInf$.

Let $\Sigma = \{a, b\}$ and let $(A_j)_{j \in \mathbb{N}}$ be the canonical enumeration of all finite subsets of $\{a\}^+$. Now \mathcal{C} contains the language $L = \{a\}^+$ and, for all $j, k \in \mathbb{N}$, the finite language $L_{\langle j, k \rangle} = A_j \cup A_k \cup \{b^j, b^k\}$.

Claim 1. $\mathcal{C} \in ItTxt$.

The required iterative learner M may work as follows. As long as only strings from $\{a\}^+$ appear, M guesses L . If a string of form b^z appears for the first time, M guesses $L_{\langle z, z \rangle}$. Past that point, M refines its recent guess, say $L_{\langle j', k' \rangle}$, when receiving a string v as follows. If $j' = k'$ and $v = b^z$ with $z \neq j'$, M guesses $L_{\langle j', z \rangle}$. In all other cases, M repeats its guess $L_{\langle j', k' \rangle}$.

It is not hard to verify that M is an iterative learner that learns \mathcal{C} as required.

Claim 2. $\mathcal{C} \notin ItGConsInf$.

Suppose to the contrary that there is an indexing $(L'_j)_{j \in \mathbb{N}}$ comprising \mathcal{C} and a learner M witnessing $\mathcal{C} \in ItGConsInf$ with respect to $(L'_j)_{j \in \mathbb{N}}$.

Consider an informant $i = ((w_n, b_n)_{n \in \mathbb{N}})$ for L such that $|w_n| \leq n$ for all $n \in \mathbb{N}$. Since M is a learner for \mathcal{C} , there has to be some n such that $M[init, i[n]] = M[init, i[n+m]]$ for all $m \geq 1$. (* Note that $M[init, i[n]] = M[init, i[n](a^z, +)]$ for all $z > n+1$.*)

Let j be fixed such that $\text{content}^+(i[n]) \subseteq A_j$ and $b^j \notin \text{content}^-(i[n])$. Now consider M when processing an informant \hat{i} for $L_{\langle j, j \rangle}$ with $\hat{i}[n] = i[n]$. Since M is a learner for \mathcal{C} , there has to be some $n' > n$ such that $\text{content}(\hat{i}[n']) = L_{\langle j, j \rangle}$ and $L'_k = L_{\langle j, j \rangle}$ for $k = M[\text{init}, \hat{i}[n']]$. (* Note that there is some finite sequence σ such that $\hat{i}[n'] = i[n]\sigma$. *)

Now let $k' > j$ be fixed such that $A_j \subset A_{k'}$, $\text{content}^-(\hat{i}[n]) \cap A_{k'} = \emptyset$, and $b^{k'} \notin \text{content}^-(\hat{i}[n])$. Let a^z be any string in $A_{k'} \setminus A_j$. (* Note that $z > n + 1$ and $a^z \notin L_{\langle j, j \rangle}$. *) Consider M when processing any informant \tilde{i} for the language $L_{\langle j, k' \rangle}$ with $\tilde{i}[n' + 1] = i[n](a^z, +)\sigma$. Since $M[\text{init}, i[n]] = M[\text{init}, i[n](a^z, +)]$, one obtains $M[\text{init}, \tilde{i}[n + 1]] = M[\text{init}, \hat{i}[n]]$. Finally since M is an iterative learner, $\hat{i}[n'] = \hat{i}[n]\sigma$, and $\tilde{i}[n' + 1] = \tilde{i}[n + 1]\sigma$, one may conclude that $M[\text{init}, \tilde{i}[n' + 1]] = M[\text{init}, \hat{i}[n']] = k$. But $L'_k = L_{\langle j, j \rangle}$, and therefore $a^z \notin L'_k$. The latter implies $\text{content}^+(\tilde{i}[n' + 1]) \not\subseteq L'_k$, contradicting the assumption that M is an iterative and globally consistent learner for \mathcal{C} . \square

Obviously $\text{ItTtxt} \subseteq \text{ItInf}$, and thus we obtain the following corollary.

Corollary 1. $\text{ItGConsInf} \subset \text{ItInf}$.

5 Incremental learning and conservativeness

This section deals with conservativeness in the context of iterative learning. Here the results for learning from text differ from those for the informant case.

5.1 The case of learning from text

Let us first discuss the different conservativeness definitions in the context of learning from positive examples only. By definition, local conservativeness is a stronger demand, since the learner is required to maintain a hypothesis if it is consistent with the most recent piece of information, even if it contradicts some previously processed examples. However, it turns out that this demand does not have any negative effect on the capabilities of iterative learners. Intuitively, a globally conservative learner may change mind depending on inconsistency with only a limited set of examples, which can be coded within the hypothesis.

