A General Theory of Deduction, Induction, and Learning *

Eric Martin[1], Arun Sharma[1], and Frank Stephan[2]

[1] School of Computer Science and Engineering, The University of New South Wales, Sydney, NSW 2052, Australia,
{emartin, arun}@cse.unsw.edu.au
[2] Universität Heidelberg, 69121 Heidelberg, Germany,
fstephan@math.uni-heidelberg.de

Abstract. Deduction, induction, learning, are various aspects of a more general scientific activity: the discovery of truth. We propose to embed them in a common, logical framework. First, we define a generalized notion of "logical consequence." Alternating compact and "weakly compact" consequences, we stratify the set of generalized logical consequences of a given theory in a hierarchy. Classical first-order logic is a particular case of this framework; the fact that it is all about deduction is due to the compactness theorem, and this is reflected by the collapsing of the corresponding hierarchy to the first level. Classical learning paradigms in the inductive inference literature provide other particular cases. Finite learning corresponds exactly to the first level (or level Σ_1) of the hierarchy, whereas learning in the limit corresponds to another level (namely Σ_2). More generally, strong and natural connections exist between our hierarchy of generalized logical consequences, the Borel hierarchy, and the hierarchy which measures the complexity of a formula in terms of alternations of quantifiers. It is hoped that this framework provides the foundation of a unified logic of deduction and induction, and highlights the inductive nature of learning. An essential motivation for our work is to apply the theory presented here to the design of "Inductive Prolog", a system with both deductive and inductive capabilities, based on a natural extension of the resolution principle.

1 Introduction

Let us first make a few remarks about the nature of deduction and the nature of induction, before we turn to the nature of learning. If a formula φ is a deductive consequence of a set of formulas T, it is clear to anyone that φ is a logical consequence of T, in the sense that φ is true in every model of T. Many would also agree that we can substitute "inductive" for "deductive" in the previous sentence. What is then the difference between φ being a deductive and φ being

* Eric Martin is supported by the Australian Research Council Grant A49803051. Frank Stephan is supported by the Deutsche Forschungsgemeinschaft (DFG) Heisenberg Grant Ste 967/1-1.

an inductive consequence of T? Well, it should be possible to discover with certainty that φ is a deductive consequence of T, if this is indeed the case. Whereas it should not be possible to discover with certainty that φ is an inductive consequence of T, if φ is not in fact a deductive consequence of T. How can we discover with certainty that φ is true on the basis of T? A natural answer is: if and only if φ is actually a logical consequence of a finite subset of T. In other words, if and only if φ is a compact logical consequence of T. On the other hand, if φ is an inductive, but not a deductive, consequence of T, then we need an infinite part of T, if not the whole of T, in order to be able to establish this fact. At this point, two questions emerge:

1. What should count as a model of T? If it is any structure (as defined in classical logic) in which every member of T is true, and if T consists of first-order formulas, then the compactness theorem shows that every logical consequence of T is actually a deductive consequence of T, and there is no scope for a proper notion of induction. Hence, we should be able to consider not all structures, but some of them. This would result in a generalized notion of logical consequence that might not be compact. Are there natural candidates for such sets of structures?
2. Suppose that the class of models of T that have been retained is such that some generalized logical consequence φ of T is not a deductive (compact) consequence of T. Is then φ automatically "promoted" to the status of inductive consequence of T? The fact that every model of T is a model of φ involves infinitely many members of T. But how difficult is it to conclude that φ is true on the basis of T? If we can define difficulty levels, should one of them be considered as "the inductive level?"

Let us now consider learning (for more details on the notions mentioned below, see [7]). We claim that the classical paradigms in the inductive inference literature are also about discovering the truth. Suppose that the underlying logical vocabulary consists of a unary predicate symbol P, together with a constant $\overline{n}$ for each natural number n. A language L can be identified with $T = \{P\overline{n} \,|\, n \in L\}$, and its complement with $\overline{T} = \{\neg P\overline{n} \,|\, n \in \mathbb{N} \setminus L\}$. A text (respect. informant) for L can then be identified with an enumeration of T (respect. $T \cup \overline{T}$). The task of discovering an r.e. index for (respect. the characteristic function of) L can be identified with the task of discovering the infinitary formula $\bigwedge T$ (respect. $\bigwedge T \wedge \bigwedge \overline{T}$). Clearly $\bigwedge T$ is a logical consequence of T—in the classical sense, hence also in any more general sense. Retaining only the structures that correspond to languages *i.e.*, the intended possible realities, will make $\bigwedge T \wedge \bigwedge \overline{T}$ a generalized consequence of T. So in both cases, identification in the limit is about discovering a particular generalized logical consequence of T, namely a formula that can be viewed upon as a description of the language to be learned. On the other hand, the task of discovering the truth of an arbitrary formula φ from a background theory is equivalent to partial classification (see [4]): a formula φ represents the class $\mathcal{C}$ of all theories T which logically imply φ in a general sense, and the partial classifier has to find out, on the basis of data from background theory T, that φ is a generalized logical consequence of T, whenever this is true. Note the following:

1. Considering the infinite formula $P\overline{n_0} \wedge P\overline{n_1} \wedge P\overline{n_2} \wedge \ldots$ rather than the index of a Turing machine which generates $n_0, n_1, n_2 \ldots$ provides a logical representation equivalent to a representation in terms r.e. indexes. But infinite formulas are not only a technical way of embedding learning paradigms into a logical framework. It turns out that the extension of first-order languages to countable fragments of $\mathcal{L}_{\omega_1\omega}$ (see below) are the natural logical languages of our framework.
2. When learning from positive data only, there is an implicit assumption that *all* positive data are enumerated in a text for a language L. This means that the models of $\{P\overline{n} \mid n \in L\}$ to be considered should not be consistent with $P\overline{n}$ for any $n \in \mathbb{N} \setminus L$. The notion of generalized logical consequence, together with the right set of structures, should be able to accommodate this kind of property.

The previous considerations go beyond epistemological concerns about the nature of induction or learning. Indeed, the aim of this work is also to investigate the foundations of induction in AI. Current work in Inductive Logic Programming (see [14]) focuses on a very specific inductive task: discovering the minimal model of a potentially infinite set of data. We take a more general view and investigate general inductive abilities. Considering both deduction and induction as particular expressions of the art of discovering the truth opens the door to a unified framework which can provide the basis of an "Inductive Prolog". If Prolog is the deductive engine of AI, giving an agent the ability to compute solutions to existential queries, Inductive Prolog should be the deductivo-inductive engine of AI, giving an agent the ability to compute solutions to existential or Σ_2 queries, such as: does there exist a chemical compound which has this effect on all molecules having this and that property?

We proceed as follows. In Section 2, we introduce the necessary notation and in Section 3, we describe the components of our framework. In Section 4, we define hierarchies of generalized logical consequences by the alternating the use of compactness and weak compactness, and show some of their properties relevant to learning paradigms. In Section 5, we investigate the relationship between the hierarchies of generalized logical consequences and formula complexity. As additional evidence of their naturalness we also demonstrate links with the Borel hierarchy. Finally, in Section 6, we show how a number of classical learning paradigms can be cast into our framework.

2 Notation

A *vocabulary* is a countable set of function symbols (possibly including constants) and predicate symbols. A vocabulary can, but does not have to, contain equality. If it does not, it is said to be *equality free*. From now on, S denotes an arbitrary countable vocabulary. For some results, assumptions on S will be made. We denote by $\mathfrak{L}^S_{\omega\omega}$ the set of all first-order S-formulas, and by $\mathfrak{L}^S_{\omega_1\omega}$ the extension of $\mathfrak{L}^S_{\omega\omega}$ that accepts countable nonempty conjunctions and disjunctions.[1] So for all countable nonempty $T \subseteq \mathfrak{L}^S_{\omega_1\omega}$, the disjunction of all members of T, written $\bigvee T$, and the conjunction of all members of T, written $\bigwedge T$, both belong to $\mathfrak{L}^S_{\omega_1\omega}$. Note that the occurrence or nonoccurrence of $=$ in S determines whether $\mathfrak{L}^S_{\omega\omega}$ and $\mathfrak{L}^S_{\omega_1\omega}$ are languages with or without equality. A *countable fragment of* $\mathfrak{L}^S_{\omega_1\omega}$ is a countable subset $\mathfrak{L}$ of $\mathfrak{L}^S_{\omega_1\omega}$ which contains $\mathfrak{L}^S_{\omega\omega}$, is closed under subformulas, boolean operators, and quantification.[2] From now on, $\mathfrak{L}$ denotes a countable fragment of $\mathfrak{L}^S_{\omega_1\omega}$. It represents the language on the basis of which the core of the theory is developed. Clearly, $\mathfrak{L}^S_{\omega\omega}$ is the smallest countable fragment of $\mathfrak{L}^S_{\omega_1\omega}$. The members of $\mathfrak{L}^S_{\omega_1\omega}$ which are in Σ_0 or Π_0 *prenex form* are the quantifier free members of $\mathfrak{L}^S_{\omega\omega}$. Let nonnull ordinal α and $\varphi \in \mathfrak{L}^S_{\omega_1\omega}$ be given. We say that φ is in Σ_α (respect. Π_α) *prenex form* just in case one of the following holds:

1. φ is in Σ_β or Π_β prenex form for some $\beta < \alpha$, or
2. φ is of the form $\exists x\psi$ (respect. $\forall x\psi$) for some $\psi \in \mathfrak{L}^S_{\omega_1\omega}$ which is in Σ_α (respect. Π_α) prenex form, or
3. φ is of the form $\bigvee X$ (respect. $\bigwedge X$) for some (countable) $X \subseteq \mathfrak{L}^S_{\omega_1\omega}$ all of whose members are in Σ_α (respect. Π_α) prenex form.

