

8-2014

Problems in the Theory of Convergence Spaces

Daniel R. Patten
Syracuse University

Follow this and additional works at: <http://surface.syr.edu/etd>

 Part of the [Computer Sciences Commons](#), and the [Mathematics Commons](#)

Recommended Citation

Patten, Daniel R., "Problems in the Theory of Convergence Spaces" (2014). *Dissertations - ALL*. Paper 152.

This Dissertation is brought to you for free and open access by the SURFACE at SURFACE. It has been accepted for inclusion in Dissertations - ALL by an authorized administrator of SURFACE. For more information, please contact surface@syr.edu.

Abstract. We investigate several problems in the theory of convergence spaces: generalization of Kolmogorov separation from topological spaces to convergence spaces, representation of reflexive digraphs as convergence spaces, construction of differential calculi on convergence spaces, mereology on convergence spaces, and construction of a universal homogeneous pretopological space. First, we generalize Kolmogorov separation from topological spaces to convergence spaces; we then study properties of Kolmogorov spaces. Second, we develop a theory of reflexive digraphs as convergence spaces, which we then specialize to Cayley graphs. Third, we conservatively extend the concept of differential from the spaces of classical analysis to arbitrary convergence spaces; we then use this extension to obtain differential calculi for finite convergence spaces, finite Kolmogorov spaces, finite groups, Boolean hypercubes, labeled graphs, the Cantor tree, and real and binary sequences. Fourth, we show that a standard axiomatization of mereology is equivalent to the condition that a topological space is discrete, and consequently, any model of general extensional mereology is indistinguishable from a model of set theory; we then generalize these results to the cartesian closed category of convergence spaces. Finally, we show that every convergence space can be embedded into a homogeneous convergence space; we then use this result to construct a universal homogeneous pretopological space.

Problems in the Theory of Convergence Spaces

by

Daniel R. Patten

B.A., Utica College of Syracuse University, 2001

M.S., Syracuse University, 2004

M.S., Syracuse University, 2013

Dissertation

Submitted in partial fulfillment of the requirements for the degree of Doctor of Philosophy
in Computer and Information Science and Engineering in the Graduate School of Syracuse
University.

August 2014

©Daniel R. Patten 2014

All Rights Reserved

To Our Lady of Fatima.

Contents

1	Introduction	1
1.1	Motivation	1
1.2	Methodology	2
1.3	Overview and Contributions	3
2	Theory of Convergence Spaces	7
2.1	Filters	7
2.2	Convergence Spaces	16
2.3	Continuity and Limits	18
2.4	Open and Closed Sets	22
2.5	Common Types of Convergence Structures	26
2.5.1	Initial Convergence Structures	27
2.5.2	Topological Convergence Structures	30
2.5.3	Pretopological Convergence Structures	33
2.5.4	Pseudotopological Convergence Structures	36
2.5.5	Discrete and Indiscrete Convergence Structures	39
2.6	Compactness and Separation	43
2.6.1	Compact Spaces	43
2.6.2	Kolmogorov Spaces	45
2.6.3	Fréchet Spaces	52

2.6.4	Hausdorff Spaces	54
3	Reflexive Digraphs as Convergence Spaces	59
3.1	Reflexive Digraphs	60
3.2	Cayley Graphs	68
4	Differential Calculus on Convergence Spaces	71
4.1	Generalized Differentials	71
4.2	Differential Calculus on Finite Convergence Spaces	76
4.2.1	Differential Calculus on Finite Kolmogorov Spaces	77
4.2.2	Differential Calculus on Finite Groups	81
4.2.3	Differential Calculus on Boolean Hypercubes	83
4.3	Differential Calculus on Labeled Graphs	86
4.4	Differential Calculus on the Cantor Tree	89
4.5	Differential Calculus on Sequences	91
4.5.1	Differentials of Real Sequences	93
4.5.2	Differentials of Binary Sequences	95
5	Mereology on Topological and Convergence Spaces	101
6	A Universal Homogeneous Pretopological Space	111
6.1	Uspenskiĭ Spaces	112
6.2	Construction and Analysis	115
7	Conclusion and Future Work	119
	Bibliography	122
	Biographical Data	127

List of Figures

2.1	The Principal Filter $[\{a, c\}]$ on the Set $\{a, b, c, d\}$	8
2.2	The Point Filter $[c]$ on the Set $\{a, b, c, d\}$	9
3.1	The Pentacle.	61
3.2	A Partial Ordering on \mathbb{N}	67
3.3	Another Partial Ordering on \mathbb{N}	68
3.4	A Cayley Graph of the Cyclic Group C_3	69
4.1	A Cayley graph of the Symmetric Group S_3	82
4.2	A Cayley Graph of \mathbb{Z}	83
4.3	The 3-Boolean Hypercube.	84
4.4	The Cantor Tree.	90
6.1	A Universal Pretopological Space.	116

Chapter 1

Introduction

We investigate several problems in the theory of convergence spaces: generalization of Kolmogorov separation from topological spaces to convergence spaces, representation of reflexive digraphs as convergence spaces, construction of differential calculi on convergence spaces, mereology on convergence spaces, and construction of a universal homogeneous pretopological space.

1.1 Motivation

Despite its profound, nearly ubiquitous, impact on modern science, elementary differential calculus is not without severe limitations: its applications are ostensibly restricted to differentiable manifolds. Every hybrid dynamical system, however, is precisely a system of *differential* equations, each of which is a relation between functions on continuous, discrete, or hybrid structures and their derivatives. Since hybrid dynamical systems model many significant phenomena¹ of current interest, it is now opportune to obtain a seamless extension of elementary differential calculus from the continuous spaces of classical analysis to hybrid structures.

¹These phenomena include verification and validation of software, mechanical systems, and biological systems. See Aihara and Suzuki [1] and van der Schaft and Schumacher [42].

Mereology, originally posited as an alternative to set theory, is the theoretical foundation of formal ontology, an increasingly important tool in artificial intelligence, computational linguistics, and database theory.² Mereotopology, in turn, restricts mereology to topological spaces in place of set theoretical universes. If mereotopology is of any consequence to computer science, then the topological consequences of its axioms must be deduced and understood.

1.2 Methodology

We now justify our use of the category of convergence spaces **CONV**, rather than the well-known category of topological spaces **TOP** or Scott's category of equilogical³ spaces **EQU**, to solve the aforementioned problems.

TOP is not a Cartesian closed category, that is, there is no canonical topology for the set of continuous functions between two topological spaces.⁴ On the other hand, **CONV** is Cartesian closed and includes **TOP** as a full subcategory. Moreover, all reflexive digraphs are convergence spaces whereas only the transitive reflexive digraphs are topological spaces.⁵ Thus, not only does the theory of convergence spaces subsume the theory of topological spaces, but it also provides a finer structural framework.

EQU is Cartesian closed, but it includes only the category of Kolmogorov topological spaces as a full subcategory, and thereby precludes investigation of many spaces of interest in the present work. Furthermore, the topology of an equilogical space is generally unrelated to the equivalence relation on its carrier; in contrast, the convergence structure of a topological space is completely determined by the topology on its carrier.

If these categorical considerations were not sufficient, then we would still have recourse to one salient fact: **CONV** unifies the ostensibly disparate concepts of *homomorphism* and

²See Guarino [19].

³See Bauer, et al., [2].

⁴See Mac Lane [26] on category theory and, in particular, Cartesian closed categories.

⁵See Definition 3.1 and Proposition 3.11.

continuity. In the sequel, we will observe that preservation of filter convergence exactly characterizes both of these notions.⁶

1.3 Overview and Contributions

Since the theory of convergence spaces is widely unknown, even to specialists in related fields, we provide a thorough overview of definitions and results essential to our work. In Chapter 2 we discuss the most fundamental of these concepts: filters, convergence structures, continuity, open sets, closed sets, initial convergence structures, compactness, and separation. Most of the results in this chapter are not original; there are, however, several important exceptions.⁷

In Section 2.5, we present common types of convergence structures. Although many properties of initial, topological, pretopological, pseudotopological, discrete, and indiscrete convergence structures are well-known, others are conspicuously absent from the literature. Our contributions here include the following results:

1. Every finite convergence space is a pretopological space,⁸
2. Pseudotopological convergence structures are completely characterized by the convergence of ultrafilters,⁹
3. Functions into pseudotopological spaces are continuous if they preserve convergence of ultrafilters,¹⁰ and
4. Convergence of point filters in a pseudotopological continuous convergence structure depends only on the convergence of ultrafilters in the domain and codomain spaces.¹¹

⁶See Propositions 3.3 and 3.4. It is slightly less noteworthy, albeit not below mentioning, that the apparatus provided by **CONV** simplifies many topological proofs (most famously of Tychonoff's theorem).

⁷All proofs throughout this entire work are original. All results, excluding Proposition 3.3, in Chapters 3 through 6 are original.

⁸See Theorem 2.73.

⁹See Theorem 2.75.

¹⁰See Theorem 2.76.

¹¹See Theorem 2.81.

In Section 2.6, we survey compactness and separation. The results of Section 2.6.2, in which we generalize Kolmogorov separation from topological spaces to convergence spaces,¹² are entirely original. Among the most significant of these results is the relationship between spaces of automorphisms and Kolmogorov spaces:

1. If a finite convergence space is Kolmogorov, then its space of automorphisms is discrete,¹³ and
2. Any topological space that has a discrete space of automorphisms must be Kolmogorov.¹⁴

In Chapter 3, we develop a theory of reflexive digraphs as convergence spaces. Although Blair, et al., [8] and Mynard [31] have formulated definitions similar to Definition 3.1, none of the results in this chapter have been published, to the best of our knowledge and with the exception of Proposition 3.3, elsewhere. We then specialize these general results to Cayley graphs:¹⁵ in particular, we show that every Cayley graph, except C_1 and C_2 , is Kolmogorov.¹⁶

We regard Chapter 4 as the centerpiece of our work. Here, we restrict the definition of *differential* given in [8], verify that this restriction conservatively extends the concept of *differential* from classical analysis to arbitrary convergence spaces, and establish a “chain rule” for differentials.¹⁷ We then use this definition to construct theories of differential calculus on several general and particular convergence spaces. In Section 4.2, we obtain an equivalent condition for differentiability on finite convergence spaces; we then specialize this condition to construct differential calculi on finite Kolmogorov spaces, finite groups, and Boolean hypercubes. In Section 4.3, we develop principles for constructing differential calculi on labeled graphs. In Section 4.4, we obtain a differential calculus on the Cantor tree. Finally, in Section 4.5, we construct differential calculi for real and binary sequences.

¹²See Definitions 2.92 and 2.95 and Theorems 2.94 and 2.100.

¹³See Theorem 2.103.

¹⁴See Theorem 2.104.

¹⁵See Definition 3.20.

¹⁶See Theorems 3.23 and 3.24.

¹⁷See Definition 4.1 and Theorems 4.6 and 4.7.

In Chapter 5, much of which has appeared in Patten [33], we show that a standard axiomatization of mereology is equivalent to the condition that a topological space is discrete,¹⁸ and consequently, any model of general extensional mereology is indistinguishable from a model of set theory.¹⁹ We generalize these results to the Cartesian closed category of convergence spaces.²⁰

Finally, in Chapter 6 we show that every convergence space can be embedded into a homogeneous convergence space;²¹ we then use this result to construct a universal homogeneous pretopological space.²²

With a view toward pedagogy, we have evolved a self-contained work, which proves everything, excluding elementary facts from Zermelo-Fraenkel-Skolem²³ set theory and point-set topology, required to establish our main results.²⁴ Many of the papers, and even textbooks, on the theory of convergence spaces omit proofs of elementary results; although we expect such omissions, conducive to concision, in papers on widely-known topics, we object to this practice in the literature of emergent specialties because it obscures methodology, and so disadvantages the uninitiated. To amend these gaps, we provide an overview of convergence spaces: beginning by proving basic facts about filters and ultrafilters, we then define *convergence space* and adapt concepts native to topology—such as continuity, limits, open and closed sets, compactness, and separation—to convergence spaces.²⁵ Furthermore, we include many examples in the hope of making some of the abstract concepts discussed herein accessible to an audience wider than present. Indeed, despite its advantages over topology,

¹⁸See Theorem 5.7.

¹⁹See Theorem 5.18.

²⁰See Theorem 5.19.

²¹See Theorem 6.7.

²²See Theorem 6.10.

²³Although it is well-known that the axiom of separation is necessary for what is conventionally called Zermelo-Fraenkel set theory, it is less-known that Skolem's work on definite statements contributed to its formulation. See Suppes [40] for further discussion.

²⁴With the exception of the proof of Proposition 4.2, for which we found an elementary yet long and tedious proof, but which we have excluded in favor of a less prolix argument that depends on several results proved outside of this present work.

²⁵While we seek thoroughness in this overview of convergence spaces, it is infeasible to be complete. Thus, for wider coverage of convergence spaces we refer the reader to Beattie and Butzmann [4], which has been an invaluable resource in developing this present work.

the theory of convergence spaces receives little notice from the research community. It is our hope that this work heightens, however modestly, the profile of convergence spaces among researchers in theoretical computer science and mathematics.²⁶

²⁶It appears that some theoretical computer scientists, for example Heckmann [21], recognize the utility of convergence spaces.

Chapter 2

Theory of Convergence Spaces

2.1 Filters

In 1937, Cartan [11] introduced filters to study the convergence of sequences without recourse to countability.¹ Filters are fundamental to the study of convergence spaces. In the sequel, we use filters to obtain many of our main results, including the extension of differential from Euclidean spaces to arbitrary convergence spaces and, in particular, to discrete structures such as Cayley graphs.