Theorem 9. $\text{ItGConvTtxt} = \text{ItLConvTtxt}$.

Proof. By definition, $\text{ItLConvTtxt} \subseteq \text{ItGConvTtxt}$. Fix an indexable class $\mathcal{C} \in \text{ItGConvTtxt}$; let $(L_j)_{j \in \mathbb{N}}$ be an indexing and M an iterative IIM identifying \mathcal{C} globally conservatively with respect to $(L_j)_{j \in \mathbb{N}}$. It remains to prove $\mathcal{C} \in \text{ItLConvTtxt}$. For that purpose, we need the following notion and technical claim.

Notion. For any text t and any $n \in \mathbb{N}$, let $mc(t[n], M)$ denote the set $\{t(0)\} \cup \{t(m) \mid 1 \leq m \leq n \text{ and } M[\text{init}, t[m-1]] \neq M[\text{init}, t[m]]\}$ of all strings in $\text{content}(t[n])$, which force M to change its mind when processing $t[n]$.

Technical claim. Let $L \in \mathcal{C}$, t a text for L , and $n \in \mathbb{N}$. Let $j = M[\text{init}, t[n]]$. If $t(n+1) \cup mc(t[n], M) \subseteq L_j$, then $M[\text{init}, t[n+1]] = M[\text{init}, t[n]]$.

Proof. Let $W = \text{content}(t[n+1]) \setminus L_j$. As $t(n+1) \cup mc(t[n], M) \subseteq L_j$, then $M[\text{init}, t[m+1]] = M[\text{init}, t[m]]$ for all $m < n$ with $t(m+1) \in W$. Now let τ be the subsequence of $t[n]$ obtained by deleting all $w \in W$ from $t[n]$. Obviously, $M[\text{init}, \tau] = M[\text{init}, t[n]]$ and $mc(t[n], M) \subseteq \text{content}(\tau) \subseteq L_j$. This implies

$$M[\text{init}, t[n+1]] = M[\text{init}, \tau t(n+1)] = M[\text{init}, \tau] = M[\text{init}, t[n]],$$

because M is globally conservative for L . (QED, technical claim).

Define an indexing $(L'_j)_{j \in \mathbb{N}}$ by $L'_{2\langle j, k \rangle} = L_j$ and $L'_{2\langle j, k \rangle + 1} = \emptyset$ for all $j, k \in \mathbb{N}$.

We now define an IIM M' (witnessing $\mathcal{C} \in \text{ItLConvTxt}$ using $(L'_j)_{j \in \mathbb{N}}$), such that, on any finite text segment σ for some $L \in \mathcal{C}$, the following invariant holds:

- $M'[\text{init}, \sigma] = 2\langle M[\text{init}, \sigma], k \rangle + y$ for some $k \in \mathbb{N}$, $y \in \{0, 1\}$, such that
- $D_k = mc(\sigma, M)$ (* and thus $D_k \subseteq \text{content}(\sigma)$ *).
- If $y = 0$, then $D_k \subseteq L_{M[\text{init}, \sigma]}$.

The reader may check that this invariant holds, if M' is defined as follows:

Definition of $M'(\text{init}, w)$, for $w \in \Sigma^$:* Let $j = M(\text{init}, w)$.

- If $w \in L_j$, let $M'(\text{init}, w) = 2\langle j, k \rangle$, where $D_k = \{w\}$.
- If $w \notin L_j$, let $M'(\text{init}, w) = 2\langle j, k \rangle + 1$, where $D_k = \{w\}$.

Definition of $M'(2\langle j, k \rangle + 1, w)$, for $w \in \Sigma^$, $j, k \in \mathbb{N}$:* Let $j' = M(j, w)$.

- If $j = j'$ and $D_k \subseteq L_j$, let $M'(2\langle j, k \rangle + 1, w) = 2\langle j, k \rangle$.
- If $j = j'$ and $D_k \not\subseteq L_j$, let $M'(2\langle j, k \rangle + 1, w) = 2\langle j, k \rangle + 1$.
- If $j \neq j'$, let $M'(2\langle j, k \rangle + 1, w) = 2\langle j', k' \rangle + 1$, where $D_{k'} = D_k \cup \{w\}$.