It is easy to verify that every member of $\mathfrak{L}^S_{\omega_1\omega}$ is logically equivalent to a member of $\mathfrak{L}^S_{\omega_1\omega}$ which is in Σ_α prenex form for some α. If $\varphi \in \mathfrak{L}^S_{\omega_1\omega}$ is logically equivalent to a closed member of $\mathfrak{L}^S_{\omega_1\omega}$ which is in Σ_α (respect. Π_α) prenex form, then we say that φ is Σ_α (respect. Π_α). Note that the classical definition of a member of $\mathfrak{L}^S_{\omega\omega}$ being Σ_n (respect. Π_n) for some $n \in \mathbb{N}$ is a particular case of the former.

The $\sim$ *operator* is the function $\sim\colon \mathfrak{L}^S_{\omega_1\omega} \to \mathfrak{L}^S_{\omega_1\omega}$ which is defined as follows. If $\varphi \in \mathfrak{L}^S_{\omega_1\omega}$ is of form $\neg\psi$ for some $\psi \in \mathfrak{L}^S_{\omega_1\omega}$, then $\sim(\varphi) = \psi$; otherwise $\sim(\varphi) = \neg\varphi$. Given $\Gamma \subseteq \mathfrak{L}^S_{\omega_1\omega}$ and S-structure $\mathfrak{M}$, the Γ*-diagram of* $\mathfrak{M}$, denoted $D_\Gamma(\mathfrak{M})$, is the set of all members of Γ that are true in $\mathfrak{M}$. *Terms* will refer to S-terms, *formulas* to members of $\mathfrak{L}$ (not $\mathfrak{L}^S_{\omega_1\omega}$), *sentences* to closed formulas, and *structures* to S-structures. A *Henkin structure* is a structure all of whose individuals interpret closed terms.[3] A *Herbrand structure* is a structure each of whose individuals interprets a unique closed term.[4] Hence Herbrand structures are Henkin. When we consider a Henkin or a Herbrand structure, or a nonempty class of Henkin or Herbrand structures, we tacitly assume that S contains at least one constant.

[1] Given regular cardinal κ, $\mathfrak{L}^S_{\kappa\omega}$ denotes the set of all S-formulas built from atomic S-formulas using boolean operators, quantifiers, and disjunctions or conjunctions over nonempty sets of cardinality smaller than κ. See [10].

[2] For more details on this definition, see [2] or [12].

[3] Henkin structures should not be confused with *Henkin models*, which is the name often given to the *general models* defined in [6]. Our notion of Henkin structure is closer to the *canonical structures* defined in [17] for Henkin's proof of the completeness of first-order logic.

[4] Herbrand structures are close to the *Herbrand models* considered in Logic Programming. See [3] or [11].

3 Components of the theory

We denote by $\mathcal{W}$ a class of structures, the class of *possible worlds*. Classical first-order logic would take for $\mathcal{W}$ the class of all structures. We have explained that in order to address questions such as deduction versus induction, we need to be free to choose a more restrictive class of possible worlds. The discussion about learning suggests the consideration of $\mathcal{W}$ to be the class of all Henkin structures, or the class of all Herbrand structures. Henkin and Herbrand structures are interesting in many respects, and play a prominent role in Logic Programming ([11]). Given $T \subseteq \mathcal{L}$, we denote by $\mathrm{Mod}_{\mathcal{W}}(T)$ the class of all members of $\mathcal{W}$ that are models of T.

We denote by $\mathcal{O}$ a set of sentences, that we call the class of *possible observations*. For classical first-order logic, the choice of $\mathcal{O}$ would be irrelevant. Suppose we want to cast learning paradigms into this framework. For learning from positive data only, $\mathcal{O}$ will be equal to the set of all atomic sentences; for learning from both positive and negative data, $\mathcal{O}$ will be equal to the set of all basic sentences. Other examples can also be found in the literature (see for example [9]).

We denote by $\mathcal{T}$ a set of sets of sentences, that we call class of *possible theories*. This corresponds roughly to the class of possible texts in the inductive inference literature. Classical first-order logic would take for $\mathcal{T}$ the set of all sets of closed members of $\mathcal{L}^S_{\omega\omega}$.

The quintuple $(S, \mathcal{L}, \mathcal{W}, \mathcal{O}, \mathcal{T})$ contains all we need to define the fundamental concepts of this framework. We call this quintuple the *paradigm* under investigation, and we denote it by $\mathcal{P}$.

Definition 1. *Let $T \subseteq \mathcal{L}$ and $\mathfrak{M} \in \mathcal{W}$ be given. We say that $\mathfrak{M}$ is an* $\mathcal{O}$-minimal model of T in $\mathcal{W}$ *iff $\mathfrak{M} \in \mathrm{Mod}_{\mathcal{W}}(T)$ and for all $\mathfrak{N} \in \mathrm{Mod}_{\mathcal{W}}(T)$,*

$$\{\varphi \in \mathcal{O} \mid \mathfrak{N} \models \varphi\} \not\subset \{\varphi \in \mathcal{O} \mid \mathfrak{M} \models \varphi\}.$$

The discussion above about learning should justify the previous definition. A similar notion is also encountered in AI in the form of the *closed-world assumption* defined in [16] (for an overview see [5]), and of course in Logic Programming with the *least Herbrand models* (see [3, 11]). Let $T \subseteq \mathcal{L}$ be given. Then T can have exactly one $\mathcal{O}$-minimal model in $\mathcal{W}$, or none, or many. We denote by $\mathrm{Mod}^{\mathcal{O}}_{\mathcal{W}}(T)$ the class of all $\mathcal{O}$-minimal models of T in $\mathcal{W}$. Note the following:

Lemma 1. *If $\mathcal{O}$ is closed under $\sim$ then for all $T \subseteq \mathcal{L}$, $\mathrm{Mod}^{\mathcal{O}}_{\mathcal{W}}(T) = \mathrm{Mod}_{\mathcal{W}}(T)$.*

We can now generalize the notion of logical consequence:

Definition 2. *Let $T \subseteq \mathcal{L}$ and $\varphi \in \mathcal{L}$ be given. We say that φ is a* logical consequence of T in $\mathcal{W}$, *and we write $T \models_{\mathcal{W}} \varphi$, iff every member of $\mathrm{Mod}_{\mathcal{W}}(T)$ is a model of φ. We say that φ is an* $\mathcal{O}$-minimal logical consequence of T in $\mathcal{W}$, *and we write $T \models^{\mathcal{O}}_{\mathcal{W}} \varphi$, iff every member of $\mathrm{Mod}^{\mathcal{O}}_{\mathcal{W}}(T)$ is a model of φ.*

The notion of $\mathcal{O}$-minimal logical consequence in $\mathcal{W}$ is the notion of generalized logical consequence we investigate; the other just proves useful. Although we develop the theory on a very broad basis, here we consider almost exclusively two cases of paradigms, that we now define.

Definition 3. *We say that $\mathcal{P}$ is* standard *iff $\mathcal{T} = \{D_{\mathcal{O}}(\mathfrak{M}) \mid \mathfrak{M} \in \mathcal{W}\}$. If $\mathcal{P}$ is standard and for all $T \in \mathcal{T}$ and sentences φ, either $T \models^{\mathcal{O}}_{\mathcal{W}} \varphi$ or $T \models^{\mathcal{O}}_{\mathcal{W}} \neg\varphi$, then we say that $\mathcal{P}$ is* ideal.

Standard paradigms are the analogues of the classical paradigms in the inductive inference literature. When no data are missing, the latter even correspond to ideal paradigms.

4 The hierarchies of generalized logical consequences

We now define the hierarchies of generalized logical consequences that are basically the fundamental object of study of this framework.[5] First we set, for all $T \in \mathcal{T}$, $\Sigma^{\mathcal{P}}_0(T) = \Pi^{\mathcal{P}}_0(T) = T$.