Informally, a filter is a nonempty collection of nonempty sets, closed under reverse inclusion and finite intersection.

Definition 2.1. Let X be a set. A nonempty collection \mathcal{F} of nonempty subsets of X is a *filter* if and only if

1. $A \in \mathcal{F}$ and $A \subseteq B$ implies $B \in \mathcal{F}$, and
2. $A \in \mathcal{F}$ and $B \in \mathcal{F}$ implies $A \cap B \in \mathcal{F}$.

If \mathcal{B} is a nonempty collection of nonempty subsets of X that is closed under finite intersection, then the set $\{C : B \supseteq C \subseteq X \text{ for some } B \in \mathcal{B}\}$ is the filter generated by \mathcal{B} and is denoted

¹For further discussion on filters, see Bourbaki [9]; for a historical treatment of their development, see Mashaal [27].




Figure 2.1: The Principal Filter $[\{a, c\}]$ on the Set $\{a, b, c, d\}$.

by $[\mathcal{B}]$. In particular, if A is a nonempty subset of X , then the set $\{B : A \subseteq B \subseteq X\}$ is a filter, denoted by $[A]$, and called the *principal filter generated by A* ; the *point filter at x* is the principal filter $[\{x\}]$, which we abbreviate as $[x]$. We denote the set of all filters on X by $\Phi(X)$. When $\mathcal{F} \supset \mathcal{G}$, we say that \mathcal{F} is *finer* than \mathcal{G} and \mathcal{G} is *coarser* than \mathcal{F} . An *ultrafilter* is a filter not coarser than any other filter; an ultrafilter that is not a point filter is a *free filter*.

Some authors² allow filters to contain the empty set; consequently, they usually distinguish between the *improper* filter, which is the unique filter that contains the empty set, and *proper* filters, which do not contain the empty set.

Example 2.2. If $X = \{a, b, c, d\}$, then $[\{a, c\}] = \{\{a, c\}, \{a, b, c\}, \{a, c, d\}, X\}$ is a principal filter on X . The red quadrilateral in the Hasse Diagram of Figure 2.1 indicates $[\{a, c\}]$ on X . On the other hand $\{\{a, c\}, \{a, b, c\}, X\}$ is not a filter on X because it is not closed under upward inclusion; likewise $\{\{a, b\}, \{a, c\}, \{a, b, c\}, \{a, c, d\}, X\}$ is not a filter on X because it

²For example, in [8] and [21].




Figure 2.2: The Point Filter $[c]$ on the Set $\{a, b, c, d\}$.

is not closed under finite intersections. The blue prism in the Hasse Diagram of Figure 2.2 indicates the point filter at c on X .

Example 2.3. Let X be an infinite set. The *Fréchet filter* on X is the collection of all subsets of X that have finite complement.³ Every free filter includes the Fréchet filter.⁴

It is not obvious that free filters exist. To establish the existence of free filters, we first show that every filter is included in some ultrafilter.

Proposition 2.4. *Every filter is included in some ultrafilter.*

Proof. Let \mathcal{F} be a filter on a set X , let $(\mathcal{F}_i)_{i \in I}$ be the collection of all filters on X that include \mathcal{F} , and let \mathbf{C} be a chain in $(\mathcal{F}_i)_{i \in I}$. It is clear that inclusion partially orders $(\mathcal{F}_i)_{i \in I}$.

First, we show that $\bigcup_{\mathcal{C} \in \mathbf{C}} \mathcal{C}$ is a filter on X . Since X belongs to every filter on X , it follows that $X \in \bigcup_{\mathcal{C} \in \mathbf{C}} \mathcal{C}$; similarly $\emptyset \notin \bigcup_{\mathcal{C} \in \mathbf{C}} \mathcal{C}$, otherwise $\emptyset \in \mathcal{C}$ for some $\mathcal{C} \in \mathbf{C}$. If $C \in \bigcup_{\mathcal{C} \in \mathbf{C}} \mathcal{C}$ and $C \subseteq C'$, then $C \in \mathcal{C}$ for some $\mathcal{C} \in \mathbf{C}$, and so $C' \in \mathcal{C}$, which implies that $C' \in \bigcup_{\mathcal{C} \in \mathbf{C}} \mathcal{C}$. If

³Hence, another name for the Fréchet filter on X is the *cofinite filter* on X .

⁴See Proposition 2.21.

C and C' both belong to $\bigcup_{\mathcal{C} \in \mathbf{C}} \mathcal{C}$, then there exist \mathcal{C} and \mathcal{D} in $\bigcup_{\mathcal{C} \in \mathbf{C}} \mathcal{C}$ such that $C \in \mathcal{C}$ and $C' \in \mathcal{D}$. Since \mathbf{C} is a chain, either $\mathcal{C} \subseteq \mathcal{D}$ or $\mathcal{D} \subseteq \mathcal{C}$. In the former case $C \cap C' \in \mathcal{D}$; in the latter case $C \cap C' \in \mathcal{C}$. Thus $C \cap C' \in \bigcup_{\mathcal{C} \in \mathbf{C}} \mathcal{C}$.

Since $\mathcal{D} \subseteq \bigcup_{\mathcal{C} \in \mathbf{C}} \mathcal{C}$ for each $\mathcal{D} \in \mathbf{C}$, it follows that $\bigcup_{\mathcal{C} \in \mathbf{C}} \mathcal{C}$ is an upper bound on \mathbf{C} . Thus, by Zorn's lemma⁵ $(\mathcal{F}_i)_{i \in I}$ has a maximal element \mathcal{U} . Therefore \mathcal{U} is an ultrafilter that contains \mathcal{F} . □

Proposition 2.5. *For every infinite set X , there exists a free filter \mathcal{U} on X .*

Proof. By Proposition 2.4, there exists an ultrafilter \mathcal{U} that includes the Fréchet filter \mathcal{F} on X . Since $X - \{x\} \in \mathcal{F}$ for each $x \in X$, it follows that \mathcal{U} is not a point filter; thus \mathcal{U} is a free filter. □

The second condition of Definition 2.1 implies a frequently useful formulation of the subset relation between filters.

Proposition 2.6. *If \mathcal{F} and \mathcal{G} are filters on a set X , then $\mathcal{F} \subseteq \mathcal{G}$ if and only if for every $F \in \mathcal{F}$ there exists $G \in \mathcal{G}$ such that $G \subseteq F$.*

Proof. [Necessity]. If $F \in \mathcal{F}$, then by hypothesis there exists $G \in \mathcal{G}$ such that $G \subseteq F$, and so $F \in \mathcal{G}$. Therefore $\mathcal{F} \subseteq \mathcal{G}$.

[Sufficiency]. If $F \in \mathcal{F}$, then by hypothesis $F \in \mathcal{G}$. Since \mathcal{G} is nonempty, there exists $G \in \mathcal{G}$. Thus $F \cap G$, a subset of F , belongs to \mathcal{G} . □

Proposition 2.7. *If A and B are subsets of X , then $[A] \subseteq [B]$ if and only if $B \subseteq A$.*

Proof. [Necessity]. If $C \in [A]$, then $A \subseteq C$, and so by hypothesis $B \subseteq C$; thus $C \in [B]$.

[Sufficiency]. By hypothesis, there exists $C \in [B]$ such that $C \subseteq A$. Since $C \in [B]$, it follows that $B \subseteq C$; thus $B \subseteq A$. □

⁵Zorn's lemma, which is equivalent to the axiom of choice, states that a partially ordered set has a maximal element whenever each of its chains has an upper bound in it.

Although we use the same notation for both the image (and likewise, preimage) of a set and a filter under a function, the notation for filters and sets should prevent any reasonable confusion from arising.

Definition 2.8. Let $f : X \rightarrow Y$ be a function, let $A \subseteq X$ and $B \subseteq Y$, and let \mathcal{F} and \mathcal{G} be filters on X and Y , respectively. The *image of A under f* , denoted by $f(A)$, is the set $\{f(a) : a \in A\}$; the *image of \mathcal{F} under f* , denoted by $f(\mathcal{F})$, is the filter $[f(F) : F \in \mathcal{F}]$. The *preimage of B under f* , denoted by $f^{-1}(B)$, is the set $\{x : f(x) \in B\}$; the *preimage of \mathcal{G} under f* , denoted by $f^{-1}(\mathcal{G})$, is the filter $[f^{-1}(G) : G \in \mathcal{G}]$.

Images of filters are well-behaved under composition of functions.

Proposition 2.9. *If $f : X \rightarrow Y$ and $g : Y \rightarrow Z$ are functions, then $(f \circ g)(\mathcal{F}) = f(g(\mathcal{F}))$ for every $\mathcal{F} \in \Phi(X)$.*

Proof. If $K \in (f \circ g)(\mathcal{F})$, then $K \supseteq (f \circ g)(H) = f(g(H))$ for some $H \in \mathcal{F}$. Since $g(H) \in g(\mathcal{F})$, it follows that $f(g(H)) \in f(g(\mathcal{F}))$; because filters are closed under supersets, we conclude that $K \in f(g(\mathcal{F}))$. Thus $(f \circ g)(\mathcal{F}) \subseteq f(g(\mathcal{F}))$.

Conversely, if $K \in f(g(\mathcal{F}))$, then $K \supseteq f(H)$ for some $H \in g(\mathcal{F})$. It follows that $H \supseteq g(F)$ for some $F \in \mathcal{F}$, and so $K \supseteq f(g(F)) = (f \circ g)(F)$. Thus $K \in (f \circ g)(\mathcal{F})$, and therefore $f(g(\mathcal{F})) \subseteq (f \circ g)(\mathcal{F})$. \square

We can construct filters on Cartesian products of sets from filters on the factor sets.

Definition 2.10. Let $(X_i)_{i \in I}$ be a collection of sets and let $(\mathcal{A}_i)_{i \in I}$ be a collection of filters such that $\mathcal{A}_i \in \Phi(X_i)$ for each $i \in I$. The *product filter* of $(\mathcal{A}_i)_{i \in I}$ on $\prod_{i \in I} X_i$ is the filter $\prod_{i \in I} \mathcal{A}_i = [\prod_{i \in I} A_i : A_i \in \mathcal{A}_i]$.

In studying convergence spaces, we often are interested in whether a given filter converges to some point.⁶ Because convergence structures are closed under upward inclusion, it is useful to construct a filter finer than a filter known to converge to a given point. We provide two such constructions.

⁶See Definition 2.24.

Proposition 2.11. *Let \mathcal{F} and \mathcal{G} be filters on a set X . If every element of \mathcal{F} intersects every element of \mathcal{G} , then the set $\mathcal{H} = \{F \cap G : F \in \mathcal{F} \wedge G \in \mathcal{G}\}$ is a filter finer than both \mathcal{F} and \mathcal{G} .*

Proof. It is clear that \mathcal{H} contains X but not \emptyset . If H and H' both belong to \mathcal{H} , then there exist F and F' in \mathcal{F} and G and G' in \mathcal{G} such that $H = F \cap G$ and $H' = F' \cap G'$. Since $F \cap F' \in \mathcal{F}$ and $G \cap G' \in \mathcal{G}$, it follows that $H \cap H' = (F \cap G) \cap (F' \cap G') = (F \cap F') \cap (G \cap G') \in \mathcal{H}$. If $H \in \mathcal{H}$ and $H \subseteq K$, then $F \cap G \subseteq K$ for some $F \in \mathcal{F}$ and $G \in \mathcal{G}$. Since $K \cup F \in \mathcal{F}$ and $K \cup G \in \mathcal{G}$, it follows that $K = (K \cup F) \cap (K \cup G) \in \mathcal{H}$. Thus \mathcal{H} is a filter. If $F \in \mathcal{F}$, then $F = F \cap X \in \mathcal{H}$, and so $\mathcal{F} \subseteq \mathcal{H}$; likewise $\mathcal{G} \subseteq \mathcal{H}$. \square

Proposition 2.12. *Let \mathcal{F} be a filter on a set X and let $A \subseteq X$. If $A \notin \mathcal{F}$, then the set $\mathcal{G} = \{B : B \supseteq F - A \wedge F \in \mathcal{F}\}$ is a filter finer than \mathcal{F} .*

Proof. If $F - A = \emptyset$, then $F \subseteq A$; since \mathcal{F} is a filter, it follows that $A \in \mathcal{F}$, in contradiction to the hypothesis. Thus $\emptyset \notin \mathcal{G}$. Likewise, it is clear that $X \in \mathcal{G}$. If $G_1 \in \mathcal{G}$ and $G_1 \subseteq G_2$, then $F - A \subseteq G_1 \subseteq G_2$ for some $F \in \mathcal{F}$; thus $G_2 \in \mathcal{G}$. If G_1 and G_2 belong to \mathcal{G} , then $G_1 \supseteq F_1 - A$ and $G_2 \supseteq F_2 - A$ for some $F_1, F_2 \in \mathcal{F}$; thus $G_1 \cap G_2 \supseteq (F_1 - A) \cap (F_2 - A) = (F_1 \cap F_2) - A$ belongs to \mathcal{G} . Therefore \mathcal{G} is a filter. Finally, if $F \in \mathcal{F}$, then $F - A \subseteq F$, and so $F \in \mathcal{G}$; therefore \mathcal{G} is finer than \mathcal{F} . \square

We frequently will use the constructions in Propositions 2.11 and 2.12 throughout this work. One application of Proposition is a generalization of Proposition 2.5; another is an equivalent formulation of ultrafilters.