Definition of $M'(2\langle j, k \rangle, w)$, for $w \in \Sigma^$, $j, k \in \mathbb{N}$:* Let $j' = M(j, w)$.

- If $w \notin L_j$ and $j = j'$, let $M'(2\langle j, k \rangle, w) = 2\langle j, k \rangle + 1$.
- If $w \notin L_j$ and $j \neq j'$, let $M'(2\langle j, k \rangle, w) = 2\langle j', k' \rangle + 1$, where $D_{k'} = D_k \cup \{w\}$.
- If $w \in L_j$ (* by the invariant, there is some text segment σ with $M[\text{init}, \sigma] = j$ and $D_k = mc(\sigma, M) \subseteq L_j$; hence $D_k \cup \{w\} \subseteq L_j$ and $j = j'$ by the technical claim *), let $M'(2\langle j, k \rangle, w) = 2\langle j, k \rangle$.

By definition, M' is locally conservative with respect to $(L'_j)_{j \in \mathbb{N}}$. Since M is globally conservative for \mathcal{C} with respect to $(L_j)_{j \in \mathbb{N}}$ and because of the invariant, it is not hard to verify that M' learns \mathcal{C} iteratively. Thus $\mathcal{C} \in \text{ItLConvTxt}$. \square

So local and global conservativeness are equal constraints for iterative text learners. Whether they reduce the capabilities of iterative text learners in general, i. e., whether ItGConvTxt and ItTxt coincide, remains an open question.

5.2 The case of learning from informant

First, comparing the two versions of conservativeness, the informant case yields results different from those in the text case, namely that globally conservative iterative learners cannot be normalised to being locally conservative. In particular, the property that globally conservative learners can code all previously seen examples, for which their current hypothesis is inconsistent, no longer holds in the informant case.

Theorem 10. $ItLConvInf \subset ItGConvInf$.

Proof. By definition, $ItLConvInf \subseteq ItGConvInf$. Thus it remains to provide a separating class \mathcal{C} that witnesses $ItGConvInf \setminus ItLConvInf \neq \emptyset$.

Let $\Sigma = \{a\}$ and $(D_j)_{j \in \mathbb{N}}$ the canonical enumeration of all finite subsets of $\{a\}^+$. Assume $D_0 = \emptyset$. For all $j \in \mathbb{N}$, set $L_j = \{a^0\} \cup D_j$ and $L'_j = \{a\}^+ \setminus D_j$. Let \mathcal{C} be the collection of all finite languages L_j and all co-finite languages L'_j .

Claim 1. $\mathcal{C} \in ItGConvInf$.

For all $j, k, z \in \mathbb{N}$, let $H_{2\langle j, k, z \rangle} = \{a\}^+ \setminus \{a^z\}$ and $H_{2\langle j, k, z \rangle + 1} = \{a^z\}$. Now the required iterative learner M , processing an informant $i = ((w_n, b_n))_{n \in \mathbb{N}}$ for some $L \in \mathcal{C}$ may work as follows.

- (i) As long as neither $(a^0, +)$ nor $(a^0, -)$ appear, M guesses — depending on whether or not $(w_0, b_0) = (a^z, +)$ or $(w_0, b_0) = (a^z, -)$ — in the first case $H_{2\langle j, k, z \rangle}$, in the second case $H_{2\langle j, k, z \rangle + 1}$, where $D_j = \text{content}^+(i[n])$ and $D_k = \text{content}^-(i[n])$ (* The recent guess of M is inconsistent, so M can change its mind without violating the global conservativeness demand. *)
- (ii) If $(a^0, +)$ or $(a^0, -)$ appears for the first time, the following cases will be distinguished. If $w_0 = a^0$ and $b_0 = +$, M guesses L_0 . If $w_0 = a^0$ and $b_0 = -$, M guesses L'_0 . Otherwise, let $j' = 2\langle j, k, z \rangle + y$, $y \in \{0, 1\}$, denote the recent guess of M . If $(a^0, +)$ appears, M' guesses the finite language L_j . If $(a^0, -)$ appears, M' guesses the co-finite language L'_k .
- (iii) Then M refines its recent guess as follows. If a positive example $(a^z, +)$ appears, the recent guess of M is $L_{j'}$, and $a^z \notin L_{j'}$, M guesses $L_j = L_{j'} \cup \{a^z\}$. If a negative example $(a^z, -)$ appears, the recent guess of M is $L'_{k'}$, and $a^z \in L'_{k'}$, M guesses $L'_k = L'_{k'} \setminus \{a^z\}$. Else M repeats its recent guess.