Definition 4. *Let nonnull ordinal α and $T \in \mathcal{T}$ be given. Suppose that $\Pi^{\mathcal{P}}_\beta(T)$ has been defined for all $\beta < \alpha$. A sentence φ belongs to $\Sigma^{\mathcal{P}}_\alpha(T)$ iff there exists finite $E \subseteq T$ and finite $H \subseteq \bigcup_{\beta<\alpha} \Pi^{\mathcal{P}}_\beta(T)$ such that for all $T' \in \mathcal{T}$:*

$$\text{if } E \subseteq T' \text{ and } T' \models^{\mathcal{O}}_{\mathcal{W}} H \text{ then } T' \models^{\mathcal{O}}_{\mathcal{W}} \varphi.$$

Definition 5. *Let nonnull ordinal α and $T \in \mathcal{T}$ be given. Suppose that $\Sigma^{\mathcal{P}}_\alpha(T')$ has been defined for all $T' \in \mathcal{T}$. A sentence φ belongs to $\Pi^{\mathcal{P}}_\alpha(T)$ iff $T \models^{\mathcal{O}}_{\mathcal{W}} \varphi$ and there exists finite $E \subseteq T$ and finite $H \subseteq \Sigma^{\mathcal{P}}_\alpha(T)$ such that for all $T' \in \mathcal{T}$:*

$$\text{if } E \subseteq T',\ T' \models^{\mathcal{O}}_{\mathcal{W}} H, \text{ and } T' \not\models^{\mathcal{O}}_{\mathcal{W}} \varphi \text{ then } \neg\varphi \in \Sigma^{\mathcal{P}}_\alpha(T').$$

The description of the Σ_α level is based on a compactness property: a finite information—evidence (subset of the theory) and hypotheses (sentences that have been "discovered" before)—enables to conclude that φ is an $\mathcal{O}$-minimal logical consequence of T in $\mathcal{W}$. The description of the Π_α level is based on a property of weak compactness: a finite information—of the same kind as before—enables to conclude that φ could not belong to the set of sentences that are not $\mathcal{O}$-minimal logical consequences of T in $\mathcal{W}$, without this fact being already discovered. The compactness property enables to conclude with certainty that φ is true. The property of weak compactness enables to believe confidently in φ. Of course, the certain conclusion and the confident belief in question are not absolute, they are relative to that part of the hierarchy below the level which is currently built. Only the sentences in $\Sigma^{\mathcal{P}}_1(T)$ can be discovered to be true with absolute certainty, and only the sentences in $\Pi^{\mathcal{P}}_1(T)$ can be believed to be true with absolute confidence, as far as T is reliable, for instance, in the case of standard paradigms, as far as T really contains all possible observations true in the underlying world. Note the relationship between weak compactness and the notion of refutability defined by Popper ([15]).

There are many equivalents to Definitions 4 and 5, some of them simpler, particularly when $\mathcal{P}$ is standard or ideal. For instance, it is easy to verify the following, to be used in the sequel.

[5] Since member of $\mathcal{L}^S_{\omega_1\omega}$ can contain infinitely many free variables, and due to the use of negation, it is simpler to accept only sentences in the hierarchies.

Lemma 2. *Suppose that $\mathcal{P}$ is standard. Let $T \in \mathcal{T}$ be given. A sentence φ belongs to $\Sigma_1^{\mathcal{P}}(T)$ iff there exists finite $E \subseteq T$ such that $E \models_{\mathcal{W}} \varphi$.*

Lemma 3. *Suppose that $\mathcal{P}$ is standard. Let $\alpha \neq 0$ and T in $\mathcal{T}$ be given. A sentence φ belongs to $\Pi_\alpha^{\mathcal{P}}(T)$ iff $T \models_{\mathcal{W}}^{\mathcal{O}} \varphi$ and there is $\psi \in \Sigma_\alpha^{\mathcal{P}}(T)$ such that for all $T' \in \mathcal{T}$, if $T' \models_{\mathcal{W}}^{\mathcal{O}} \psi$ and $T' \not\models_{\mathcal{W}}^{\mathcal{O}} \varphi$, then $\neg\varphi \in \Sigma_\alpha^{\mathcal{P}}(T')$.*

Lemma 4. *Suppose that $\mathcal{P}$ is ideal. Let $\alpha > 1$ and $T \in \mathcal{T}$ be given. A sentence φ belongs to $\Sigma_\alpha^{\mathcal{P}}(T)$ iff there is $\psi \in \bigcup_{\beta<\alpha} \Pi_\beta^{\mathcal{P}}(T)$ such that $\psi \models_{\mathcal{W}} \varphi$.*

Also note the following closure property of the Π_α levels of the hierarchies:

Lemma 5. *For all $\alpha \neq 0$ and $T \in \mathcal{T}$, $\Pi_\alpha^{\mathcal{P}}(T)$ is closed under (finite) disjunction and (finite) conjunction.*

When we expand the set of possible observations, we cannot lose any generalized logical consequence on any level of the hierarchy, assuming that we are dealing with ideal paradigms. Recall that in standard paradigms, the set of possible theories is uniquely determined by the set of possible worlds and the set of possible observations, and that ideal paradigms are standard.

Proposition 1. *Suppose that $\mathcal{P}$ is ideal. Let $\mathcal{P}'$ be a standard paradigm of the form $(S, \mathcal{L}, \mathcal{W}, \mathcal{O}', \mathcal{T}')$ with $\mathcal{O}$ included in $\mathcal{O}'$. Then for all ordinals α and $\mathfrak{M} \in \mathcal{W}$, $\Sigma_\alpha^{\mathcal{P}}(D_{\mathcal{O}}(\mathfrak{M})) \subseteq \Sigma_\alpha^{\mathcal{P}'}(D_{\mathcal{O}'}(\mathfrak{M}))$ and $\Pi_\alpha^{\mathcal{P}}(D_{\mathcal{O}}(\mathfrak{M})) \subseteq \Pi_\alpha^{\mathcal{P}'}(D_{\mathcal{O}'}(\mathfrak{M}))$.*

Proof. Proof is by induction. Since $\mathcal{O} \subseteq \mathcal{O}'$, $D_{\mathcal{O}}(\mathfrak{M}) \subseteq D_{\mathcal{O}'}(\mathfrak{M})$ for all $\mathfrak{M} \in \mathcal{W}$. Let $\alpha \neq 0$ be given, and suppose that $\Pi_\beta^{\mathcal{P}}(D_{\mathcal{O}}(\mathfrak{M})) \subseteq \Pi_\beta^{\mathcal{P}'}(D_{\mathcal{O}'}(\mathfrak{M}))$ for all $\beta < \alpha$ and $\mathfrak{M} \in \mathcal{W}$. Definition 4 implies immediately that for all $\mathfrak{M} \in \mathcal{W}$, $\Sigma_\alpha^{\mathcal{P}}(D_{\mathcal{O}}(\mathfrak{M})) \subseteq \Sigma_\alpha^{\mathcal{P}'}(D_{\mathcal{O}'}(\mathfrak{M}))$. Let $\mathfrak{M} \in \mathcal{W}$ and $\varphi \in \Pi_\alpha^{\mathcal{P}}(D_{\mathcal{O}}(\mathfrak{M}))$ be given. By Lemma 3, choose $\psi \in \Sigma_\alpha^{\mathcal{P}}(D_{\mathcal{O}}(\mathfrak{M}))$ such that for all $T \in \mathcal{T}$ with $T \models_{\mathcal{W}}^{\mathcal{O}} \psi$ and $T \not\models_{\mathcal{W}}^{\mathcal{O}} \varphi$, $\neg\varphi \in \Sigma_\alpha^{\mathcal{P}}(T)$. Let $\mathfrak{N} \in \mathcal{W}$ with $D_{\mathcal{O}'}(\mathfrak{N}) \models_{\mathcal{W}}^{\mathcal{O}'} \psi$ and $D_{\mathcal{O}'}(\mathfrak{N}) \not\models_{\mathcal{W}}^{\mathcal{O}'} \varphi$ be given. Since $\mathcal{P}$ is ideal and $\mathcal{O} \subseteq \mathcal{O}'$, $\mathcal{P}'$ is ideal as well, and we infer that $D_{\mathcal{O}}(\mathfrak{N}) \models_{\mathcal{W}}^{\mathcal{O}} \psi$ and $D_{\mathcal{O}}(\mathfrak{N}) \not\models_{\mathcal{W}}^{\mathcal{O}} \varphi$. Hence $\neg\varphi \in \Sigma_\alpha^{\mathcal{P}}(D_{\mathcal{O}}(\mathfrak{N}))$ which, we have seen, implies that $\neg\varphi \in \Sigma_\alpha^{\mathcal{P}'}(D_{\mathcal{O}'}(\mathfrak{N}))$. Since the same part of the proof implies that $\psi \in \Sigma_\alpha^{\mathcal{P}'}(D_{\mathcal{O}'}(\mathfrak{M}))$, we conclude with Lemma 3 that $\varphi \in \Pi_\alpha^{\mathcal{P}'}(D_{\mathcal{O}'}(\mathfrak{M}))$.

When learning from just positive data versus learning from both positive and negative data, the corresponding paradigms are still more strongly related. First we need a definition.

Definition 6. *Let ordinal α and $\varphi \in \mathcal{L}$ be given. We say that φ is Σ_α* in $\mathcal{P}$ *just in case for all $T \in \mathcal{T}$, if $T \models_{\mathcal{W}}^{\mathcal{O}} \varphi$ then $\varphi \in \Sigma_\alpha^{\mathcal{P}}(T)$. We say that φ is Π_α* in $\mathcal{P}$ *just in case for all $T \in \mathcal{T}$, if $T \models_{\mathcal{W}}^{\mathcal{O}} \varphi$ then $\varphi \in \Pi_\alpha^{\mathcal{P}}(T)$.*

Trivially, all sentences which are logically equivalent to the negation of a member of $\mathcal{O}$ are Π_1 in $\mathcal{P}$. We can then apply the following to $\mathcal{O}$ equal to the set of positive data only, and $\mathcal{O}'$ equal to the set of both positive and negative data, to see how formulas that are generalized logical consequences of a given theory T can be raised from some level in the hierarchy over T defined from $\mathcal{O}'$ to some level above in the hierarchy over T defined from $\mathcal{O}$.