Proposition 2.13. *If \mathcal{S} is a nonempty collection of subsets of X such that the intersection of every finite subcollection of \mathcal{S} is infinite, then there exists a free filter on X that includes \mathcal{S} .*

Proof. Let \mathcal{F} denote the Fréchet filter on X and let

$$\mathcal{T} = \left\{ \bigcap \mathcal{S}' : \mathcal{S}' \text{ is a nonempty finite subcollection of } \mathcal{S} \right\}.$$

By Proposition 2.11, the set $\mathcal{G} = \{F \cap T : F \in \mathcal{F} \wedge T \in \mathcal{T}\}$ is a filter finer than both \mathcal{F} and \mathcal{T} . By Proposition 2.4, there exists an ultrafilter \mathcal{U} finer than \mathcal{G} . Since $X - \{x\} \in \mathcal{F} \subseteq \mathcal{G}$ for each $x \in X$, it follows that \mathcal{U} is not a point filter; thus \mathcal{U} is a free filter. If $S \in \mathcal{S}$, then $S \in \mathcal{T}$; thus $S = X \cap S \in \mathcal{G} \subseteq \mathcal{T}$. Therefore $\mathcal{S} \subseteq \mathcal{T}$, as desired. \square

Proposition 2.14. *A filter \mathcal{U} on X is an ultrafilter if and only if either $A \in \mathcal{U}$ or $X - A \in \mathcal{U}$ for every $A \subseteq X$.*

Proof. [Necessity]. Let \mathcal{U} be a filter on X such that $A \in \mathcal{U}$ or $X - A \in \mathcal{U}$ for every $A \subseteq X$. If \mathcal{U} is properly contained in a filter \mathcal{V} , then there exists $V \in \mathcal{V}$ such that $V \notin \mathcal{U}$. By hypothesis $X - V \in \mathcal{U}$, and so $X - V \in \mathcal{V}$; but $(X - V) \cap V = \emptyset$. Therefore, we conclude that \mathcal{U} is not properly contained in any filter.

[Sufficiency]. Let \mathcal{U} be an ultrafilter on X . Suppose that $A \notin \mathcal{U}$ for some subset A of X . By Proposition 2.12 it follows that $\mathcal{V} = \{V : V \supseteq U - A \wedge U \in \mathcal{U}\}$ is a filter finer than \mathcal{U} . Since \mathcal{U} is an ultrafilter, it follows that $\mathcal{U} = \mathcal{V}$. Because $X \in \mathcal{U}$, we conclude that $X - A \in \mathcal{U}$. \square

In general, if an ultrafilter contains a finite union of sets, it must contain at least one of those sets; if the sets are pairwise disjoint, then the ultrafilter contains exactly one of the sets. Similarly, any ultrafilter finer than a finite intersection of ultrafilters is identical to one of them.

Proposition 2.15. *Let \mathcal{U} be an ultrafilter on a set X and let $(S_i)_{i \in I}$ be a finite collection of subsets of X . If $\bigcup_{i \in I} S_i \in \mathcal{U}$, then $S_i \in \mathcal{U}$ for some $i \in I$.*

Proof. We proceed by induction on $|I|$. Let $|I| = 2$. If neither S_0 nor S_1 belong to \mathcal{U} , then $X - S_0$ and $X - S_1$ both belong to \mathcal{U} , and so $X - (S_0 \cup S_1) = (X - S_0) \cap (X - S_1) \in \mathcal{U}$, in contradiction to the hypothesis that $S_0 \cup S_1 \in \mathcal{U}$. Now suppose that $|I| = n + 1$ for some integer $n \geq 2$. By hypothesis $(S_1 \cup \dots \cup S_n) \cup S_{n+1} \in \mathcal{U}$. If $S_{n+1} \notin \mathcal{U}$, then $S_1 \cup \dots \cup S_n \in \mathcal{U}$, and so by the inductive hypothesis $S_k \in \mathcal{U}$ for some $1 \leq k \leq n$. \square

Proposition 2.16. *Let $(\mathcal{U}_i)_{i \in I}$ be a finite collection of ultrafilters on a set X . If \mathcal{V} is an ultrafilter finer than $\bigcap_{i \in I} \mathcal{U}_i$, then $\mathcal{V} = \mathcal{U}_i$ for some $i \in I$.*

Proof. If $\mathcal{V} \neq \mathcal{U}_i$ for each $i \in I$, then for every $i \in I$, there exists $U_i \in \mathcal{U}_i$ such that $X - U_i \in \mathcal{V}$; thus $X - \bigcup_{i \in I} U_i = \bigcap_{i \in I} (X - U_i) \in \mathcal{V}$, which implies that $\bigcup_{i \in I} U_i \notin \mathcal{V}$; but $\bigcup_{i \in I} U_i \in \mathcal{U}_i$ for each $i \in I$, it follows that $\bigcup_{i \in I} U_i \in \bigcup_{i \in I} \mathcal{U}_i \subseteq \mathcal{V}$, in contradiction to the previous result that $\bigcup_{i \in I} U_i \notin \mathcal{V}$. Therefore, we conclude that $\mathcal{V} = \mathcal{U}_i$ for some $i \in I$. \square

In Example 2.2, the only ultrafilters on X are the point filters and all other filters are principal filters. This is not a peculiarity of X ; rather it is true of all finite sets.

Proposition 2.17. *Let \mathcal{U} be an ultrafilter. If \mathcal{U} contains a finite set, then \mathcal{U} is a point filter.*

Proof. Let F be the smallest finite set in \mathcal{U} . If $G \subset F$, then $G \notin \mathcal{U}$, which implies that $\mathcal{V} = \{V : V \supseteq U - G \wedge U \in \mathcal{U}\}$ is a filter finer than \mathcal{U} . Since \mathcal{U} is an ultrafilter, it follows that $\mathcal{U} = \mathcal{V}$. In particular $F - G \in \mathcal{U}$. But $|F - G| = |F| - |G| < |F|$, in contradiction to the hypothesis that F is the smallest finite set in \mathcal{U} . Thus F contains no proper subsets, and so F is either empty or a singleton set. Since F cannot be empty, we conclude that F is a singleton set, and therefore \mathcal{U} is a point filter. \square

Proposition 2.18. *Every ultrafilter on a finite set is a point filter.*

Proof. Let \mathcal{U} be an ultrafilter on a finite set X . Since $X \in \mathcal{U}$, Proposition 2.17 implies that \mathcal{U} is a point filter. \square

Proposition 2.19. *Every filter on a finite set is a principal filter.*

Proof. Let \mathcal{F} be a filter on a finite set X . Let F be the smallest set in \mathcal{F} . If $G \in \mathcal{F}$ but $F \not\subseteq G$, then $|F \cap G| < |F|$; since $F \cap G \in \mathcal{F}$, this contradicts the hypothesis that F is the smallest set in \mathcal{F} . Thus $F \subseteq G$, and therefore $\mathcal{F} = [F]$. \square

Each filter is the intersection of those ultrafilters including it. Consequently, the intersection of all free filters on an infinite set is the Fréchet filter.

Proposition 2.20. *Every filter is the intersection of those ultrafilters including it.*

Proof. Let \mathcal{F} be a filter on a set X and let $(\mathcal{U}_i)_{i \in I}$ be the collection of all ultrafilters including \mathcal{F} . It suffices to show that $\bigcap_{i \in I} \mathcal{U}_i \subseteq \mathcal{F}$. To the contrary, if $\bigcap_{i \in I} \mathcal{U}_i \not\subseteq \mathcal{F}$, then there exists $U \in \bigcap_{i \in I} \mathcal{U}_i$ that does not include any $F \in \mathcal{F}$. Thus, the set $\mathcal{G} = \{G \mid G \supseteq F - U \wedge F \in \mathcal{F}\}$ is a filter finer than \mathcal{F} . If \mathcal{V} is an ultrafilter finer than \mathcal{G} , then \mathcal{V} is finer than \mathcal{F} , and so $\mathcal{V} = \mathcal{U}_i$ for some $i \in I$; but then $F - U$ and U belong to \mathcal{U}_i , which is absurd. Therefore $\bigcap_{i \in I} \mathcal{U}_i \subseteq \mathcal{F}$. \square

Proposition 2.21. *If X is an infinite set, then the intersection of all free filters on X is the Fréchet filter.*

Proof. Let $(\mathcal{U}_i)_{i \in I}$ denote the collection of all free filters on X and let \mathcal{F} denote the Fréchet filter.

First, we show that $\mathcal{F} \subseteq \bigcap_{i \in I} \mathcal{U}_i$. To the contrary, if $\mathcal{F} \not\subseteq \mathcal{U}_i$ for some $i \in I$, then there exists $F \in \mathcal{F}$ not contained in \mathcal{U}_i , and so $X - F \in \mathcal{U}_i$; but $X - F$ is finite, which implies that \mathcal{U}_i is a point filter, in contradiction to the assumption that it is free. Thus $\mathcal{F} \subseteq \bigcap_{i \in I} \mathcal{U}_i$.

Next, we show that $\bigcap_{i \in I} \mathcal{U}_i \subseteq \mathcal{F}$. To the contrary, if $\bigcap_{i \in I} \mathcal{U}_i \not\subseteq \mathcal{F}$, then there exists U that belongs to each \mathcal{U}_i but does not belong to \mathcal{F} . Thus, the set $\mathcal{G} = \{G \mid G \supseteq F - U \wedge F \in \mathcal{F}\}$ is a filter finer than \mathcal{F} . If \mathcal{V} is an ultrafilter finer than \mathcal{G} , then it is also finer than \mathcal{F} , from which it follows that \mathcal{V} is a free filter, and so $\mathcal{V} = \mathcal{U}_i$ for some $i \in I$, which implies that $U \in \mathcal{V}$, which is absurd since $X - U \in \mathcal{V}$. Therefore $\bigcap_{i \in I} \mathcal{U}_i \subseteq \mathcal{F}$. \square

We conclude our discussion on filters with two propositions concerning the image of an ultrafilter. The first establishes that the image of an ultrafilter must be an ultrafilter; the second, that every ultrafilter including the image of a given filter is itself the image of some ultrafilter finer than the given filter.

Proposition 2.22. *Any image of an ultrafilter is an ultrafilter.*

Proof. Let $f : X \rightarrow Y$ be a function. Let \mathcal{U} be an ultrafilter on X . Suppose that $V \notin f(\mathcal{U})$ for some $V \subseteq Y$. Then $f^{-1}(V) \notin \mathcal{U}$; otherwise $f(f^{-1}(V)) \in f(\mathcal{U})$, which implies that $V \in f(\mathcal{U})$

since $f(f^{-1}(V)) \subseteq V$. By hypothesis $X - f^{-1}(V) \in \mathcal{U}$. Thus $f(X - f^{-1}(V)) \in f(\mathcal{U})$; since $f(X - f^{-1}(V)) \subseteq Y - V$, it follows that $Y - V \in f(\mathcal{U})$. Therefore $f(\mathcal{U})$ is an ultrafilter on Y . \square

Proposition 2.23. *Let $f : X \rightarrow Y$ be a function and let \mathcal{F} be a filter on X . For every ultrafilter \mathcal{V} finer than $f(\mathcal{F})$, there exists an ultrafilter \mathcal{U} finer than \mathcal{F} such that $f(\mathcal{U}) = \mathcal{V}$.*

Proof. Let \mathcal{V} be an ultrafilter finer than $f(\mathcal{F})$. If $V \in \mathcal{V} - f(\mathcal{F})$, then $V \notin f(\mathcal{F})$ and $Y - V \notin f(\mathcal{F})$. If $F \cap f^{-1}(V) = \emptyset$ for some $F \in \mathcal{F}$, then $F - f^{-1}(Y - V) = \emptyset$, and so $F \subseteq f^{-1}(Y - V)$, which implies that $f(F) \subseteq Y - V$; thus $Y - V \in f(\mathcal{F})$, in contradiction to the previous result. Hence $F \cap f^{-1}(V) \neq \emptyset$ for every $F \in \mathcal{F}$, and so we can define the filter $\mathcal{G} = [F \cap W : F \in \mathcal{F} \wedge W \in f^{-1}(\mathcal{V})]$. Let \mathcal{U} be an ultrafilter finer than \mathcal{G} . Since \mathcal{G} is finer than both \mathcal{F} and $f^{-1}(\mathcal{V})$, it follows that \mathcal{U} is finer than both \mathcal{F} and $f^{-1}(\mathcal{V})$. Thus $\mathcal{V} \subseteq f(f^{-1}(\mathcal{V})) \subseteq f(\mathcal{U})$. Because \mathcal{V} is an ultrafilter, we conclude that $f(\mathcal{U}) = \mathcal{V}$. \square

2.2 Convergence Spaces

Ten years after Cartan introduced filters, Choquet [13] used them to develop the theory of convergence spaces. Unlike topological spaces, convergence spaces form a Cartesian closed category; moreover, the former is a full subcategory of the latter.

A *convergence space* is a set in which each point is associated with a collection of filters; each associated collection must be closed under both reverse inclusion and finite intersection and include the point filter.

Definition 2.24. Let X be a set. A *convergence structure* is a relation \downarrow between $\Phi(X)$ and X such that for every $x \in X$ and $\mathcal{F}, \mathcal{G} \in \Phi(X)$:

1. $[x] \downarrow x$,
2. $\mathcal{F} \downarrow x$ and $\mathcal{F} \subseteq \mathcal{G}$ implies $\mathcal{G} \downarrow x$, and

3. $\mathcal{F} \downarrow x$ and $\mathcal{G} \downarrow x$ implies $\mathcal{F} \cap \mathcal{G} \downarrow x$.