It is not hard to verify that M is an iterative learner that learns \mathcal{C} as required.

Claim 2. $\mathcal{C} \notin ItLConvInf$.

Suppose to the contrary that there is an indexing $(L_j^*)_{j \in \mathbb{N}}$ comprising \mathcal{C} and a learner M which locally conservatively identifies \mathcal{C} with respect to $(L_j^*)_{j \in \mathbb{N}}$.

Let $j = M(\text{init}, (a, +))$. We distinguish the following cases:

Case 1. $L_j^* \cap \{a\}^+$ is infinite.

Choose $a^r \in L_j^*$ with $r > 1$ and $L = \{a^0, a^1, a^r\}$. Consider M on the informant $i = (a, +), (a^r, +), (a^0, +), (a^2, -), \dots, (a^{r-1}, -), (a^{r+1}, -), (a^{r+2}, -), \dots$ for L . As M learns \mathcal{C} , there is an $n \geq 2$ with $M[\text{init}, i[n]] = M[\text{init}, i[n + m]]$ for all $m \geq 1$. (* $M[\text{init}, i[n](a^s, -)] = M[\text{init}, i[n]]$ for all a^s with $a^s \notin (\text{content}^+(i[n]) \cup \text{content}^-(i[n]))$.) * Let a^s be any string in L_j^* with $s > r + 1$, $a^s \notin (\text{content}^+(i[n]) \cup \text{content}^-(i[n]))$. As $L_j \cap \{a\}^+$ is infinite, such a^s exists. (* There is some σ with $i = (a, +), (a^r, +)\sigma(a^{s-1}, -), (a^s, -), (a^{s+1}, -), \dots$ *)

Next let $\hat{i} = (a^1, +), (a^r, +), (a^s, +)\sigma(a^{s-1}, -), (a^{s+1}, -), (a^{s+2}, -), \dots$. Consider M when processing the informant \hat{i} for $L' = \{a^0, a^1, a^r, a^s\}$. Since M is locally conservative and $a^s \in L_j^*$, $M[\text{init}, \hat{i}[2]] = M[\text{init}, \hat{i}[1]]$. As M is an iterative learner, $M[\text{init}, \hat{i}[n + 1]] = M[\text{init}, \hat{i}[n]]$. Past step $n + 1$, M receives only negative examples $(a^z, -)$ with $a^z \notin (\text{content}^+(i[n]) \cup \text{content}^-(i[n]))$. Hence M converges on \hat{i} to the same hypothesis j as on i , namely to $j = M[\text{init}, i[n]]$. Finally because $L \neq L'$, M cannot learn both finite languages L and L' .

Case 2. $L_j^* \cap \{a\}^+$ is finite.

An argumentation similar to that used in Case 1 shows that M must fail to learn some co-finite language in \mathcal{C} . We omit the relevant details. \square

The observed difference in the above theorem can now even be extended to a proper hierarchy of iterative learning from informant; globally conservative learners in general outperform locally conservative ones, but are not capable of solving all the learning tasks a general iterative learner can cope with. So there are classes in $ItInf$ which cannot be learned by any iterative, globally conservative learner.

Theorem 11. $ItGConvInf \subset ItInf$.

Proof. By definition, $ItGConvInf \subseteq ItInf$. Thus it remains to provide a separating class \mathcal{C} that witnesses $ItInf \setminus ItGConvInf \neq \emptyset$.

Let $(D_j)_{j \in \mathbb{N}}$ be the canonical enumeration of all finite subsets of \mathbb{N} .

Let $\mathcal{C} = \bigcup_{k \in \mathbb{N}} \mathcal{C}_k$, where \mathcal{C}_k are defined below based on following cases.

Case (a). If $\varphi_k(k) \uparrow$, then \mathcal{C}_k contains just one language, namely $L_k = \{a^k\}$.

Case (b). If $\varphi_k(k) \downarrow$, then \mathcal{C}_k contains infinitely many languages. Let $s = \Phi_k(k)$. For all $j \in \mathbb{N}$, \mathcal{C}_k contains the language $L_{\langle k, j \rangle} = L_k \cup \{b^s\} \cup \{c^{s+z} \mid z \in D_j\}$ as well as the language $L'_{\langle k, j \rangle} = L_k \cup \{c^{s+z} \mid z \notin D_j\}$. (* Note that $L_{\langle k, j \rangle}$ contains a finite subset of $\{c\}^*$, whereas $L'_{\langle k, j \rangle}$ contains a co-finite subset of $\{c\}^*$.)