Proposition 2. *Let paradigm $\mathcal{P}'$ be ideal and of the form $(S, \mathcal{L}, \mathcal{W}, \mathcal{O}', \mathcal{T}')$ with $\mathcal{O}'$ closed under $\sim$. Let ordinal α be given, and suppose that all members of $\mathcal{O}'$ are Π_α in $\mathcal{P}$. For all $\mathfrak{M} \in \mathcal{W}$:*

1. $\Sigma_\beta^{\mathcal{P}'}(D_{\mathcal{O}'}(\mathfrak{M})) \subseteq \Sigma_{\alpha+\beta}^{\mathcal{P}}(D_{\mathcal{O}}(\mathfrak{M}))$ *for all nonnull ordinals β;*
2. $\Pi_\beta^{\mathcal{P}'}(D_{\mathcal{O}'}(\mathfrak{M})) \subseteq \Pi_{\alpha+\beta}^{\mathcal{P}}(D_{\mathcal{O}}(\mathfrak{M}))$ *for all ordinals β.*

Proof. Without loss of generality we can assume that all members of $\mathcal{O}'$ are satisfiable in $\mathcal{W}$. If $\alpha = 0$ then $\mathcal{O}' \subseteq \mathcal{O}$, and we conclude immediately with Proposition 1, so suppose $\alpha \neq 0$. Proof is by induction on β. Since all members of $\mathcal{O}'$ are Π_α in $\mathcal{P}$, $D_{\mathcal{O}'}(\mathfrak{M}) \subseteq \Pi_\alpha^{\mathcal{P}}(D_{\mathcal{O}}(\mathfrak{M}))$ for all $\mathfrak{M} \in \mathcal{W}$, which proves the second inclusion in the statement of the proposition for $\beta = 0$.

Let us verify that $\mathcal{P}$ is ideal. Let $\mathfrak{M} \in \mathcal{W}$ be given. By the preceding relation $D_{\mathcal{O}}(\mathfrak{M}) \models_{\mathcal{W}}^{\mathcal{O}} D_{\mathcal{O}'}(\mathfrak{M})$, hence for all $\varphi \in \mathcal{L}$, if $D_{\mathcal{O}'}(\mathfrak{M}) \models_{\mathcal{W}} \varphi$ then $D_{\mathcal{O}}(\mathfrak{M}) \models_{\mathcal{W}}^{\mathcal{O}} \varphi$. Since $\mathcal{P}'$ is ideal, $D_{\mathcal{O}'}(\mathfrak{M}) \models_{\mathcal{W}}^{\mathcal{O}'} \varphi$ or $D_{\mathcal{O}'}(\mathfrak{M}) \models_{\mathcal{W}}^{\mathcal{O}'} \neg\varphi$ for all sentences φ. Since $\mathcal{O}'$ is closed under $\sim$, $\{\varphi \in \mathcal{L} \mid D_{\mathcal{O}'}(\mathfrak{M}) \models_{\mathcal{W}} \varphi\} = \{\varphi \in \mathcal{L} \mid D_{\mathcal{O}'}(\mathfrak{M}) \models_{\mathcal{W}}^{\mathcal{O}'} \varphi\}$ by Lemma 1. We infer immediately that $D_{\mathcal{O}}(\mathfrak{M}) \models_{\mathcal{W}}^{\mathcal{O}} \varphi$ or $D_{\mathcal{O}}(\mathfrak{M}) \models_{\mathcal{W}}^{\mathcal{O}} \neg\varphi$ for all sentences φ. Hence $\mathcal{P}$ is ideal.

Let $\beta \neq 0$ be given. Suppose that for all $\gamma < \beta$ and $\mathfrak{M} \in \mathcal{W}$, $\Pi_\gamma^{\mathcal{P}'}(D_{\mathcal{O}'}(\mathfrak{M}))$ is a subset of $\Pi_{\alpha+\gamma}^{\mathcal{P}}(D_{\mathcal{O}}(\mathfrak{M}))$. Let $\mathfrak{M} \in \mathcal{W}$ and $\varphi \in \Sigma_\beta^{\mathcal{P}'}(D_{\mathcal{O}'}(\mathfrak{M}))$ be given. Assume that $\beta = 1$. By Lemma 2, there exists finite $E \subseteq D_{\mathcal{O}'}(\mathfrak{M})$ such that $E \models_{\mathcal{W}} \varphi$. Since $\mathcal{P}$ is ideal, $D_{\mathcal{O}}(\mathfrak{M}) \models_{\mathcal{W}}^{\mathcal{O}} E$. As all members of E are Π_α in $\mathcal{P}$, we infer that $E \subseteq \Pi_\alpha^{\mathcal{P}}(D_{\mathcal{O}}(\mathfrak{M}))$, hence $\bigwedge E \in \Pi_\alpha^{\mathcal{P}}(D_{\mathcal{O}}(\mathfrak{M}))$ by Lemma 5. This, the fact that $E \models_{\mathcal{W}} \varphi$, and Lemma 4 then imply that $\varphi \in \Sigma_{\alpha+\beta}^{\mathcal{P}}(D_{\mathcal{O}}(\mathfrak{M}))$. If $\beta > 1$, if follows easily from Lemma 4 and the induction hypothesis that $\varphi \in \Sigma_{\alpha+\beta}^{\mathcal{P}}(D_{\mathcal{O}}(\mathfrak{M}))$. So we have shown that for all $\mathfrak{M} \in \mathcal{W}$, $\Sigma_\beta^{\mathcal{P}'}(D_{\mathcal{O}'}(\mathfrak{M})) \subseteq \Sigma_{\alpha+\beta}^{\mathcal{P}}(D_{\mathcal{O}}(\mathfrak{M}))$. Let $\mathfrak{M} \in \mathcal{W}$ and $\varphi \in \Pi_\beta^{\mathcal{P}'}(D_{\mathcal{O}'}(\mathfrak{M}))$ be given. To complete the proof we show that φ belongs to $\Pi_{\alpha+\beta}^{\mathcal{P}}(D_{\mathcal{O}}(\mathfrak{M}))$. By Lemma 3, choose $\psi \in \Sigma_\beta^{\mathcal{P}'}(D_{\mathcal{O}'}(\mathfrak{M}))$ such that for all $T \in \mathcal{T}'$ with $T \models_{\mathcal{W}}^{\mathcal{O}'} \psi$ and $T \not\models_{\mathcal{W}}^{\mathcal{O}'} \varphi$, $\neg\varphi \in \Sigma_\beta^{\mathcal{P}'}(T)$. Let $\mathfrak{N} \in \mathcal{W}$ with $D_{\mathcal{O}}(\mathfrak{N}) \models_{\mathcal{W}}^{\mathcal{O}} \psi$ and $D_{\mathcal{O}}(\mathfrak{N}) \not\models_{\mathcal{W}}^{\mathcal{O}} \varphi$ be given. Since $\mathcal{P}$ and $\mathcal{P}'$ are ideal, we infer that $D_{\mathcal{O}'}(\mathfrak{N}) \models_{\mathcal{W}}^{\mathcal{O}'} \psi$ and $D_{\mathcal{O}'}(\mathfrak{N}) \not\models_{\mathcal{W}}^{\mathcal{O}'} \varphi$. Hence $\neg\varphi \in \Sigma_\beta^{\mathcal{P}'}(D_{\mathcal{O}'}(\mathfrak{N}))$, so $\neg\varphi \in \Sigma_{\alpha+\beta}^{\mathcal{P}}(D_{\mathcal{O}}(\mathfrak{N}))$ as proved above. Since the same part of the proof also shows that $\psi \in \Sigma_{\alpha+\beta}^{\mathcal{P}}(D_{\mathcal{O}}(\mathfrak{M}))$, we conclude with Lemma 3 that $\varphi \in \Pi_{\alpha+\beta}^{\mathcal{P}}(D_{\mathcal{O}}(\mathfrak{M}))$.

5 Connections with other hierarchies

In order to be able to establish relations between the hierarchies of generalized logical consequences and other hierarchies, we define still a new hierarchy, where the Σ_α's levels are better behaved:

Definition 7. *For all ordinals α, the sets of sentences $\widetilde{\Sigma}_\alpha^{\mathcal{P}}$ and $\widetilde{\Pi}_\alpha^{\mathcal{P}}$ are defined by induction on α, as follows.*

1. $\widetilde{\Sigma}_0^{\mathcal{P}} = \mathcal{O} \cup \{\sim \varphi \mid \varphi \in \mathcal{O}\}$.