A *convergence space* is a pair (X, \downarrow) in which X is a set and \downarrow is a convergence structure. When no reasonable confusion is likely, we refer to (X, \downarrow) by X . We read $\mathcal{F} \downarrow x$ as “ \mathcal{F} converges to x ”.

There are several extant notions of convergence structures. For example, in [8] and [21] convergence structures are not required to satisfy the finite intersection property. To distinguish between the spaces associated with such structures and those associated with the convergence structures of Definition 2.24, the former are called *generalized convergence spaces*; the latter, *limit spaces*.⁷ There are some distinctions between generalized convergence spaces and limit spaces. One notable example is that finite limit spaces are pretopological, but finite generalized convergence spaces need not be pretopological.⁸

Example 2.25. Let $X = \{0, 1, 2\}$. Define a convergence structure on X by the equivalences:

$$\mathcal{F} \downarrow 0 \text{ if and only if } \{0, 1\} \in \mathcal{F}$$

$$\mathcal{F} \downarrow 1 \text{ if and only if } \{1, 2\} \in \mathcal{F}$$

$$\mathcal{F} \downarrow 2 \text{ if and only if } \{2\} \in \mathcal{F}$$

Notice that neither $[0, 2]$ nor $[X]$ converge.

Example 2.26. Consider \mathbb{R} with the standard topology. For each $x \in \mathbb{R}$, define \mathcal{U}_x to be collection of all topological neighborhoods of x . Define a convergence structure on \mathbb{R} by the equivalence $\mathcal{F} \downarrow x$ if and only if $\mathcal{F} \supseteq \mathcal{U}_x$. Unless noted otherwise, whenever we refer to \mathbb{R} we assume that it has this convergence structure.

⁷Definition 2.24 also coincides with that of Binz [7]. Nel [32] credits the convergence spaces of Definition 2.24 to Kowalsky [24] and Fischer [17]. See Preuß [34] for a taxonomy of convergence spaces.

⁸See Definition 2.64 and Theorem 2.73.

2.3 Continuity and Limits

Continuous functions preserve structure. For example, a function on the Euclidean line is continuous at a point if and only if its inverse preserves open intervals, that is, given some open interval in the range that contains the image of a particular point, there is an open interval in the domain, which the function maps into the specified open interval in the range, that contains the point. Since the open sets of topological spaces generalize the open intervals of the Euclidean line, a function between topological spaces is continuous at a point if and only if its inverse preserves open sets, that is, given some open set in the range that contains the image of a particular point, there is an open set in the domain, which the function maps into the specified open set in the range, that contains the point. Although continuous functions between convergence spaces also preserve structure, it is not by means of preimages: the image of convergent filter under a continuous function is again a convergent filter.⁹

Definition 2.27. Let X and Y be convergence spaces and let $x \in X$. A function $f : X \rightarrow Y$ is *continuous at x* if and only if $f(\mathcal{F}) \downarrow f(x)$ in Y whenever $\mathcal{F} \downarrow x$ in X for each filter \mathcal{F} on X ; it is *continuous* if and only if it is continuous at every point of X .

Definition 2.27 coincides with the usual concept of continuity when the convergence space is topological.¹⁰ As observed in [8], continuity has an equivalent formulation, which is occasionally more useful than Definition 2.27.¹¹

Proposition 2.28. *Let X and Y be convergence spaces. The function $f : X \rightarrow Y$ is continuous at $x \in X$ if and only if for every filter \mathcal{F} converging to $x \in X$, there exists a filter \mathcal{G} converging to $f(x)$ in Y such that for every element G in \mathcal{G} , there exists an element F in \mathcal{F} such that $f(F) \subseteq G$.*

⁹As we shall see in Chapter 3, the continuous functions on a reflexive digraph are exactly the graph automorphisms, and so preserve the structure of the graph.

¹⁰See Definition 2.59 and Proposition 2.61.

¹¹Or at least noteworthy in its formal similarity to the definition of *differential*. See Definition 4.1.

Proof. [Necessity]. If \mathcal{F} converges to x in X , then there exists a filter \mathcal{G} converging to $f(x)$ in Y such that for every element G in \mathcal{G} , there exists an element F in \mathcal{F} such that $f(F) \subseteq G$. If $G \in \mathcal{G}$, then there exists $F \in \mathcal{F}$ such that $f(F) \subseteq G$; since $f(F) \in f(\mathcal{F})$, it follows that $G \in f(\mathcal{F})$. Thus $\mathcal{G} \subseteq f(\mathcal{F})$, and therefore $f(\mathcal{F})$ converges to $f(x)$ in Y .

[Sufficiency]. If \mathcal{F} converges to x in X , then $f(\mathcal{F})$ converges to $f(x)$ in Y . Moreover, if G belongs to $f(\mathcal{F})$, then $f(F) \subseteq G$, as desired. \square

Of course, composition of functions preserves continuity.

Proposition 2.29. *Let X , Y , and Z be convergence spaces. If $f : X \rightarrow Y$ and $g : Y \rightarrow Z$ are continuous, then $g \circ f$ is continuous.*

Proof. Let $\mathcal{F} \downarrow x$ in X . By hypothesis $f(\mathcal{F}) \downarrow f(x)$ in Y , and so $g(f(\mathcal{F})) \downarrow g(f(x))$ in Z , or equivalently $(g \circ f)(\mathcal{F}) \downarrow (g \circ f)(x)$ in Z . Therefore, we conclude that $g \circ f$ is continuous. \square

Much of the present work involves structure preserving functions, that is, homeomorphisms.

Definition 2.30. Let X and Y be convergence spaces. A *homeomorphism* is a continuous bijection $f : X \rightarrow Y$, the inverse of which is also continuous; the spaces are *homeomorphic* if and only if there exists a homeomorphism between them. An *automorphism* is a homeomorphism on a convergence space to itself; we denote the set of all automorphisms on a convergence space X by $\text{Aut}(X)$.¹² A convergence space is *homogeneous* if and only if for each pair of its elements there exists an automorphism that maps one of the pair to the other; the space is *rigid* if and only if its only automorphism is the identity function.

Proposition 2.31. *If X and Y are convergence spaces and $f : X \rightarrow Y$ is a bijection, then f is a homeomorphism if and only if $\mathcal{F} \downarrow x$ in X if and only if $f(\mathcal{F}) \downarrow f(x)$ in Y .*

¹²We always equip $\text{Aut}(X)$ with the subspace convergence structure inherited from $\mathcal{C}(X, X)$. See Definitions 2.34 and 2.53.

Proof. [Necessity]. Since $f(\mathcal{F}) \downarrow f(x)$ in Y whenever $\mathcal{F} \downarrow x$ in X , it follows that f is continuous. Since $\mathcal{F} \downarrow x$ in X whenever $f(\mathcal{F}) \downarrow f(x)$ in Y , it follows that $f^{-1}(\mathcal{F}) \downarrow f^{-1}(x)$ in X whenever $f(f^{-1}(\mathcal{F})) = \mathcal{F}$ converges to $x = f(f^{-1}(x))$ in Y .

[Sufficiency]. Since f is continuous, it follows that $f(\mathcal{F}) \downarrow f(x)$ in Y whenever $\mathcal{F} \downarrow x$ in X . Conversely, since f^{-1} is continuous, it follows that $f^{-1}(\mathcal{F}) \downarrow f^{-1}(x)$ in X whenever $\mathcal{F} \downarrow x$ in Y ; thus $\mathcal{F} = f^{-1}(f(\mathcal{F}))$ converges to $f^{-1}(f(x)) = x$ in X whenever $f(\mathcal{F})$ converges to $f(x)$ in Y □

Example 2.32. Consider the convergence space X of Example 2.25. There are exactly six continuous functions on X : the identity function, the constant functions, the function that maps 0 and 1 to 0 and 2 to 1, and the function that maps 1 and 2 to 1 and 0 to 0. Only the identity function is an automorphism; thus X is not a homogeneous space, and in particular, it is rigid.

Example 2.33. The Euclidean line \mathbb{R} is homogeneous: if a and b belong to \mathbb{R} , then the function $\lambda x.x + b - a$ is an automorphism that maps a to b .

Definition 2.34. Let X and Y be convergence spaces and let $\mathcal{C}(X, Y)$ be the set of all continuous functions from X to Y . For each filter \mathcal{F} on $\mathcal{C}(X, Y)$ and each filter \mathcal{A} on X , define the filter

$$\mathcal{F} \cdot \mathcal{A} = [F \cdot A : F \in \mathcal{F} \text{ and } A \in \mathcal{A}],$$

in which

$$F \cdot A = \{f(a) : f \in F \text{ and } a \in A\}.$$

The *continuous convergence structure* on $\mathcal{C}(X, Y)$ is defined by the equivalence.

$$\mathcal{F} \downarrow f \text{ if and only if } \mathcal{F} \cdot \mathcal{A} \downarrow f(x) \text{ whenever } \mathcal{A} \downarrow x.$$

Equipped with the continuous convergence structure,¹³ the category of convergence spaces,

¹³The notion of continuous convergence structure is due to Binz and Keller [5].

unlike the category of topological spaces, is Cartesian closed.¹⁴ In other words, there is a canonical convergence structure for the set of continuous functions between two convergence spaces.

Sometimes¹⁵ it is inconvenient to use Definition 2.34 directly; fortunately, it has an equivalent formulation—one which uses the evaluation map and product filter.

Proposition 2.35. *Let X and Y be convergence spaces. Define $E : (\mathcal{C}(X, Y) \times X) \rightarrow Y$ by $E(f, x) = f(x)$ for each $f \in \mathcal{C}(X, Y)$ and $x \in X$. Then a filter \mathcal{F} converges to f in $\mathcal{C}(X, Y)$ if and only if $E(\mathcal{F} \times \mathcal{A})$ converges to $f(x)$ in Y whenever \mathcal{A} converges to x in X .*

Proof. Since $F \cdot A = \{f(a) : f \in F \wedge a \in A\} = \{E(f, a) : f \in F \wedge a \in A\} = E(F \times A)$ for each $F \in \mathcal{F}$ and $A \in \mathcal{A}$, it follows that $\mathcal{F} \cdot \mathcal{A} = E(\mathcal{F} \times \mathcal{A})$. The desired result follows immediately from this observation. \square

Often in classical analysis, continuity is defined in terms of limits.¹⁶ This is not done out of logical necessity; rather it is likely motivated by pedagogical considerations or, perhaps, simple neglect.

Definition 2.36. Let X and Y be convergence spaces, let $f : X \rightarrow Y$ be a partial function, and let $p \in X$. A *limit* of f at p is a point $l \in Y$ such that the function $f_{p \rightarrow l} : X \rightarrow Y$, defined by

$$f_{p \rightarrow l}(x) = \begin{cases} l, & \text{if } x = p; \\ f(x), & \text{otherwise,} \end{cases}$$

is continuous at p . We write $l = \lim_{x \rightarrow p} f(x)$ whenever l is unique and $l \in \lim_{x \rightarrow p} f(x)$ otherwise.

Definition 2.36 coincides with the standard definition of *limit* for functions on \mathbb{R} . We postpone, however, the proof of this assertion until Section 2.5, in which we discuss pretopological spaces.

¹⁴The category of pseudotopological spaces, however, is Cartesian closed. See Definition 2.74 and Proposition 2.80.

¹⁵For example, in Proposition 2.80.

¹⁶For example, in Rudin [38].

Example 2.37. Consider the convergence space X of Example 2.25. Let $f : X \rightarrow X$ be the function $(1, 2)$. Although f is not continuous at 0 and 1, it is continuous at 2. Moreover, we calculate that $\lim_{x \rightarrow 0} f(x) = \{1, 2\}$, $\lim_{x \rightarrow 1} f(x) = \{0, 1\}$, and $\lim_{x \rightarrow 2} f(x) = \{0, 1, 2\}$. Thus, the limit of a function at a point need not be unique even if the function is continuous at that point.

2.4 Open and Closed Sets

The fundamental objects of a topological space are its open sets: to know the open sets of a topological space *is* to know that space.¹⁷ By contrast, knowledge about a convergence space requires knowledge about which filters converge to which points. Thus, filters supersede the importance of open sets in the study of convergence spaces. Nevertheless, since neighborhoods, of which open sets are a special type, are important when examining pretopological spaces and closed sets are central to our discussion of mereology in Chapter 5, we will devote this chapter to a survey of the essential facts about open and closed sets.

Definition 2.38. Let X be a convergence space. For every $x \in X$, the *neighborhood filter of x* is the set

$$\mathcal{N}_x = \bigcap \{ \mathcal{F} \in \Phi(X) : \mathcal{F} \downarrow x \}.$$

Every element of \mathcal{N}_x is a *neighborhood of x* . A set is *open* if and only if it is a neighborhood of each of its members.

Definition 2.39. Let X be a convergence space. The *closure* of a subset A of X is the set

$$\text{cl}(A) = \{x \in X \mid (\exists \mathcal{F})(\mathcal{F} \downarrow x \wedge A \in \mathcal{F})\}.$$

A set A is *closed* if and only if $\text{cl}(A) = A$.

¹⁷Equivalently, to know the closed sets of a topological space is to know that space.