It is not hard to verify that \mathcal{C} constitutes an indexable class.

Claim 1. $\mathcal{C} \in ItInf$.

Let $i = ((w_n, b_n))_{n \in \mathbb{N}}$ be an informant for some $L \in \mathcal{C}$. A corresponding iterative learner M may be informally defined as follows:

- (i) As long as no positive example $(a^k, +)$ appears, M' encodes in its guess all examples seen so far.
- (ii) If some positive example $(a^k, +)$ appears, M' tests whether or not $\Phi_k(k) \leq |w|$, where w is the longest string seen so far. In case that $\varphi_k(k) \downarrow$ has been verified, M' guesses L_k , where in its hypothesis all examples seen so far are encoded. Subsequently, M' behaves according to (iv). In case that $\Phi_k(k) > |w|$, M' guesses L_k , where the encoded examples can be simply ignored. Afterwards, M' behaves according to (iii).
- (iii) As long as M' guesses L_k , M' uses the recent example (w_n, b_n) to check whether or not $\Phi_k(k) \leq |w_n|$. In the positive case, M' behaves as in (iv). Else M' repeats its recent guess, without encoding any further example.
- (iv) Let $s = \Phi_k(k)$. As long as $(b^s, +)$ and $(b^s, -)$ neither appear nor belong to the examples encoded in the recent guess, M' adds the new example into the encoding of examples in the recent guess. If $(b^s, +)$ (or $(b^s, -)$) appears or is encoded, M' guesses a language $L_{\langle k, j \rangle}$ (or $L'_{\langle k, j \rangle}$, respectively) that is consistent with all examples encoded. Past that point, M' works like the iterative learner M used in the proof of Theorem 10, Claim 1.

It is not hard to see that M' is an iterative learner for \mathcal{C} .

Claim 2. $\mathcal{C} \notin ItGConvInf$.

Suppose the converse. That is, there is an indexing $(L_j^*)_{j \in \mathbb{N}}$ comprising \mathcal{C} and an iterative learner M which globally conservatively identifies \mathcal{C} with respect to $(L_j^*)_{j \in \mathbb{N}}$. We shall show that M can be utilised to solve the halting problem.

Algorithm \mathcal{A} : Let k be given. Let $i = (w_n, b_n)_{n \in \mathbb{N}}$ be a repetition-free informant for L_k with $w_0 = a^k$ and $b_0 = +$ such that, for all $n \in \mathbb{N}$, $w_m = b^n$ implies $m < n$. For $m = 0, 1, 2, \dots$ test in parallel whether $(\alpha 1)$ or $(\alpha 2)$ happens.

$(\alpha 1)$ $\Phi_k(k) \leq m$.

$(\alpha 2)$ An index $j_m = M(\text{init}, i[m])$ is output such that $\text{content}^+(i[m]) \subseteq L_{j_m}^*$ and $\text{content}^-(i[m]) \cap L_{j_m}^* = \emptyset$.

If $(\alpha 1)$ happens first, output “ $\varphi_k(k) \downarrow$.” Otherwise, i.e., $(\alpha 2)$ happens first, output “ $\varphi_k(k) \uparrow$.”

Fact 1. On every input k , algorithm \mathcal{A} terminates.

It suffices to show that either $(\alpha 1)$ or $(\alpha 2)$ happens. Suppose, $(\alpha 1)$ does not happen, and thus $\varphi_k(k) \uparrow$. Hence, $L_k \in \mathcal{C}_k \subseteq \mathcal{C}$. Consequently, M , when processing the informant i for L_k , eventually returns a hypothesis $j_m = M(\text{init}, i[m])$ such that $L_{j_m}^* = L_k$. Thus, $(\alpha 2)$ must happen.

Fact 2. Algorithm \mathcal{A} decides the halting problem.

Obviously, if $(\alpha 1)$ happens then $\varphi_k(k)$ is indeed defined. Suppose $(\alpha 2)$ happens. We have to show that $\varphi_k(k) \uparrow$. Assume $\varphi_k(k) \downarrow$. Then, $\Phi_k(k) = s$ for some $s \in \mathbb{N}$. Since $(\alpha 2)$ happens, there is an $m < s$ such that $j_m = M(\text{init}, i[m])$ as well as $\text{content}^+(i[m]) \subseteq L_{j_m}^*$ and $\text{content}^-(i[m]) \cap L_{j_m}^* = \emptyset$. (* Note that neither $(b^s, +)$ nor $(b^s, -)$ appears in the initial segment $i[m]$.*)