2. *For all ordinals* α, $\widetilde{\Pi}^{\mathcal{P}}_{\alpha} = \{\sim \varphi \mid \varphi \in \widetilde{\Sigma}^{\mathcal{P}}_{\alpha}\}$.
3. *Let* $\alpha \neq 0$ *be given. A sentence* φ *belongs to* $\widetilde{\Sigma}^{\mathcal{P}}_{\alpha}$ *iff there exists* $\beta < \alpha$ *with the following property. For all* $\mathfrak{M} \in \mathcal{W}$ *with* $\mathfrak{M} \models \varphi$, *there exists finite* $D \subseteq \widetilde{\Pi}^{\mathcal{P}}_{\beta}$ *such that* $\mathfrak{M} \models D$ *and* $D \models_{\mathcal{W}} \varphi$.

Informally, we will refer to the hierarchy defined above as the *uniform* hierarchy. We will need the following properties. First note that the levels of the uniform hierarchy are ordered as expected in a hierarchy of a Borel type.

Proposition 3. *For all ordinals* α, β *if* $\alpha < \beta$ *then* $\widetilde{\Sigma}^{\mathcal{P}}_{\alpha} \cup \widetilde{\Pi}^{\mathcal{P}}_{\alpha} \subseteq \widetilde{\Sigma}^{\mathcal{P}}_{\beta} \cap \widetilde{\Pi}^{\mathcal{P}}_{\beta}$.

Proof. Trivially, $\widetilde{\Sigma}^{\mathcal{P}}_{0} = \widetilde{\Pi}^{\mathcal{P}}_{0}$, which implies immediately that for all ordinals β, $\widetilde{\Sigma}^{\mathcal{P}}_{0} \cup \widetilde{\Pi}^{\mathcal{P}}_{0} \subseteq \widetilde{\Sigma}^{\mathcal{P}}_{\beta} \cap \widetilde{\Pi}^{\mathcal{P}}_{\beta}$. Let $\alpha \neq 0$ be given. For all $\beta > \alpha$, the inclusions $\widetilde{\Sigma}^{\mathcal{P}}_{\alpha} \subseteq \widetilde{\Sigma}^{\mathcal{P}}_{\beta}$ and $\widetilde{\Pi}^{\mathcal{P}}_{\alpha} \subseteq \widetilde{\Pi}^{\mathcal{P}}_{\beta}$ are straightforward. It is easily verified that $\widetilde{\Sigma}^{\mathcal{P}}_{\alpha} \subseteq \widetilde{\Pi}^{\mathcal{P}}_{\alpha+1}$. We infer that any $\varphi \in \widetilde{\Pi}^{\mathcal{P}}_{\alpha}$ is equal to $\sim \psi$ for some $\psi \in \widetilde{\Sigma}^{\mathcal{P}}_{\alpha}$, hence $\sim \varphi \in \widetilde{\Pi}^{\mathcal{P}}_{\alpha+1}$, hence $\varphi \in \widetilde{\Sigma}^{\mathcal{P}}_{\alpha+1}$. So $\widetilde{\Pi}^{\mathcal{P}}_{\alpha} \subseteq \widetilde{\Sigma}^{\mathcal{P}}_{\alpha+1}$. The result follows.

Remember that $\mathcal{L}$ is just a fragment of $\mathcal{L}^{S}_{\omega_1\omega}$, hence does not contain the disjunction or conjunction of any of its countable subsets. So $\bigvee X$ and $\bigwedge X$ in the closure property below are members of $\mathcal{L}^{S}_{\omega_1\omega}$, but not necessarily members of $\mathcal{L}$.

Lemma 6. *Let* $\alpha \neq 0$, *sentence* φ, *and countable* $X \subseteq \mathcal{L}$ *be given.*

1. *If* $X \subseteq \widetilde{\Sigma}^{\mathcal{P}}_{\alpha}$ *and* $\models_{\mathcal{W}} (\varphi \leftrightarrow \bigvee X)$ *then* $\varphi \in \widetilde{\Sigma}^{\mathcal{P}}_{\alpha}$.
2. *If* $X \subseteq \widetilde{\Pi}^{\mathcal{P}}_{\alpha}$ *and* $\models_{\mathcal{W}} (\varphi \leftrightarrow \bigwedge X)$ *then* $\varphi \in \widetilde{\Pi}^{\mathcal{P}}_{\alpha}$.

Adding the requirement that $\mathcal{W}$ consists exclusively of Henkin structures enables to treat existential quantifiers as countable disjunctions, and universal quantifiers as countable conjunctions.

Corollary 1. *Suppose that* $\mathcal{W}$ *is a set of Henkin structures. Let formula* φ *with free variables* $x_1, \ldots, x_n$ *be given. Denote by* X *the set of all sentences of the form* $\varphi[t_1/x_1, \ldots t_n/x_n]$ *for some closed terms* $t_1, \ldots t_n$. *Let* $\alpha \neq 0$ *be given.*

1. *If* $X \subseteq \widetilde{\Sigma}^{\mathcal{P}}_{\alpha}$ *then* $\exists x_1 \ldots \exists x_n \varphi \in \widetilde{\Sigma}^{\mathcal{P}}_{\alpha}$.
2. *If* $X \subseteq \widetilde{\Pi}^{\mathcal{P}}_{\alpha}$ *then* $\forall x_1 \ldots \forall x_n \varphi \in \widetilde{\Pi}^{\mathcal{P}}_{\alpha}$.

We characterize $\widetilde{\Sigma}^{\mathcal{P}}_{1}$ and $\widetilde{\Pi}^{\mathcal{P}}_{1}$, assuming that $\mathcal{O}$ is closed under the $\sim$ operator.

Proposition 4. *Suppose that* $\mathcal{O}$ *is closed under* $\sim$. *Let sentence* φ *be given.*

1. $\varphi \in \widetilde{\Sigma}^{\mathcal{P}}_{1}$ *if and only if* $\models_{\mathcal{W}} \varphi$, *or* $\models_{\mathcal{W}} \neg\varphi$, *or there is a nonempty set* X *of finite, nonempty subsets of* $\mathcal{O}$ *such that* $\models_{\mathcal{W}} \bigvee\{\bigwedge D \mid D \in X\} \leftrightarrow \varphi$.
2. $\varphi \in \widetilde{\Pi}^{\mathcal{P}}_{1}$ *if and only if* $\models_{\mathcal{W}} \varphi$, *or* $\models_{\mathcal{W}} \neg\varphi$, *or there is a nonempty set* X *of finite, nonempty subsets of* $\mathcal{O}$ *such that* $\models_{\mathcal{W}} \bigwedge\{\bigvee D \mid D \in X\} \leftrightarrow \varphi$.

Proof. The proof is trivial if $\models_{\mathcal{W}} \varphi$ or $\models_{\mathcal{W}} \neg\varphi$, so suppose otherwise. Assume that $\varphi \in \widetilde{\Sigma}^{\mathcal{P}}_1$. Let $\mathfrak{M} \in \mathcal{W}$ be such that $\mathfrak{M} \models \varphi$. Choose a nonempty subset $D_{\mathfrak{M}}$ of $D_{\mathcal{O}}(\mathfrak{M})$ such that $D_{\mathfrak{M}} \models_{\mathcal{W}} \varphi$. Set $X = \{D_{\mathfrak{M}} \mid \mathfrak{M} \in \mathcal{W} \text{ and } \mathfrak{M} \models \varphi\}$. Then X is nonempty, and it is easy to verify that $\models_{\mathcal{W}} \bigvee\{\bigwedge D \mid D \in X\} \leftrightarrow \varphi$. Conversely, let nonempty set X of finite, nonempty subsets of $\mathcal{O}$ be such that $\models_{\mathcal{W}} \bigvee\{\bigwedge D \mid D \in X\} \leftrightarrow \varphi$. Since $\bigwedge D \in \widetilde{\Sigma}^{\mathcal{P}}_1$ for all $D \in X$, it follows from Lemma 6 that $\varphi \in \widetilde{\Sigma}^{\mathcal{P}}_1$. We conclude that *1.* holds, and *2.* is an immediate consequence.

Then we characterize the other levels:

Proposition 5. *Let $\alpha > 1$ and sentence φ be given.*

1. *$\varphi \in \widetilde{\Sigma}^{\mathcal{P}}_\alpha$ iff there is nonempty $X \subseteq \bigcup_{\beta<\alpha}(\widetilde{\Sigma}^{\mathcal{P}}_\beta \cup \widetilde{\Pi}^{\mathcal{P}}_\beta)$ such that $\models_{\mathcal{W}} \bigvee X \leftrightarrow \varphi$.*
2. *$\varphi \in \widetilde{\Pi}^{\mathcal{P}}_\alpha$ iff there is nonempty $X \subseteq \bigcup_{\beta<\alpha}(\widetilde{\Sigma}^{\mathcal{P}}_\beta \cup \widetilde{\Pi}^{\mathcal{P}}_\beta)$ such that $\models_{\mathcal{W}} \bigwedge X \leftrightarrow \varphi$.*

Proof. We assume that $\not\models_{\mathcal{W}} \neg\varphi$ (otherwise, the proof is trivial). Suppose that $\varphi \in \widetilde{\Sigma}^{\mathcal{P}}_\alpha$. Let $\mathfrak{M} \in \mathcal{W}$ be such that $\mathfrak{M} \models \varphi$. By Proposition 3, and since $\widetilde{\Pi}^{\mathcal{P}}_\beta$ is trivially closed under conjunction, we can choose $\beta < \alpha$ and $\psi_{\mathfrak{M}} \in \widetilde{\Pi}^{\mathcal{P}}_\beta$ such that $\psi_{\mathfrak{M}} \models_{\mathcal{W}} \varphi$. Set $X = \{\psi_{\mathfrak{M}} \mid \mathfrak{M} \in \mathcal{W} \text{ and } \mathfrak{M} \models \varphi\}$. Then $X \neq \emptyset$, and it is easy to verify that $\models_{\mathcal{W}} \bigvee X \leftrightarrow \varphi$. Conversely, let nonempty $X \subseteq \bigcup_{\beta<\alpha}(\widetilde{\Sigma}^{\mathcal{P}}_\beta \cup \widetilde{\Pi}^{\mathcal{P}}_\beta)$ be such that $\models_{\mathcal{W}} \bigvee X \leftrightarrow \varphi$. Then $X \subseteq \widetilde{\Sigma}^{\mathcal{P}}_\alpha$ by Proposition 3 again, and Lemma 6 implies that $\varphi \in \widetilde{\Sigma}^{\mathcal{P}}_\alpha$. Hence *1.* holds, and *2.* is an immediate consequence.