Definitions 2.38 and 2.39 coincide with the usual topological notions when the convergence space is itself topological.¹⁸ Although closure, as defined in Definition 2.39, is extensive and preserves both binary and nullary unions, it does not satisfy all of the Kuratowski closure axioms; in particular, closure is not an idempotent operation.¹⁹ Thus, some authors²⁰ write of the *adherence* of a subset rather than its closure. Propositions 2.40 through 2.42 confirm those Kuratowski closure axioms satisfied by closure as defined above; Example 2.43 shows that closure is not idempotent.

Proposition 2.40. *The empty set is closed.*

Proof. If $x \in \text{cl}(\emptyset)$, then there exists a filter converging to x that contains \emptyset ; no filter, however, contains the empty set. Thus $x \notin \text{cl}(\emptyset)$, and therefore $\text{cl}(\emptyset) = \emptyset$. \square

Proposition 2.41. *If A is a subset of a convergence space X , then $A \subseteq \text{cl}(A)$.*

Proof. If x belongs to A , then A belongs to $[x]$, which converges to x . Therefore x belongs to $\text{cl}(A)$. \square

Proposition 2.42. *If A and B are subsets of a convergence space X , then $\text{cl}(A \cup B) = \text{cl}(A) \cup \text{cl}(B)$.*

Proof. If $x \in \text{cl}(A \cup B)$, then there exists a filter \mathcal{F} converging to x that contains $A \cup B$. If $A \in \mathcal{F}$, then $x \in \text{cl}(A) \subseteq \text{cl}(A) \cup \text{cl}(B)$. If $A \notin \mathcal{F}$, then $\mathcal{G} = \{G : G \supseteq \mathcal{F} - A \wedge G \cap A \neq \emptyset\}$ is a filter finer than \mathcal{F} . Since $(A \cup B) - A \subseteq B$, it follows that $B \in \mathcal{G}$. Since $\mathcal{G} \supseteq \mathcal{F}$, it follows that $\mathcal{G} \downarrow x$, and so $x \in \text{cl}(B) \subseteq \text{cl}(A) \cup \text{cl}(B)$. Therefore $\text{cl}(A \cup B) \subseteq \text{cl}(A) \cup \text{cl}(B)$.

Conversely, suppose that $x \in \text{cl}(A) \cup \text{cl}(B)$. If $x \in \text{cl}(A)$, then there exists a filter \mathcal{F} converging to x that contains A . Since $A \subseteq A \cup B$, it follows that $A \cup B \in \mathcal{F}$. Thus $x \in \text{cl}(A \cup B)$. An exactly similar argument shows that $x \in \text{cl}(B)$ implies $x \in \text{cl}(A \cup B)$. Therefore $\text{cl}(A) \cup \text{cl}(B) \subseteq \text{cl}(A \cup B)$. \square

¹⁸See Definition 2.59.

¹⁹For details on the Kuratowski closure axioms, see Kelley [23].

²⁰For example, in [4].

Example 2.43. Define a convergence structure on $Y = \{0, 1, 2\}$ by the equivalences

$$\mathcal{F} \downarrow 0 \text{ if and only if } \{0, 1\} \in \mathcal{F}$$

$$\mathcal{F} \downarrow 1 \text{ if and only if } \{1, 2\} \in \mathcal{F}$$

$$\mathcal{F} \downarrow 2 \text{ if and only if } \{0, 2\} \in \mathcal{F}$$

This space has no nontrivial open or closed sets. Since $\text{cl}(\text{cl}(\{0\})) = \text{cl}(\{0, 2\}) = Y$, it follows that the closure operator of Definition 2.39, called the Katětov closure operator, is distinct from the Kuratowski closure operator of topological spaces.²¹

As in topological spaces, the complement of an open subset of a convergence space is closed; conversely, the complement of a closed subset is open. Moreover, the empty set and the convergence space itself are both open and closed.

Proposition 2.44. *If X is a convergence space, then a subset A of X is open if and only if $X - A$ is closed.*

Proof. [Necessity]. We proceed by contraposition. If A is not open, then there exists $x \in A$ and $\mathcal{F} \in \Phi(X)$ such that $\mathcal{F} \downarrow x$ but $A \notin \mathcal{F}$. Thus $\mathcal{G} = \{G : G \supseteq F - A \wedge F \in \mathcal{F}\}$ is a filter finer than \mathcal{F} that contains $X - A$. Since \mathcal{G} is finer than \mathcal{F} , it converges to x . Thus $x \in \text{cl}(X - A)$. Because $x \notin X - A$, we conclude that $X - A$ is not closed.

[Sufficiency]. If $y \in \text{cl}(X - A)$, then there exists a filter \mathcal{F} converging to y that contains $X - A$. Now $A \notin \mathcal{F}$; otherwise $\emptyset = (X - A) \cap A \in \mathcal{F}$. By hypothesis $y \notin A$, and so $y \in X - A$. Therefore $X - A$ is closed. \square

Proposition 2.45. *If X is a convergence space, then \emptyset and X are both open and closed.*

Proof. By Proposition 2.40, the empty set is closed, and so by Proposition 2.44 it follows that X is open. Since the empty set has no members, it is vacuously true that it is a

²¹For properties of the Katětov closure operator—also known as the preclosure or Čech closure operator—see Dikranjan [16].

neighborhood of each of its members; thus the empty set is open, and so by Proposition 2.44 it follows that X is closed. \square

Like open and closed sets, neighborhoods and closures are related by complementation: the neighborhoods of a point are precisely those sets the closures of the complements of which do not contain that point.

Proposition 2.46. *A subset N of a convergence space X is a neighborhood of $x \in X$ if and only if $x \notin \text{cl}(X - N)$.*

Proof. [Necessity]. We proceed by contraposition. If N is not a neighborhood of x , then there exists a filter \mathcal{F} that converges to x but does not contain N . Since the set $\mathcal{G} = \{B : B \supseteq F - N \wedge F \in \mathcal{F}\}$ is a filter finer than \mathcal{F} , it follows that \mathcal{G} converges to x . Since $X - N \in \mathcal{G}$, we conclude that $x \in \text{cl}(X - N)$.

[Sufficiency]. We proceed by contraposition. If $x \in \text{cl}(X - N)$, then there exists a filter \mathcal{F} that converges to x and contains $X - N$. Since $N \cap (X - N) = \emptyset$, it follows that $N \notin \mathcal{F}$. Therefore N is not a neighborhood of x . \square

For functions between topological spaces, continuity is equivalent to preservation of open sets under preimages; for functions between convergence spaces, however, preservation of open sets under preimages is necessary but not sufficient for continuity.

Proposition 2.47. *Let X and Y be convergence spaces. If $f : X \rightarrow Y$ is continuous, then $f^{-1}(U)$ is open in X whenever U is open in Y .*

Proof. If $x \in f^{-1}(U)$, then $f(x) \in f(f^{-1}(U)) \subseteq U$. If $\mathcal{F} \downarrow x$ in X , then by hypothesis $f(\mathcal{F}) \downarrow f(x)$. Since U is open in Y , it follows that $U \in f(\mathcal{F})$, which implies that there exists $V \in \mathcal{F}$ such that $f(V) \subseteq U$; thus $V \subseteq f^{-1}(U)$, and so $f^{-1}(U) \in \mathcal{F}$. Since $f^{-1}(U)$ belongs to each filter converging to one of its elements, we conclude that $f^{-1}(U)$ is open in X . \square

Proposition 2.48. *Let X and Y be convergence spaces. If $f : X \rightarrow Y$ is continuous, then $f^{-1}(V)$ is closed in X whenever V is closed in Y .*

Proof. Suppose that V is closed in Y . Then $Y - V$ is open in Y , and so $X - f^{-1}(V) = f^{-1}(Y - V)$ is open in X . Therefore $f^{-1}(V)$ is closed in X . \square

Example 2.49. The converse of Proposition 2.47 is not true. The only nonempty open set of the convergence space Y of Example 2.43 is Y itself. Consider the function $f : Y \rightarrow Y$ defined by $(0, 1)$. It is vacuously true that the preimage under f of each open set is also an open set. This function, however, is not continuous since $[0] \downarrow 2$ but $f([0]) = [1]$, which does not converge to 2. In fact, there are only six continuous functions on Y : the constant functions, the identity function, $(0, 1, 2)$, and $(0, 2, 1)$. Of these functions, the last three are automorphisms; thus Y is a homogeneous space.

We conclude this chapter with the observation that a continuous function has the property, stated informally, that the closure of its image includes its image of the closure

Proposition 2.50. *If $f : X \rightarrow Y$ is continuous, then $f(\text{cl}(A)) \subseteq \text{cl}(f(A))$ for each $A \subseteq X$.*

Proof. If $y \in f(\text{cl}(A))$, then there exists $x \in \text{cl}(A)$ such that $f(x) = y$. Since $x \in \text{cl}(A)$, there exists a filter \mathcal{F} converging to x that contains A . Thus $f(A) \in f(\mathcal{F})$ and by hypothesis $f(\mathcal{F}) \downarrow y$. Therefore $y \in \text{cl}(f(A))$, as desired. \square

2.5 Common Types of Convergence Structures

In this chapter, we present several common types of convergence structures. First, we study initial convergence structures with particular attention to the subspace and product convergence structures. Next, we examine topological, pretopological, and pseudotopological spaces. These are closely related structures since every topological space is a pretopological space and, in turn, every pretopological space is a pseudotopological space. Pseudotopological spaces are particularly interesting because they, like convergence spaces but unlike topological or pretopological spaces, form a Cartesian closed category.²² Finally, we dis-

²²See Proposition 2.80.

cuss discrete and indiscrete structures, which despite their simplicity play a pivotal role throughout this work.²³

2.5.1 Initial Convergence Structures

Definition 2.51. Let X be a set, let $(X_i)_{i \in I}$ be a collection of convergence spaces, and let $(f_i : X \rightarrow X_i)_{i \in I}$ be a family of functions. The *initial convergence structure on X with respect to $(f_i)_{i \in I}$* is defined by

$$\mathcal{F} \downarrow x \text{ if and only if } f_i(\mathcal{F}) \downarrow f_i(x) \text{ for each } i \in I.$$

For a function into a convergence space initial with respect to a family of functions, it is sufficient that its composition on the right with each member of the family is also continuous.

Proposition 2.52. *Let $(X_i)_{i \in I}$ be a collection of convergence spaces, let $(f_i : X \rightarrow X_i)_{i \in I}$ be a family of functions, and let X have the initial convergence structure with respect to $(f_i)_{i \in I}$. A function $g : Y \rightarrow X$ is continuous if and only if $f_i \circ g$ is continuous for each $i \in I$.*

Proof. [Necessity]. If $\mathcal{F} \downarrow y$ in Y , then $(f_i \circ g)(\mathcal{F}) \downarrow (f_i \circ g)(y)$ in X_i for each $i \in I$; equivalently $f_i(g(\mathcal{F})) \downarrow f_i(g(y))$. Since X has the initial convergence structure with respect to $(f_i)_{i \in I}$, it follows that $g(\mathcal{F}) \downarrow g(y)$. Therefore g is continuous.

[Sufficiency]. Since each f_i is continuous, it follows that $f_i \circ g$ is also continuous. \square

Two types of initial convergence structures are the subspace convergence structure and the product convergence structure. The subspace convergence structure is the coarsest convergence structure on a subset of a convergence space for which the inclusion function is continuous; likewise, the product convergence structure is the coarsest convergence structure on the Cartesian product of convergence spaces for which the projections are continuous.

²³See Sections 4.2.2 and 4.5 as well as Chapter 5.

Definition 2.53. Let X be a convergence space and let Y be a subset of X . The *subspace convergence structure* on Y is the initial convergence structure on Y with respect to the inclusion function $\iota : Y \rightarrow X$.

Definition 2.54. Let $(X_i)_{i \in I}$ be a collection of convergence spaces. The *product convergence structure* on $\prod_{i \in I} X_i$ is the initial convergence structure on $\prod_{i \in I} X_i$ with respect to the projections $\pi_i : \prod_{i \in I} X_i \rightarrow X_i$.

In some circumstances²⁴ we wish to “put” one space into another. To formalize this intuitive desire, we must show that the one space is homeomorphic to a subspace of the other. We call the relevant homeomorphism an embedding.

Definition 2.55. An injection $e : X \rightarrow Y$ is an *embedding* if and only if X and $e(X)$ are homeomorphic; we say that Y *embeds* X if and only if there exists an embedding $e : X \rightarrow Y$.

Given a family of embeddings $(f_i)_{i \in I}$ from a convergence space X onto a convergence space Y , we can construct an embedding from $\prod_I X$ onto $\prod_I Y$.

Proposition 2.56. *Let X and Y be convergence spaces, let I be an index set, and let $\pi_i : \prod_{i \in I} X \rightarrow X$ and $\rho_i : \prod_{i \in I} Y \rightarrow Y$ be projections. If $f_i : X \rightarrow Y$ is an embedding for each $i \in I$, then $f : \prod_I X \rightarrow \prod_I Y$ defined by $\rho_i(f(p)) = f_i(\pi_i(p))$ for each $i \in I$ and $p \in \prod_I X$ is an embedding.*

Proof. To see that f is injective, suppose that $f(p) = f(q)$ for some $p, q \in \prod_I X$. It follows that $f_i(\pi_i(p)) = \rho_i(f(p)) = \rho_i(f(q)) = f_i(\pi_i(q))$. Since each f_i is injective, we infer that $\pi_i(p) = \pi_i(q)$ for each $i \in I$, and so $p = q$.