Now, similarly to the proof of Theorem 10, Claim 2 one has to distinguish two cases: (i) $L_{j_m}^*$ contains infinitely many strings from $\{c\}^*$ and (ii) $L_{j_m}^*$ contains only finitely many strings of from $\{c\}^*$. In both cases, an argumentation similar to that used in the proof of Theorem 10, Claim 2 can be utilised to show that M fails to learn at least one language in \mathcal{C}_k which contain a finite (co-finite) subset of $\{c\}^*$. We omit the relevant details. Since M is supposed to learn \mathcal{C} , the latter contradicts our assumption that $\varphi_k(k) \downarrow$, and thus Fact 2 follows.

Since the halting problem is undecidable, $\mathcal{C} \notin \text{ItGConvInf}$. □

6 Some concluding remarks

We have studied iterative learning with two versions of consistency and conservativeness. In fact, a third version is conceivable. Note that an iterative learner M may use a redundant hypothesis space for coding in its current hypothesis all examples, upon which M has previously changed its guess. So one may think of mind changes as ‘memorising examples’ and repeating hypotheses as ‘forgetting examples’. One might call a hypothesis consistent with the examples seen, if it does not contradict the ‘memorised’ examples, i. e., those upon which M has changed its hypothesis. Similarly, M may be considered conservative, if M sticks to its recent hypothesis, as long as it agrees with the ‘memorised’ examples.

Obviously, this version of consistency is equivalent to local consistency – the proof is essentially the same as for Theorem 5 and the fact is not surprising.

However, the third version of conservativeness is worth considering a little closer. For iterative learning from text Theorem 9 immediately implies that this notion is equivalent to both global and local conservativeness. The idea is quite simple: a conservative learner really has to ‘know’ that it is allowed to change its hypothesis! Thus being inconsistent with forgotten positive examples doesn’t help at all, because the learner cannot memorise the forgotten examples and thus not justify its mind change. In this sense, ‘forgotten’ examples are really examples without any relevance for the learner on the given text. This intuition is already reflected in the technical claim used in the proof of Theorem 9.

Many similar insights may be taken from the proofs above to obtain further results. For instance, the separating classes provided in the proofs of Theorems 6 and 8, additionally lift our results to a more general case of incremental learning, where the learner has a k -bounded memory, i. e., the capacity for memorising up to k examples during the learning process, cf. [9]. Note that among our results we did not have a characterisation of the structure of classes learnable iteratively, however, our analysis will hopefully serve as a first step into this direction.

References

1. Angluin, D., Inductive inference of formal languages from positive data, *Information and Control* **45**, 117–135, 1980.
2. Angluin, D., Queries and concept learning, *Machine Learning* **2**, 319–342, 1988.
3. Blum, M., A machine independent theory of the complexity of recursive functions, *Journal of the ACM* **14**, 322–336, 1967.
4. Case, J., Jain, S., Lange, S., and Zeugmann, T., Incremental concept learning for bounded data mining, *Information and Computation* **152**, 74–110, 1999.
5. Gennari, J.H., Langle, P., and Fisher, D., Models of incremental concept formation, *Artificial Intelligence* **40**, 11–61, 1989.
6. Gold, E.M., Language identification in the limit, *Information and Control* **10**, 447–474, 1967.
7. Kinber, E. and Stephan, F., Language learning from texts: Mind changes, limited memory and monotonicity, *Information and Computation* **123**, 224–241, 1995.
8. Lange, S. and Grieser, G., On the power of incremental learning, *Theoretical Computer Science* **288**, 277–307, 2002.
9. Lange, S. and Zeugmann, T., Incremental learning from positive data, *Journal of Computer and System Sciences* **53**, 88–103, 1996.
10. Shinohara, T., Polynomial time inference of extended regular pattern languages, in: *Proc. RIMS Symposium on Software Science and Engineering*, LNCS, Vol. 147, pp. 115–127, Springer-Verlag, 1983.
11. Valiant, L.G., A theory of the learnable, *Communications of the ACM* **27**, 1134–1142, 1984.
12. Wiehagen, R., Limes-Erkennung rekursiver Funktionen durch spezielle Strategien, *Journal of Information Processing and Cybernetics (EIK)* **12**, 93–99, 1976.
13. Zeugmann, T. and Lange, S., A guided tour across the boundaries of learning recursive languages, in: *Algorithmic Learning for Knowledge-Based Systems*, LNAI, Vol. 961, pp. 190–258, Springer-Verlag, 1995.