The uniform hierarchy is more easily investigated than the hierarchies of generalized logical consequences. But our aim is to gain more insights into the latter, by transferring properties of the former. For this to be possible, we need to establish relations between the uniform hierarchy and the hierarchies of generalized logical consequences. In one direction we have:

Proposition 6. *Suppose that $\mathcal{O}$ is closed under $\sim$ and $\mathcal{P}$ is ideal. Then for all ordinals α and $T \in \mathcal{T}$, $\widetilde{\Sigma}^{\mathcal{P}}_\alpha(T) \subseteq \Sigma^{\mathcal{P}}_\alpha(T)$ and $\widetilde{\Pi}^{\mathcal{P}}_\alpha(T) \subseteq \Pi^{\mathcal{P}}_\alpha(T)$.*

Proof. Proof is by induction. Since $\mathcal{O}$ is closed under $\sim$ and $\mathcal{P}$ is standard, $\widetilde{\Sigma}^{\mathcal{P}}_0(T) = \Sigma^{\mathcal{P}}_0(T) = \widetilde{\Pi}^{\mathcal{P}}_0(T) = \Pi^{\mathcal{P}}_0(T) = T$, for all $T \in \mathcal{T}$. Let $\alpha \neq 0$ be given. Let $T \in \mathcal{T}$ be given, and suppose that $\widetilde{\Pi}^{\mathcal{P}}_\beta(T)$ is included in $\Pi^{\mathcal{P}}_\beta(T)$ for all $\beta < \alpha$. Let $\varphi \in \widetilde{\Sigma}^{\mathcal{P}}_\alpha(T)$ be given. Choose $\beta < \alpha$ and finite $D \subseteq \widetilde{\Pi}^{\mathcal{P}}_\beta$ such that $T \models^{\mathcal{O}}_{\mathcal{W}} D$ and $D \models_{\mathcal{W}} \varphi$. By induction hypothesis, $D \subseteq \Pi^{\mathcal{P}}_\beta(T)$, which implies immediately that $\varphi \in \Sigma^{\mathcal{P}}_\alpha(T)$. Now suppose that for all $T \in \mathcal{T}$, $\widetilde{\Sigma}^{\mathcal{P}}_\alpha(T) \subseteq \Sigma^{\mathcal{P}}_\alpha(T)$. Let $T \in \mathcal{T}$ and $\varphi \in \widetilde{\Pi}^{\mathcal{P}}_\alpha(T)$ be given. If for all $T' \in \mathcal{T}$, $T' \models^{\mathcal{O}}_{\mathcal{W}} \varphi$, then trivially $\varphi \in \Pi^{\mathcal{P}}_\alpha(T)$. Suppose there exists $T' \in \mathcal{T}$ such that $T' \not\models^{\mathcal{O}}_{\mathcal{W}} \varphi$. Since $\mathcal{P}$ is ideal, $T' \models^{\mathcal{O}}_{\mathcal{W}} \neg\varphi$. Hence $\neg\varphi \in \widetilde{\Sigma}^{\mathcal{P}}_\alpha(T')$. So by inductive hypothesis, $\neg\varphi$ belongs to $\Sigma^{\mathcal{P}}_\alpha(T')$, and we conclude that $\varphi \in \Pi^{\mathcal{P}}_\alpha(T)$.

The other direction of the relationship between the uniform hierarchy and the hierarchies of generalized logical consequences is more subtle. It is given by following proposition, left without proof for lack of space.

Proposition 7. *Suppose that $\mathcal{P}$ is ideal. Let $\alpha > 1$, $T \in \mathcal{T}$, and $\varphi \in \Sigma^{\mathcal{P}}_{\alpha}(T)$ be given. There exists $\beta < \alpha$ and $\psi \in \widetilde{\Pi}^{\mathcal{P}}_{\beta}(T)$ such that $\psi \models_{\mathcal{W}} \varphi$.*

The natural connection between the hierarchies of generalized logical consequences (for paradigms of a special form) and formula complexity is given by the result below, together with Proposition 6. Recall that a basic formula is an atomic formula or the negation of an atomic formula. Also note that the assumptions of the proposition below inply that $\mathcal{P}$ is ideal.

Proposition 8. *Suppose that $\mathcal{W}$ is a set of Henkin structures, $\mathcal{O}$ is the set of all basic sentences, and $\mathcal{P}$ is standard. Let $\alpha \neq 0$ be given. Every sentence in Σ_{α} (respect. Π_{α}) prenex form belongs to $\widetilde{\Sigma}^{\mathcal{P}}_{\alpha}$ (respect. $\widetilde{\Pi}^{\mathcal{P}}_{\alpha}$).*

Proof. Proof is by double induction on ordinals. Suppose that for all nonnull $\gamma < \alpha$, every sentence in Σ_{γ} prenex form belongs to $\widetilde{\Sigma}^{\mathcal{P}}_{\gamma}$ (hence every sentence in Π_{γ} prenex form belongs to $\widetilde{\Pi}^{\mathcal{P}}_{\gamma}$). Let sentence φ and ordinal β be such that β is the height of φ and φ is in Σ_{α} prenex form. Suppose that for all $\gamma < \beta$, every sentence of height γ which is in Σ_{α} prenex form belongs to $\widetilde{\Sigma}^{\mathcal{P}}_{\alpha}$. We now distinguish the cases corresponding to the definition of a formula being in Σ_{α} prenex form. Assume that φ is in Σ_{γ} or Π_{γ} prenex form for some $\gamma < \alpha$. If γ is nonnull then φ belongs to $\widetilde{\Sigma}^{\mathcal{P}}_{\gamma}$ or $\widetilde{\Pi}^{\mathcal{P}}_{\gamma}$ by inductive hypothesis, hence to $\widetilde{\Sigma}^{\mathcal{P}}_{\alpha}$. Suppose that $\gamma = 0$. Then φ is a quantifier free member of $\mathcal{L}^{S}_{\omega\omega}$. Let $\mathfrak{M} \in \mathcal{W}$ be such that $\mathfrak{M} \models \varphi$. Let D be the set of all basic sentences ψ such that either ψ or $\sim \psi$ occurs in φ, and $\mathfrak{M} \models \psi$. By the choice of $\mathcal{O}$ and the fact that $\mathcal{P}$ is standard, it is clear that $D \subseteq D_{\mathcal{O}}(\mathfrak{M})$ and $D \models \varphi$. Hence φ belongs to $\widetilde{\Sigma}^{\mathcal{P}}_{1}$, hence to $\widetilde{\Sigma}^{\mathcal{P}}_{\alpha}$. Assume that φ is of form $\exists x\psi$ for some variable x and $\psi \in \mathcal{L}$ which is in Σ_{α} prenex form. Set $X = \{\psi[t/x] \mid t \text{ closed term}\}$. Then X consists of sentences whose heights are smaller than β and which are in Σ_{α} prenex form, and it follows from Corollary 1 that $\varphi \in \widetilde{\Sigma}^{\mathcal{P}}_{\alpha}$. Suppose that φ is of form $\bigvee X$ for some (countable) set X of formulas which are in Σ_{α} prenex form. Then X consists of sentences whose height is smaller than β and which are in Σ_{α} prenex form, and it follows from Lemma 6 that $\varphi \in \widetilde{\Sigma}^{\mathcal{P}}_{\alpha}$.