To see that f is continuous, suppose that \mathcal{F} converges to p in $\prod_I X$. Thus $\pi_i(\mathcal{F})$ converges to $\pi_i(p)$ in X for each $i \in I$. Since each f_i is continuous, it follows that $\rho_i(f(\mathcal{F})) = f_i(\pi_i(\mathcal{F}))$ converges to $f_i(\pi_i(p)) = \rho_i(f(p))$ for each $i \in I$. Thus $f(\mathcal{F})$ converges to $f(p)$.

To see that f^{-1} is continuous, suppose that \mathcal{G} converges to q in $f(\prod_I X)$. Thus $\rho_i(\mathcal{G})$ converges to $\rho_i(q)$ in Y for each $i \in I$. Observe that $\pi_i(f^{-1}(q)) = f_i^{-1}(f_i(\pi_i(f^{-1}(q)))) =$

²⁴For example, in Chapter 6

$f_i^{-1}(\rho_i(f(f^{-1}(q)))) = f_i^{-1}(\rho_i(q))$ for each $i \in I$. Since each f_i^{-1} is continuous, it follows that $\pi_i(f^{-1}(\mathcal{G})) = f_i^{-1}(\rho_i(\mathcal{G}))$ converges to $f_i^{-1}(\rho_i(q)) = \pi_i(f^{-1}(q))$ for each $i \in I$. Thus $f^{-1}(\mathcal{G}) = f^{-1}(\mathcal{G})$ converges to $f^{-1}(q)$. \square

It is a consequence of Proposition 2.56 that $\prod_I H$ is homogeneous whenever H is homogeneous.

Proposition 2.57. *If H is a homogeneous space and I is an index set, then $\prod_I H$ is a homogeneous space.*

Proof. If $x, y \in \prod_I H$, then for each $i \in I$, there exists an automorphism $f_i : H \rightarrow H$ such that $f_i(\pi_i(x)) = \pi_i(y)$. By Proposition 2.56, the function $f : \prod_I H \rightarrow \prod_I H$ defined by $\pi_i(f(p)) = f_i(\pi_i(p))$ for each $i \in I$ and $p \in \prod_I X$ is an embedding. Since each f_i is an automorphism, it follows that f is an automorphism. Since $\pi_i(f(x)) = f_i(\pi_i(x)) = \pi_i(y)$ for each $i \in I$, it follows that $f(x) = y$. Thus f is automorphism that maps x to y . Therefore $\prod_I H$ is homogeneous. \square

An interesting property of the continuous convergence structure is that it has a subspace homeomorphic to the codomain.

Proposition 2.58. *If X and Y are convergence spaces, then $\mathcal{C}(X, Y)$ embeds Y .*

Proof. For each $y \in Y$, define $\bar{y} : X \rightarrow Y$ by $\bar{y}(x) = y$ for every $x \in X$. Define $\phi : Y \rightarrow \mathcal{C}(X, Y)$ by $\phi(y) = \bar{y}$. It is clear that ϕ is an injection. We will show that Y and $\phi(Y)$ are homeomorphic.

To see that ϕ is continuous, suppose that \mathcal{A} converges to y in Y and \mathcal{B} converges to x in X . If $A \in \mathcal{A}$, then $A = \phi(A) \cdot B \in \phi(\mathcal{A}) \cdot \mathcal{B}$ for every $B \in \mathcal{B}$; thus $\mathcal{A} \subseteq \phi(\mathcal{A}) \cdot \mathcal{B}$, which implies that $\phi(\mathcal{A}) \cdot \mathcal{B}$ converges to $y = \bar{y}(x) = (\phi(y))(x)$. Therefore $\phi(\mathcal{A})$ converges to $\phi(y)$.

To see that ϕ^{-1} is continuous, suppose that \mathcal{F} converges to \bar{y} in $\phi(Y)$. If \mathcal{G} converges to x in X , then $\mathcal{F} \cdot \mathcal{G}$ converges to $\bar{y}(x) = y$ in Y . Moreover, if $H \in \mathcal{F} \cdot \mathcal{G}$, then there exist $F \in \mathcal{F}$ and $G \in \mathcal{G}$ such that $F \cdot G \subseteq H$. Since $F \cdot G = \phi^{-1}(F) \in \phi^{-1}(\mathcal{F})$, it follows that $H \in \phi^{-1}(\mathcal{F})$; thus $\mathcal{F} \cdot \mathcal{G} \subseteq \phi^{-1}(\mathcal{F})$. Therefore $\phi^{-1}(\mathcal{F})$ converges to $y = \phi^{-1}(\bar{y})$. \square

2.5.2 Topological Convergence Structures

Definition 2.59. Let (X, \mathcal{T}) be a topological space. For every $x \in X$, define \mathcal{U}_x , the set of all topological neighborhoods of x . The *topological convergence structure* on X is defined by

$$\mathcal{F} \downarrow x \text{ if and only if } \mathcal{F} \supseteq \mathcal{U}_x.$$

A convergence space is *topological* if and only if it has the topological convergence structure.

Every topological space is a topological convergence space. To see this, observe that if (X, \mathcal{T}) is a topological space and (X, \downarrow) is its corresponding topological convergence space, then $U \in \mathcal{T}$ if and only if U is open; that is, the open sets of (X, \mathcal{T}) and (X, \downarrow) are identical. Thus, a topological convergence space is defined by the open sets of the inducing topology.

Example 2.60. A convergence space need not be a topological space. Suppose that the convergence space Y of Example 2.43 is a topological convergence space; that is, some topology \mathcal{T} induces the convergence structure \downarrow . It follows that $\mathcal{U}_0 = [0, 1]$ and $\mathcal{U}_1 = [1, 2]$. Thus $\{0, 1\}$ and $\{1, 2\}$ must both belong to \mathcal{T} , and so $\{1\}$ also belongs to \mathcal{T} , which implies that $\{1\} \in \mathcal{U}_1$; but $\{1\} \notin [1, 2]$. Therefore Y is not a topological convergence space.²⁵

Next we verify that Definition 2.27 correctly generalizes continuity.

Proposition 2.61. *Let X and Y be topological convergence spaces. A function $f : X \rightarrow Y$ is continuous if and only if $f^{-1}(U)$ is open in X whenever U is open in Y .*

Proof. [Necessity]. The desired result follows immediately from Proposition 2.47.

[Sufficiency]. If $\mathcal{F} \downarrow x$, then $\mathcal{F} \supseteq \mathcal{U}_x$. If $U \in \mathcal{U}_{f(x)}$, then U includes a set V that is open in Y and contains $f(x)$. Thus by hypothesis $f^{-1}(V)$ is open in X . Since $f(x) \in V$, it follows that $x \in f^{-1}(V)$, which implies that $f^{-1}(V) \in \mathcal{U}_x$; since $f^{-1}(V) \subseteq f^{-1}(U)$, it follows that

²⁵Another argument is that since the Katětov closure operator is not idempotent on Y , the convergence space Y cannot be topological. Yet another is that as represented by a reflexive digraph Y is not transitive, and so not topological—see Definition 3.1 and Proposition 3.11.

$f^{-1}(U) \in \mathcal{U}_x$. Because $f(f^{-1}(U)) \subseteq U$ we infer that $U \in f(\mathcal{U}_x) \subseteq f(\mathcal{F})$; thus $\mathcal{U}_{f(x)} \subseteq f(\mathcal{F})$, and so $f(\mathcal{F}) \downarrow f(x)$. Therefore, we conclude that f is continuous. \square

The topological convergence structure is preserved by homeomorphisms.

Proposition 2.62. *Any convergence space homeomorphic to a topological space is also topological.*

Proof. Let X be a convergence space, let Y be a topological space, and let $f : X \rightarrow Y$ be a homeomorphism. We must show that there exists a topology \mathcal{T} on X such that $\mathcal{F} \downarrow x$ if and only if $\mathcal{F} \supseteq \mathcal{U}_x$.

Let $U \in \mathcal{T}$ if and only if $f(U)$ is included by a set open in Y .

First, we verify that \mathcal{T} is a topology on X . Since $f(X) = Y$ is open in Y , it follows that $X \in \mathcal{T}$; likewise, since $f(\emptyset) = \emptyset$ is open in Y , it follows that $\emptyset \in \mathcal{T}$. Suppose that V_α belongs to \mathcal{T} for each α of some index set A . Since $f(\bigcup_{\alpha \in A} V_\alpha) = \bigcup_{\alpha \in A} f(V_\alpha)$ is open in Y , it follows that $f(\bigcup_{\alpha \in A} V_\alpha) \in \mathcal{T}$. Suppose that V and V' belong to \mathcal{T} . Since $f(V \cap V') \subseteq f(V) \cap f(V')$ and $f(V) \cap f(V')$ is open in Y , it follows that $V \cap V' \in \mathcal{T}$.

Now we show that $\mathcal{F} \downarrow x$ if and only if $\mathcal{F} \supseteq \mathcal{U}_x$.

[Necessity]. If U is a neighborhood of $f(x)$, then $f^{-1}(U)$ is a neighborhood of x . By hypothesis, there exists $F \in \mathcal{F}$ such that $F \subseteq f^{-1}(U)$, which implies that $f(F) \subseteq U$. Thus $\mathcal{U}_{f(x)} \subseteq f(\mathcal{F})$, and so $f(\mathcal{F}) \downarrow f(x)$. Therefore $\mathcal{F} \downarrow x$.

[Sufficiency]. If $V \in f(\mathcal{U}_x)$, then there exists $U \in \mathcal{U}_x$ such that $f(U) \subseteq V$. By construction $f(U) \in \mathcal{U}_{f(x)}$; thus $V \in \mathcal{U}_{f(x)}$, and so $f(\mathcal{U}_x) \subseteq \mathcal{U}_{f(x)}$. Since f is continuous and $\mathcal{F} \downarrow x$, it follows that $f(\mathcal{F}) \downarrow f(x)$, which implies that $\mathcal{U}_{f(x)} \subseteq f(\mathcal{F})$. Thus $f(\mathcal{U}_x) \subseteq f(\mathcal{F})$, and therefore $\mathcal{U}_x = f^{-1}(f(\mathcal{U}_x)) \subseteq f^{-1}(f(\mathcal{F})) = \mathcal{F}$. \square

If a space is initial with respect to a family of functions into a collection of topological spaces, then it is also topological.

Proposition 2.63. *Let $(X_i)_{i \in I}$ be a collection of convergence spaces, let $(f_i : X \rightarrow X_i)_{i \in I}$ be a family of functions, and let X have the initial convergence structure with respect to $(f_i)_{i \in I}$. If each $(X_i)_{i \in I}$ is topological, then X is topological.*

Proof. We must show that there exists a topology \mathcal{T} on X such that $\mathcal{F} \downarrow x$ if and only if $\mathcal{F} \supseteq \mathcal{U}_x$.

Let $V \in \mathcal{T}$ if and only if for every $x \in V$ and for every $i \in I$ there exists an open set U_i of X_i such that $x \in f_i^{-1}(U_i) \subseteq V$.

First, we verify that \mathcal{T} is a topology on X . Since X_i is open in X_i and $X = f_i^{-1}(X_i)$, it follows that $X \in \mathcal{T}$; likewise, since \emptyset is open in X_i and $\emptyset = f_i^{-1}(\emptyset)$, it follows that $\emptyset \in \mathcal{T}$. Suppose that V_α belongs to \mathcal{T} for each α of some index set A . If $x \in \bigcup_{\alpha \in A} V_\alpha$, then $x \in V_\alpha$ for some $\alpha \in A$, and so for every $i \in I$ there exists an open set U_i of X_i such that $x \in f_i^{-1}(U_i) \subseteq V_\alpha \subseteq \bigcup_{\alpha \in A} V_\alpha$. Thus $\bigcup_{\alpha \in A} V_\alpha \in \mathcal{T}$. Suppose that V and V' belong to \mathcal{T} . If $x \in V \cap V'$, then for every $i \in I$ there exists an open set U_i of X_i such that $x \in f_i^{-1}(U_i) \subseteq V$ and there exists an open set U_i' of X_j such that $x \in f_i^{-1}(U_i') \subseteq V'$. Since U_i and U_i' are both open in X_i , it follows that $U_i \cap U_i'$ is also open in X_i ; moreover $x \in f_i^{-1}(U_i) \cap f_i^{-1}(U_i') = f_i^{-1}(U_i \cap U_i')$ and $f_i^{-1}(U_i \cap U_i') = f_i^{-1}(U_i) \cap f_i^{-1}(U_i') \subseteq V \cap V'$. Thus $V \cap V' \in \mathcal{T}$.

Now we show that $\mathcal{F} \downarrow x$ if and only if $\mathcal{F} \supseteq \mathcal{U}_x$.

[Necessity]. If $U \in \mathcal{U}_{f_i(x)}$, then U includes an open set V that contains $f_i(x)$. By hypothesis $f_i^{-1}(V)$ is an open set that contains x , and so $f_i^{-1}(U) \in \mathcal{U}_x$. Since $f_i(f_i^{-1}(U)) \subseteq U$, it follows that $U \in f_i(\mathcal{U}_x)$. Thus $\mathcal{U}_{f_i(x)} \subseteq f_i(\mathcal{U}_x) \subseteq f_i(\mathcal{F})$, which implies that $f_i(\mathcal{F}) \downarrow f_i(x)$. Since this argument holds for each $i \in I$, we conclude that $\mathcal{F} \downarrow x$.