Using some connection with the Borel hierarchy (see [18] for definitions and properties), we can exhibit natural paradigms such that the associated uniform hierarchy does not collapse to a finite level:

Proposition 9. *There exists a finite, equality free vocabulary V and a subset T of $\mathcal{L}^{V}_{\omega\omega}$ with the following property. Suppose that $S = V$, $\mathcal{W}$ is the set of Herbrand models of T, $\mathcal{O}$ is the set of basic sentences, and $\mathcal{P}$ is standard. Then for all nonnull $n \in \mathbb{N}$, $\widetilde{\Pi}^{\mathcal{P}}_{n} \setminus \widetilde{\Sigma}^{\mathcal{P}}_{n}$ contains a Π_n formula in $\mathcal{L}^{S}_{\omega\omega}$.*

Proof. Given $\sigma \in \{0,1\}^{<\mathbb{N}}$, the basic open set of the Cantor space consisting of members $\boldsymbol{c} \in \{0,1\}^{\mathbb{N}}$ that extend σ is denoted O_{σ}. For all nonnull $n \in \mathbb{N}$, choose total mapping $h_n : \mathbb{N}^n \to \{0,1\}^{<\mathbb{N}}$ such that the set A_n equal to $\bigcap_{i_1 \in \mathbb{N}} \bigcup_{i_2 \in \mathbb{N}} \dots O_{h_n(i_1,i_2,\dots,i_n)}$ is Π_n Borel, but not Σ_n Borel, in the Cantor space. Suppose that $S = \{\overline{0}, s, \langle\rangle, P\}$ where $\overline{0}$ is a constant, s a unary function symbol, $\langle\rangle$ a binary function symbol, and P a unary predicate symbol. Given nonnull $n \in \mathbb{N}$, we denote by $\overline{n}$ the term obtained from $\overline{0}$ by n applications of s. Given $n > 2$ and terms $t_1, t_2, \ldots, t_n$, we denote by $\langle t_1, t_2, \ldots, t_n\rangle$ the term $\langle t_1, \langle t_2, \ldots, \langle t_{n-1}, t_n\rangle \ldots\rangle\rangle$. Choose a bijective mapping f from the set of closed terms into $\mathbb{N}$. Choose a surjective mapping g from the set of closed terms into $\{0,1\}^{<\mathbb{N}}$ such that for all nonnull $n \in \mathbb{N}$ and closed terms $t_1, \ldots, t_n$, $g(\langle \overline{n}, t_1, \ldots, t_n\rangle) = h_n(f(t_1), \ldots, f(t_n))$. Given $\boldsymbol{c} \in \{0,1\}^{\mathbb{N}}$, we define $\mathfrak{M}_{\boldsymbol{c}}$ to be the unique Herbrand structure which satisfies:

$$(*) \qquad \text{for every closed term } t,\ \mathfrak{M}_{\boldsymbol{c}} \models Pt \text{ iff } \boldsymbol{c} \in O_{g(t)}.$$

Suppose that $\mathcal{W} = \{\mathfrak{M}_{\boldsymbol{c}} \mid \boldsymbol{c} \in \{0,1\}^{\mathbb{N}}\}$, $\mathcal{O}$ is the set of closed basic formulas, and $\mathcal{P}$ is standard. Since g is surjective, $\mathcal{W}$ clearly consists of the Herbrand models of the set of:

1. all formulas of the form Pt, where t is a closed term with $g(t) = ()$;
2. all formulas of the form $Pt \to \neg Pt'$, where t, t' are closed terms such that $g(t) \not\subseteq g(t')$ and $g(t') \not\subseteq g(t)$;
3. all formulas of the form $Pt \to Pt_0 \vee Pt_1$, where t, t_0, t_1 are closed terms such that $g(t_0) = g(t) \star 0$ and $g(t_1) = g(t) \star 1$.

For all nonnull $n \in \mathbb{N}$, set $\varphi_n = \forall x_1 \exists x_2 \ldots P(\langle \overline{n}, x_1, x_2, \ldots, x_n\rangle)$. Since f is bijective, it follows easily from the definition of g that for all members $\boldsymbol{c}$ of $\{0,1\}^{<\mathbb{N}}$ and nonnull $n \in \mathbb{N}$, $\mathfrak{M}_{\boldsymbol{c}} \models \varphi_n$ iff $\boldsymbol{c}$ belongs to A_n. Let $\Phi : \mathcal{L}^S_{\omega_1\omega} \to \{0,1\}^{\mathbb{N}}$ be such that:

- for all closed terms t, $\Phi(P(t)) = O_{g(t)}$;
- for all $\varphi, \psi \in \mathcal{L}^S_{\omega_1\omega}$, $\Phi(\varphi \vee \psi) = \Phi(\varphi) \cup \Phi(\psi)$ and $\Phi(\varphi \wedge \psi) = \Phi(\varphi) \cap \Phi(\psi)$;
- for all nonempty sets X of members of $\mathcal{L}^S_{\omega_1\omega}$, $\Phi(\bigvee X) = \bigcup_{\varphi \in X} \Phi(\varphi)$ and $\Phi(\bigwedge X) = \bigcap_{\varphi \in X} \Phi(\varphi)$.

Using the definition of Φ and $(*)$ above, it is easy to show by induction that for all nonnull $n \in \mathbb{N}$, the following holds. Let φ be a member of $\widetilde{\Sigma}^{\mathcal{P}}_n$ (respect. $\widetilde{\Pi}^{\mathcal{P}}_n$). There exists a closed member φ^* of $\mathcal{L}^S_{\omega_1\omega}$ built from $\mathcal{O}$ using $\vee, \wedge, \bigvee, \bigwedge$ only, and such that:

- $\Phi(\varphi^*)$ is Σ_n (respect. Π_n) Borel in the Cantor space;
- for all $\boldsymbol{c} \in \{0,1\}^{\mathbb{N}}$, $\mathfrak{M}_{\boldsymbol{c}} \models \varphi$ iff $\boldsymbol{c} \in \Phi(\varphi^*)$.

Let nonnull $n \in \mathbb{N}$ be given. By Proposition 8, φ_n belongs to $\widetilde{\Pi}^{\mathcal{P}}_n$. Suppose for a contradiction that $\varphi_n \in \widetilde{\Sigma}^{\mathcal{P}}_n$. Then $\Phi((\varphi^{\Sigma}_n)^*)$ is Σ_n Borel in the Cantor space. Moreover, we have seen that for every $\boldsymbol{c} \in \{0,1\}^{\mathbb{N}}$, $\mathfrak{M}_{\boldsymbol{c}} \models \varphi_n$ iff $\boldsymbol{c} \in \Phi((\varphi_n)^*)$, and for every $\boldsymbol{c} \in \{0,1\}^{\mathbb{N}}$, $\mathfrak{M}_{\boldsymbol{c}} \models \varphi_n$ iff $\boldsymbol{c} \in A_n$. Hence $A_n = \Phi((\varphi_n)^*)$, which contradicts the fact that A_n it not Σ_n Borel in the Cantor space.

The result in Proposition 9 can be transferred to the hierarchies of generalized logical consequences, thanks to the following corollary to Proposition 7.

Proposition 10. *Suppose that $\mathcal{O}$ is closed under $\sim$ and $\mathcal{P}$ is ideal. Let $\alpha \neq 0$ and $\varphi \in \mathcal{L}$ be such that $\varphi \in \widetilde{\Pi}^{\mathcal{P}}_{\alpha} \setminus \widetilde{\Sigma}^{\mathcal{P}}_{\alpha}$. Then $\varphi \in \Pi^{\mathcal{P}}_{\alpha}(T) \setminus \Sigma^{\mathcal{P}}_{\alpha}(T)$ for some $T \in \mathcal{T}$.*

Proof. Since $\varphi \notin \widetilde{\Sigma}^{\mathcal{P}}_{\alpha}$, neither $\models_{\mathcal{W}} \varphi$ nor $\models_{\mathcal{W}} \neg\varphi$. Suppose for a contradiction that for all $T \in \mathcal{T}$, if $\varphi \in \Pi^{\mathcal{P}}_{\alpha}(T)$ then $\varphi \in \Sigma^{\mathcal{P}}_{\alpha}(T)$. So for all $T \in \mathcal{T}$, if $T \models_{\mathcal{W}} \varphi$ then $\varphi \in \Sigma^{\mathcal{P}}_{\alpha}(T)$. We distinguish two cases.

1st Case: $\alpha = 1$. By Lemma 2, given $T \in \mathcal{T}$ with $\varphi \in \Sigma^{\mathcal{P}}_{\alpha}(T)$, we can choose finite, nonempty $D_T \subseteq T$ such that $D_T \models_{\mathcal{W}} \varphi$. Let X be the nonempty set of sets of form D_T, for all $T \in \mathcal{T}$ with $T \models_{\mathcal{W}} \varphi$. Clearly, $\models_{\mathcal{W}} \varphi \leftrightarrow \bigvee\{\bigwedge D_T \mid D_T \in X\}$. It then follows from Proposition 4 that $\varphi \in \widetilde{\Sigma}^{\mathcal{P}}_{\alpha}$, contradiction.

2nd Case: $\alpha > 1$. By Proposition 7, given $T \in \mathcal{T}$ such that $\varphi \in \Sigma^{\mathcal{P}}_{\alpha}(T)$, we can choose $\beta < \alpha$ and $\psi_T \in \widetilde{\Pi}^{\mathcal{P}}_{\beta}(T)$ such that $\psi_T \models_{\mathcal{W}} \varphi$. Let X be the nonempty set of formulas of form ψ_T, for all $T \in \mathcal{T}$ such that $T \models_{\mathcal{W}} \varphi$. Clearly, $\models_{\mathcal{W}} \varphi \leftrightarrow \bigvee X$. It then follows from Proposition 5 that $\varphi \in \widetilde{\Sigma}^{\mathcal{P}}_{\alpha}$, contradiction.