[Sufficiency]. If $U \in \mathcal{U}_x$, then U includes an open set V that contains x . Thus, for every $i \in I$, there exists an open set W_i of X_i such that $x \in f_i^{-1}(W_i) \subseteq V \subseteq U$. Since $\mathcal{F} \downarrow x$, it follows that $f_i(\mathcal{F}) \downarrow f_i(x)$, and so $W_i \in \mathcal{U}_{f_i(x)} \subseteq f_i(\mathcal{F})$. Thus, there exists $F \in \mathcal{F}$ such that $f_i(F) \subseteq W_i$, which implies that $F \subseteq f_i^{-1}(W_i) \subseteq U$. Therefore $\mathcal{U}_x \subseteq \mathcal{F}$. \square

2.5.3 Pretopological Convergence Structures

Definition 2.64. A convergence space X is *pretopological* if and only if $\mathcal{N}_x \downarrow x$ for every $x \in X$.

It is frequently more useful to formulate Definition 2.64 in terms of inclusion of the neighborhood filter rather than its convergence.

Proposition 2.65. *If X is a pretopological space and $x \in X$, then $\mathcal{F} \downarrow x$ if and only if $\mathcal{F} \supseteq \mathcal{N}_x$.*

Proof. [Necessity]. Since X is pretopological, it follows that $\mathcal{N}_x \downarrow x$, and so $\mathcal{F} \downarrow x$.

[Sufficiency]. If $N \in \mathcal{N}_x$, then N is a neighborhood of x ; thus $N \in \mathcal{F}$. □

By Proposition 2.65, it follows that every topological space is a pretopological space. Pretopological spaces, however, need not be topological: the convergence space Y of Example 2.43 is a pretopological space that is not topological. Thus, in view of Example 2.49, we see that Proposition 2.61 does not hold for pretopological spaces. Nevertheless, for functions into pretopological spaces, continuity may be formulated in terms of neighborhoods or neighborhood filters.

Proposition 2.66. *If X is a convergence space and Y is a pretopological space, then the function $f : X \rightarrow Y$ is continuous at $x \in X$ if and only if for every neighborhood V of $f(x)$, there exists a neighborhood U of x such that $f(U) \subseteq V$.*

Proof. [Necessity]. Let $\mathcal{F} \downarrow x$. By hypothesis for every $V \in \mathcal{N}_{f(x)}$, there exists $U \in \mathcal{N}_x \subseteq \mathcal{F}$ such that $f(U) \subseteq V$; thus $V \in f(\mathcal{F})$, and so $\mathcal{N}_{f(x)} \subseteq f(\mathcal{F})$, which implies that $f(\mathcal{F}) \downarrow f(x)$.

[Sufficiency]. Let V be a neighborhood of $f(x)$. If $\mathcal{F} \downarrow x$ in X , then by hypothesis $f(\mathcal{F}) \downarrow f(x)$. Since Y is pretopological, it follows that $V \in f(\mathcal{F})$, which implies that there exists $F \in \mathcal{F}$ such that $F \subseteq f^{-1}(V)$; thus $f^{-1}(V) \in \mathcal{F}$. Since $f^{-1}(V)$ belongs to each filter converging to x , we conclude that $f^{-1}(V)$ is a neighborhood of x . □

Proposition 2.67. *If X and Y are pretopological spaces, then $f : X \rightarrow Y$ is continuous if and only if $\mathcal{N}_{f(x)} \subseteq f(\mathcal{N}_x)$ for each $x \in X$.*

Proof. [Necessity]. Since Y is pretopological, it follows that $f(\mathcal{N}_x) \downarrow f(x)$ for each $x \in X$; thus f is continuous.

[Sufficiency]. Since X is pretopological, it follows that $\mathcal{N}_x \downarrow x$, and so by hypothesis $f(\mathcal{N}_x) \downarrow f(x)$, which implies that $\mathcal{N}_{f(x)} \subseteq f(\mathcal{N}_x)$. \square

Proposition 2.68. *If X and Y are pretopological spaces, then a bijection $f : X \rightarrow Y$ is a homeomorphism if and only if $\mathcal{N}_{f(x)} = f(\mathcal{N}_x)$ for each $x \in X$.*

Proof. [Necessity]. By Proposition 2.67, it follows that f is continuous. If $y \in Y$, then there exists $x \in X$ such that $y = f(x)$, and so by hypothesis $f(\mathcal{N}_x) \subseteq \mathcal{N}_{f(x)}$, which implies that $\mathcal{N}_{f^{-1}(y)} = \mathcal{N}_x \subseteq f^{-1}(\mathcal{N}_{f(x)}) = f^{-1}(\mathcal{N}_y)$; thus $f^{-1}(\mathcal{N}_y) \downarrow f^{-1}(y)$. Therefore f^{-1} is continuous.

[Sufficiency]. If $x \in X$, then there exists $y \in Y$ such that $x = f^{-1}(y)$. Since f^{-1} is continuous, it follows that $\mathcal{N}_{f^{-1}(y)} \subseteq f^{-1}(\mathcal{N}_y)$; thus $f(\mathcal{N}_x) = f(\mathcal{N}_{f^{-1}(y)}) \subseteq \mathcal{N}_y = \mathcal{N}_{f(x)}$. By Proposition 2.67, it follows that $\mathcal{N}_{f(x)} \subseteq f(\mathcal{N}_x)$; therefore $\mathcal{N}_{f(x)} = f(\mathcal{N}_x)$. \square

As with topologies, pretopologies are preserved by homeomorphisms.

Proposition 2.69. *Any convergence homeomorphic to a pretopological space is also pretopological.*

Proof. Let X be a convergence space, let Y be a pretopological space, and let $f : X \rightarrow Y$ be a homeomorphism. Suppose that $\mathcal{F} \supseteq \mathcal{N}_x$. If $N \in \mathcal{N}_{f(x)}$, then there exists $M \in \mathcal{N}_x \subseteq \mathcal{F}$ such that $f(M) \subseteq N$; thus $\mathcal{N}_{f(x)} \subseteq f(\mathcal{F})$, which implies that $f(\mathcal{F}) \downarrow f(x)$, and so $\mathcal{F} = f^{-1}(f(\mathcal{F}))$ converges to $f^{-1}(f(x)) = x$. \square

With Proposition 2.66 we can verify the claim of Section 2.3 that *limit* as defined in Definition 2.36 coincides with the standard definition of limit for functions on the Euclidean line.

Proposition 2.70. *If X is a convergence space, Y is a pretopological space, $f : X \rightarrow Y$ is a function, $p \in X$, and $l \in Y$, then $l \in \lim_{x \rightarrow p} f(x)$ if and only if for every neighborhood V of l , there exists a neighborhood U of p such that $f_{p \rightarrow l}(U) \subseteq V$.*

Proof. The result follows immediately from Definition 2.36 and Proposition 2.66. \square

Proposition 2.71. *Let $f : \mathbb{R} \rightarrow \mathbb{R}$ be a function and let $p \in \mathbb{R}$. If $l \in \mathbb{R}$, then l is a limit of f as x approaches p if and only if for every $\varepsilon > 0$, there exists $\delta > 0$ such that $|f(x) - l| < \varepsilon$ for every $x \in \mathbb{R}$ for which $0 < |x - p| < \delta$.*

Proof. [Necessity]. If V is a neighborhood of l , then $V = \{y \in \mathbb{R} : |y - l| < \varepsilon\}$ for some $\varepsilon > 0$. By hypothesis, there exists $\delta > 0$ such that $|f(x) - l| < \varepsilon$ for every x for which $0 < |x - p| < \delta$. Let U be the neighborhood $\{x \in \mathbb{R} : |x - p| < \delta\}$. Thus $f(x) \in V$ whenever $x \in U - \{p\}$, which implies that $f(U - \{p\}) \subseteq V$, and so $f_{p \rightarrow l}(U) \subseteq V$. Therefore l is a limit of f as x approaches p .

[Sufficiency]. Let $\varepsilon > 0$. The set $V = \{y \in \mathbb{R} : |y - l| < \varepsilon\}$ is a neighborhood of l . Thus, there exists a neighborhood U of p such that $f_{p \rightarrow l}(U) \subseteq V$, which implies that $f(U - \{p\}) \subseteq V$. Since U is a neighborhood of p , it follows that $U = \{x \in \mathbb{R} : |x - p| < \delta\}$ for some $\delta > 0$. If $0 < |x - p| < \delta$ for some $x \in \mathbb{R}$, then $x \in U - \{p\}$, and so $f(x) \in V$, which implies that $|f(x) - l| < \varepsilon$. \square

If a space is initial with respect to a family of functions into a collection of pretopological, then it is also pretopological.

Proposition 2.72. *Let $(X_i)_{i \in I}$ be a collection of convergence spaces, let $(f_i : X \rightarrow X_i)_{i \in I}$ be a family of functions, and let X have the initial convergence structure with respect to $(f_i)_{i \in I}$. If each $(X_i)_{i \in I}$ is pretopological, then X is pretopological.*

Proof. If $V \in \mathcal{N}_{f_i(x)}$, then by the continuity of f_i , there exists $U \in \mathcal{N}_x$ such that $f_i(U) \subseteq V$, which implies that $V \in f_i(\mathcal{N}_x)$; thus $\mathcal{N}_{f_i(x)} \subseteq f_i(\mathcal{N}_x)$, from which it follows that $f_i(\mathcal{N}_x) \downarrow f_i(x)$. Since this argument holds for each $i \in I$, we conclude that $\mathcal{N}_x \downarrow x$. \square

A notable distinction—albeit one unnoted in the literature—between convergence spaces as defined in Definition 2.24 and so-called generalized convergence spaces is that when finite the former are pretopological but the latter need not be.

Theorem 2.73. *Every finite convergence space is a pretopological space.*

Proof. Let X be a finite convergence space and let $x \in X$. Suppose that N is the smallest neighborhood of x . It follows that the neighborhood filter of x must be $[N]$. Suppose that there exists $y \in N$ such that $[y] \not\downarrow x$. Also suppose by way of contradiction that $[A] \downarrow x$ but $N - \{y\} \notin [A]$; then $A \not\subseteq N - \{y\}$. Since $N \in [A]$, it follows that $A \subseteq N$, which implies that $A \cap \{y\} \neq \emptyset$. Thus $y \in A$, and so $[A] \subseteq [y]$, from which it follows that $[y] \downarrow x$, in contradiction to the previous assumption. Hence, every filter converging to x must contain $N - \{y\}$. This implies that $N - \{y\}$ is a neighborhood of x , from which it follows that $N - \{y\} \in [N]$, and so $N \subseteq N - \{y\}$, which is absurd. Thus, it must be the case that $[y] \downarrow x$ for every $y \in N$. Since $[N] = \bigcap_{y \in N} [y]$ and N is finite, it follows that $[N] \downarrow x$. Therefore X is pretopological. \square

2.5.4 Pseudotopological Convergence Structures

Definition 2.74. A convergence space X is a *pseudotopological space* if and only if a filter \mathcal{F} converges to x whenever every ultrafilter finer than \mathcal{F} converges to x .

Some authors²⁶ refer to pseudotopological spaces as *Choquet spaces*. To describe the convergence structure on a pseudotopological space, it suffices to specify to which points each ultrafilter converges and that the ultrafilters satisfy the conditions of Definition 2.24.

Theorem 2.75. *Let X be a set. A relation \downarrow between $\Phi(X)$ and X is a pseudotopological convergence structure if and only if for every $x \in X$ and for every filter \mathcal{F} on X :*

1. $[x] \downarrow x$, and

²⁶For example, in [4].

2. $\mathcal{F} \downarrow x$ if and only if every ultrafilter finer than \mathcal{F} also converges to x .

Proof. [Necessity]. First, suppose that $\mathcal{F} \downarrow x$ and $\mathcal{F} \subseteq \mathcal{G}$. If \mathcal{U} is an ultrafilter finer than \mathcal{G} , then $\mathcal{U} \supseteq \mathcal{F}$, and so $\mathcal{U} \downarrow x$; thus $\mathcal{G} \downarrow x$. Second, suppose that $\mathcal{F} \downarrow x$ and $\mathcal{G} \downarrow x$. If \mathcal{U} is an ultrafilter finer than $\mathcal{F} \cap \mathcal{G}$, then \mathcal{U} is finer than the intersection of all ultrafilters finer than either \mathcal{F} or \mathcal{G} ; thus \mathcal{U} must be finer than either \mathcal{F} or \mathcal{G} , and so \mathcal{U} converges to x , which implies that $\mathcal{F} \cap \mathcal{G}$ converges to x .

[Sufficiency]. The desired result follows immediately from Definitions 2.24 and 2.74. \square

Another elementary fact, apparently neglected in the literature, is that for a function between pseudotopological spaces to be continuous it is sufficient that convergence of ultrafilters be preserved under its image. In this way, the collection of ultrafilters on a pseudotopological space are loosely analogous to the basis vectors of a vector space.

Theorem 2.76. *Let X be a convergence spaces and let Y be a pseudotopological space. A function $f : X \rightarrow Y$ is continuous if and only if $f(\mathcal{U}) \downarrow f(x)$ whenever $\mathcal{U} \downarrow x$ for every ultrafilter \mathcal{U} on X .*

Proof. [Necessity]. Suppose that $\mathcal{F} \downarrow x$. If \mathcal{U} is finer than $f(\mathcal{F})$, then there exists an ultrafilter \mathcal{V} finer than \mathcal{F} such that $f(\mathcal{V}) = \mathcal{U}$. Since $\mathcal{V} \downarrow x$, it follows that $f(\mathcal{V}) \downarrow f(x)$; thus $\mathcal{U} \downarrow f(x)$. By the hypothesis that Y is pseudotopological, we infer that $f(\mathcal{F}) \downarrow f(x)$. Therefore f is continuous.