6 Connections to Learning Theory

Many classical learning paradigms can be cast in this framework; in other words, some learning paradigms are isomorphic to some (standard) paradigms (quintuples of the form $(\mathcal{S}, \mathcal{L}, \mathcal{W}, \mathcal{O}, \mathcal{T})$), and many learnability results are equivalent to statements involving concepts defined from the notion of generalized logical consequence. We give a flavour of the connection, in the form of a few definitions and results left without proof.

The definitions usually given in the numerical setting (see [7]) are immediately adapted to the logical one. Given $T \subseteq \mathcal{L}$, a *text for* T is a sequence of members of T with at least one occurrence of every member of T. If e is a text for T and k an integer, we denote by $e[k]$ the initial segment of e of length k. A *learner* is a partial function from the set of finite sequences of members of $\mathcal{L}$, into the power set of $\mathcal{L}$. (Only uncomputable learners will be considered here, but the following material is easily relativized to computable learners.)

Definition 8 ([4]). *Let learner f and $\varphi \in \mathcal{L}$ be given. We say that f* partially identifies φ in $\mathcal{P}$ *iff for all $T \in \mathcal{T}$ and texts e for T, the following are equivalent:*

1. $T \models^{\mathcal{O}}_{\mathcal{W}} \varphi$.
2. $\{k \in \mathbb{N} \mid \varphi \notin f(e[k])\}$ *is finite.*

If f partially identifies φ in $\mathcal{P}$ and there is no $T \in \mathcal{T}$, no text e for T, and no $k_1, k_2 \in \mathbb{N}$ with $k_1 < k_2$, $\varphi \in f(e[k_1])$, and $\varphi \notin f(e[k_2])$, then we way that f partially identifies φ in $\mathcal{P}$ without mind changes.

Being at level Σ_1 of the hierarchies of generalized logical consequences (which in our view is also equivalent to being a deductive consequence of the underlying theory) and being learnable without mind changes are basically the same notion. No assumption on the paradigm $\mathcal{P}$ is needed:

Proposition 11. *Let sentence φ be given. The following are equivalent.*

1. *φ is Σ_1 in $\mathcal{P}$.*
2. *Some learner partially classifies φ in $\mathcal{P}$ without mind changes.*

Still without assumption on $\mathcal{P}$, it can be verified that level Σ_2 of our hierarchies consist of nothing but formulas that are partially identifiable:

Proposition 12. *There exists a learner f such that for all sentences φ, if φ is Σ_2 in $\mathcal{P}$ then f partially classifies φ in $\mathcal{P}$.*

The converse of Proposition 12 holds for standard paradigms (classical learning paradigms indeed correspond to standard paradigms) when the set of $\mathcal{O}$-diagrams of members of $\mathcal{W}$ is countable, as a consequence of the assumption of the following proposition together with the fact that $\mathcal{L}$ is countable.

Proposition 13. *Set $\sim \mathcal{O} = \{\sim \psi \,|\, \psi \in \mathcal{O}\}$. Suppose that $\mathcal{P}$ is standard and $\bigwedge D_{\sim\mathcal{O}}(\mathfrak{M})$ belongs to $\mathcal{L}$ for all $\mathfrak{M} \in \mathcal{W}$. Let sentence φ be given. The following are equivalent.*

1. *φ is Σ_2 in $\mathcal{P}$.*
2. *Some learner partially classifies φ in $\mathcal{P}$.*
3. *For all $T \in \mathcal{T}$ with $T \models^{\mathcal{O}}_{\mathcal{W}} \varphi$, there exists finite $D \subseteq T$ such that for all $T' \in \mathcal{T}$, if $D \subseteq T' \subseteq T$ then $T' \models^{\mathcal{O}}_{\mathcal{W}} \varphi$.*

Identification in the limit of classes of nonempty languages also has a natural analogue here. Basically, it corresponds to the hierarchies of generalized logical consequences collapsing to level Σ_2. Note the similarity of clause *3.* with the characterization of learnability of classes of nonempty recursive languages from positive data given in [1].

Proposition 14. *Set $\widetilde{\mathcal{O}} = \mathcal{O} \cup \{\sim \psi \,|\, \psi \in \mathcal{O}\}$. Suppose that $\mathcal{P}$ is standard and $\bigwedge D_{\widetilde{\mathcal{O}}}(\mathfrak{M})$ belongs to $\mathcal{L}$ for all $\mathfrak{M} \in \mathcal{W}$. The following are equivalent.*

1. *Every sentence is Σ_2 in $\mathcal{P}$.*
2. *For all $\mathfrak{M} \in \mathcal{W}$, $\bigwedge D_{\widetilde{\mathcal{O}}}(\mathfrak{M}) \in \Pi^{\mathcal{P}}_1(D_{\mathcal{O}}(\mathfrak{M}))$.*
3. *For all $T \in \mathcal{T}$, there exists finite $D \subseteq T$ such that for all $T' \in \mathcal{T}$, if $D \subseteq T'$ then $T' \not\subset T$.*

There are connections between our framework and learning paradigms with uncountably many possible worlds (as examples of such paradigms, see [8, 13]). We can get results similar to the previous ones with suitable sets of assumptions.

7 Conclusion

We still haven't answered the following question: what is an inductive consequence of a theory T? We think there are good epistemological reasons to claim that a sentence φ an inductive consequence of T (in $\mathcal{P}$, assuming that T belongs to $\mathcal{T}$) just in case φ is a member of $\Pi^{\mathcal{P}}_1(T)$. So we would identify sentences obtained from T by one application of the compactness property with deductive consequences of T, and sentences obtained from T by one application of the weak compactness property as inductive consequences of T. But the weak compactness property needs justifications coming from many directions. The fact that it enables, together with the compactness property, to build hierarchies of generalized logical consequences that have natural connections with other classical hierarchies, provides another justification. Higher levels of the hierarchies correspond to more challenging ways of discovering the truth. With this respect, level Σ_2 deserves special attention, thanks to the connection between this level and the notion of learnability in the limit. Another part of our project, an Inductive Prolog, targets precisely these formulas. Given a possible theory T and a sentence φ of the form $\exists\overline{x}\forall\overline{y}\psi(\overline{x},\overline{y})$, where ψ is a quantifier-free formula, such that $T \models^{\mathcal{O}}_{\mathcal{W}} \varphi$, the system will try to compute a sequence of terms $\overline{t}$ such that $T \models^{\mathcal{O}}_{\mathcal{W}} \forall\overline{y}\psi(\overline{t},\overline{y})$, performing top-down searches both at level Σ_2 and at level Σ_1.

References

1. D. Angluin. *Inductive Inference of Formal Languages from Positive Data.* Information and Control, 45, 1980, p. 117–135.
2. J. Barwise. *Admissible Sets and Structures.* Perspectives in Mathematical Logic, Springer-Verlag, 1975.
3. K. Doets. *From Logic to Logic Programming.* MIT Press, 1994.
4. W. Gasarch, M. Pleszkoch, F. Stephan, and M. Velauthapillai. *Classification using information.* Annals of Mathematics and Artificial Intelligence. Selected papers from ALT 1994 and AII 1994, 23:147-168, 1998.
5. M. L. Ginsberg ed. *Readings in Nonmonotonic reasoning.* Morgan Kaufmann, 1987.
6. L. Henkin. *Completeness in the theory of types.* Journal of Symbolic Logic 15:81-91, 1950.
7. S. Jain, D. N. Osherson. J. S. Royer and A. Sharma. *Systems that learn: An Introduction to Learning Theory, Second Edition.* The MIT Press, 1999.
8. K. Kelly. The Logic of Reliable Inquiry. Oxford University Press, 1996.
9. K. Kelly and C. Glymour. *Inductive inference and theory-laden data.* Journal of Philosophical Logic, 21(4), 1992.
10. H. J. Keisler. *Fundamentals of Model Theory.* In Handbook of Mathematical Logic, Ed. J. Barwise, Springer-Verlag, 1977.
11. J. W. Lloyd. Foundations of Logic Programming, 2nd edition. Springer-Verlag, 1987.
12. M. Makkai. *Admissible Sets and Infinitary Logic.* In *Handbook of Mathematical Logic*, Ed. J. Barwise, Springer-Verlag, 1977.
13. E. Martin and D. N. Osherson. *Elements of Scientific Inquiry.* The MIT Press, 1998.
14. S. H. Nienhuys-Cheng and R. de Wolf. *Foundations of Inductive Logic Programming.* Lecture Notes in Artificial Intelligence, Vol. 1228, Springer-Verlag, 1997.
15. K. Popper. *The Logic of Scientific Discovery.* Hutchinson, London, 1959.
16. R. Reiter. *On closed world data bases.* In *Logic and Data Bases*, Ed. A. Gallaire and J. Minker. Plenum, New York, 1978.
17. J. Shoenfield. *Mathematical Logic.* Addison-Wesley, 1967.
18. S. M. Srivastava. *A Course on Borel Sets.* Graduate Texts in Mathematics 180. Springer-Verlag, 1998.