[Sufficiency]. This follows immediately by definition of continuity. \square

As with topologies and pretopologies, pseudotopologies are preserved by homeomorphisms.

Proposition 2.77. *Any convergence space homeomorphic to a pseudotopological space is also pseudotopological.*

Proof. Let X be a convergence space, let Y be a pseudotopological space, and let $f : X \rightarrow Y$ be a homeomorphism. Suppose that every ultrafilter finer than \mathcal{F} converges to x . If \mathcal{F} does

not converge to x , then $f(\mathcal{F})$ does not converge to $f(x)$. Thus, there exists an ultrafilter \mathcal{V} finer than $f(\mathcal{F})$ that does not converge to $f(x)$, and so there exists an ultrafilter \mathcal{U} finer than \mathcal{F} such that $f(\mathcal{U}) = \mathcal{V}$; but by hypothesis, it follows that $\mathcal{U} \downarrow x$, and so $\mathcal{V} \downarrow f(x)$, in contradiction to the result that \mathcal{V} does not converge to $f(x)$. Therefore $\mathcal{F} \downarrow x$. \square

If a space is initial with respect to a family of functions into a collection of pseudotopological, then it is also pseudotopological.

Proposition 2.78. *Let $(X_i)_{i \in I}$ be a collection of convergence spaces, let $(f_i : X \rightarrow X_i)_{i \in I}$ be a family of functions, and let X have the initial convergence structure with respect to $(f_i)_{i \in I}$. If each $(X_i)_{i \in I}$ is pseudotopological, then X is pseudotopological.*

Proof. Suppose that every ultrafilter finer than \mathcal{F} converges to $x \in X$. If \mathcal{V} is an ultrafilter finer than $f_i(\mathcal{F})$, then there exists an ultrafilter \mathcal{U} finer than \mathcal{F} such that $f_i(\mathcal{U}) = \mathcal{V}$. Since f_i is continuous, it follows that $\mathcal{V} \downarrow f_i(x)$; thus $f_i(\mathcal{F}) \downarrow f_i(x)$. Since this argument holds for each $i \in I$, we conclude that $\mathcal{F} \downarrow x$. \square

Every pretopological space is a pseudotopological space; pseudotopological spaces, however, need not be pretopological.²⁷

Proposition 2.79. *Every pretopological space is a pseudotopological space.*

Proof. Let X be a pretopological space. Suppose that \mathcal{F} is a filter on X such that $\mathcal{U} \downarrow x$ for every ultrafilter \mathcal{U} finer than \mathcal{F} . If \mathcal{F} does not converge to x , then by hypothesis \mathcal{F} is not finer than \mathcal{N}_x ; in particular there exists $N \in \mathcal{N}_x$ such that $N \notin \mathcal{F}$, and so $X - N \in \mathcal{U}$ for every ultrafilter \mathcal{U} finer than \mathcal{F} . Since, however, any such ultrafilter converges to x , it must contain N ; thus $\emptyset = N \cap (X - N)$ also belongs to the ultrafilter, which is absurd. Therefore $\mathcal{F} \downarrow x$. \square

Like convergence spaces but unlike topological or pretopological spaces, pseudotopological spaces form a Cartesian closed category: whereas the continuous convergence space of

²⁷See Example 2.118.

two topological or pretopological need not be topological or pretopological, the continuous convergence space of two pseudotopological spaces must be pseudotopological. In fact, if the continuous convergence space of any two convergence spaces is pseudotopological, then the domain space must be pseudotopological.

Proposition 2.80. *If X and Y are convergence spaces, then $\mathcal{C}(X, Y)$ is pseudotopological if and only if Y is pseudotopological.*

Proof. [Necessity]. Suppose that every ultrafilter finer than \mathcal{F} converges to $f \in \mathcal{C}(X, Y)$. Let \mathcal{A} converge to $x \in X$. If \mathcal{V} is an ultrafilter finer than $E(\mathcal{F} \times \mathcal{A})$, then there exists an ultrafilter \mathcal{U} finer than $\mathcal{F} \times \mathcal{A}$ such that $E(\mathcal{U}) = \mathcal{V}$. Since $\pi_1(\mathcal{U}) \downarrow f$ and $\pi_2(\mathcal{U}) \downarrow x$, it follows that $\mathcal{V} = E(\mathcal{U}) = E(\pi_1(\mathcal{U}) \times \pi_2(\mathcal{U})) \downarrow f(x)$; thus $E(\mathcal{F} \times \mathcal{A}) \downarrow f(x)$, and therefore $\mathcal{F} \downarrow f$.

[Sufficiency]. By Proposition 2.58, it follows that Y is homeomorphic to a subspace of $\mathcal{C}(X, Y)$, which by Proposition 2.78 must be pseudotopological. \square

Theorem 2.81. *If $\mathcal{C}(X, Y)$ is a pseudotopological space and $f, g \in \mathcal{C}(X, Y)$, then $[f] \downarrow g$ if and only if $f(\mathcal{U}) \downarrow g(x)$ whenever \mathcal{U} is an ultrafilter that converges to x in X .*

Proof. [Necessity]. Suppose that \mathcal{A} converges to x in X . If \mathcal{V} is an ultrafilter finer than $[f] \cdot \mathcal{A}$, then $\mathcal{V} \supseteq f(\mathcal{A})$, and so there exists an ultrafilter \mathcal{U} finer than \mathcal{A} such that $f(\mathcal{U}) = \mathcal{V}$; thus by hypothesis $\mathcal{V} \downarrow g(x)$. Since Y is pseudotopological, it follows that $[f] \cdot \mathcal{A} \downarrow g(x)$. Therefore $[f] \downarrow g$.

[Sufficiency]. If \mathcal{U} is an ultrafilter that converges to x in X , then by hypothesis $[f] \cdot \mathcal{U} \downarrow g(x)$. Since $f(\mathcal{U}) \subseteq [f] \cdot \mathcal{U}$ and $f(\mathcal{U})$ is an ultrafilter, it follows that $f(\mathcal{U}) = [f] \cdot \mathcal{U}$; therefore $f(\mathcal{U}) \downarrow g(x)$. \square

2.5.5 Discrete and Indiscrete Convergence Structures

Definition 2.82. Let X be a convergence space. The *discrete convergence structure* on X is defined by the equivalence $\mathcal{F} \downarrow x$ if and only if $\mathcal{F} = [x]$; the *indiscrete convergence structure* on X is defined by $\mathcal{F} \downarrow x$.

The discrete convergence spaces are the discrete topological spaces; likewise, the indiscrete convergence spaces are the indiscrete topological spaces. Although these spaces have a simple, virtually trivial, structure, they arise frequently in the sequel.²⁸ Thus, we conclude this chapter with a survey of their properties.

Proposition 2.83. *Let X be a discrete space.*

1. *If $x \in X$, then $\mathcal{N}_x = [x]$.*
2. *The space X is topological.*
3. *Every subset of X is both open and closed.*
4. *If Y is a convergence space, then $f : X \rightarrow Y$ is continuous.*
5. *Any convergence space homeomorphic to X is also discrete.*

Proof.

1. The only filter that converges to x is $[x]$; thus $\mathcal{N}_x = [x]$.
2. Let X be a discrete convergence space. If \mathcal{T} is the discrete topology on X , then $\mathcal{U}_x = [x]$ for every $x \in X$. By hypothesis $\mathcal{F} \downarrow x$ if and only if \mathcal{U}_x . Therefore X is topological.
3. Let S be a subset of X . If $x \in S$, then $S \in [x] = \mathcal{N}_x$. Since S is a neighborhood of each of its points, it is open. By a similar argument $X - S$ is open, and so S is closed.
4. If \mathcal{F} converges to x in X , then $f(\mathcal{F}) = [f(x)]$, which converges to $f(x)$.
5. Let Y be a convergence space and let $f : Y \rightarrow X$ be a homeomorphism. If $\mathcal{F} \downarrow y$, then $f(\mathcal{F}) \downarrow f(y)$, and so $f(\mathcal{F}) = [f(y)]$, which implies that there exists $F \in \mathcal{F}$ such that $f(F) = \{f(y)\} = f(\{y\})$. Since f is a bijection, it follows that $F = \{y\}$; thus $\mathcal{F} = [y]$.

□

²⁸See Chapters 4 and 5.

Proposition 2.84. *Let X be an indiscrete space.*

1. *If $x \in X$, then $\mathcal{N}_x = [X]$.*
2. *The space X is topological*
3. *Every nonempty proper subset of X is neither open nor closed.*
4. *If Y is a convergence space, then $f : Y \rightarrow X$ is continuous.*
5. *Any convergence space homeomorphic to X is also indiscrete.*

Proof.

1. If a proper subset S of X is a neighborhood of x , then $S \in [X - S]$ since $[X - S]$ converges to x , but $[X - S]$ does not contain $\emptyset = S \cap (X - S)$; thus S is not a neighborhood of x .
2. Let X be an indiscrete convergence space. If \mathcal{T} is the indiscrete topology on X , then $\mathcal{U}_x = [X]$ for every $x \in X$. By hypothesis $\mathcal{F} \downarrow x$ if and only if \mathcal{U}_x . Thus X is topological.
3. Let S be a nonempty proper subset of X . If S is open and $x \in S$, then $S \in \mathcal{N}_x = [X]$, and so $X = S \subset X$, which is absurd; thus S is not open. By a similar argument $X - S$ is not open, and so S is not closed.
4. If \mathcal{F} converges to y in Y , then $f(\mathcal{F})$ converges to $f(y)$ in X .
5. Let Y be a convergence space and let $f : Y \rightarrow X$ be a homeomorphism. Suppose that \mathcal{F} is a filter on Y and $y \in Y$ but $\mathcal{F} \not\downarrow y$. Since f is a homeomorphism, it follows that $f(\mathcal{F}) \not\downarrow f(y)$; but this is false, since Y is indiscrete. Therefore $\mathcal{F} \downarrow y$.

□

If a space is initial with respect to a family of functions at least one of which is an injection into a discrete space, then it is also discrete; in particular, subspaces of a discrete space are

also discrete, although products need not be discrete. If a space is initial with respect to a family of functions into a collection of indiscrete spaces, then it is also indiscrete; in particular, subspaces and products of indiscrete spaces are indiscrete.

Proposition 2.85. *Let $(X_i)_{i \in I}$ be a collection of convergence spaces, let $(f_i : X \rightarrow X_i)_{i \in I}$ be a family of functions, and let X have the initial convergence structure with respect to $(f_i)_{i \in I}$.*

1. *If there exists $i \in I$ such that X_i is discrete and f_i is injective, then X is discrete.*
2. *If X_i is indiscrete for each $i \in I$, then X is indiscrete.*

Proof.

1. If $\mathcal{F} \downarrow x$, then $f_i(\mathcal{F}) \downarrow f_i(x)$, which implies that $f_i(\mathcal{F}) = [f_i(x)]$, and so there exists $F \in \mathcal{F}$ such that $f_i(F) = \{f_i(x)\}$; since f_i is injective, it follows that $F = \{x\}$; thus $\mathcal{F} = [x]$.
2. If \mathcal{F} is a filter on X and $x \in X$, then $f_i(\mathcal{F}) \downarrow f_i(x)$. Since this argument holds for each $i \in I$, it follows that $\mathcal{F} \downarrow x$.

□

Finally, we note that any continuous convergence structure with an indiscrete codomain is also indiscrete.

Proposition 2.86. *Let X be a convergence space.*

1. *If X is finite and Y is a finite discrete convergence space, then $\mathcal{C}(X, Y)$ is discrete.*
2. *If Y is an indiscrete convergence space, then $\mathcal{C}(X, Y)$ is indiscrete.*

Proof.

1. Let $f \in \mathcal{C}(X, Y)$. If $[f] \downarrow g$ in $\mathcal{C}(X, Y)$, then $[f(x)] \downarrow g(x)$ for each $x \in X$. Since Y is discrete, it follows that $f(x) = g(x)$ for each $x \in X$; thus $f = g$.

2. If \mathcal{F} converges to f in $\mathcal{C}(X, Y)$ and \mathcal{A} converges to a in X , then since Y is indiscrete $\mathcal{F} \cdot \mathcal{A}$ converges to $f(a)$ in Y ; thus \mathcal{F} converges to f in $\mathcal{C}(X, Y)$.

□

2.6 Compactness and Separation

We now attend to two concepts familiar from topology: compactness and separation. When extended to convergence spaces, compactness positively constrains convergence structures: every ultrafilter on a compact space must converge. On the other hand, separation negatively constrains convergence structures: depending on the degree of separation²⁹ convergence is restricted.

2.6.1 Compact Spaces

Definition 2.87. A convergence space X is *compact* if and only if every ultrafilter on X converges.

Every finite convergence space is compact: the only ultrafilters on a finite space are the point filters, each of which converges to at least one point.

Example 2.88. The set of non-negative integers \mathbb{N} equipped with the pretopology defined in Section 4.5 is compact.

The product convergence space induced by a collection of compact spaces is also compact; other initial convergence structures, including the subspace convergence structure, need not be compact.

Proposition 2.89. *If $(X_i)_{i \in I}$ is a collection of compact spaces, then $\prod_{i \in I} X_i$ is compact.*

²⁹For example, Kolmogorov (T_0), Fréchet (T_1), or Hausdorff (T_2). See Definitions 2.95, 2.109, and 2.119.

Proof. If \mathcal{U} is an ultrafilter on X , then $f_i(\mathcal{U})$ is an ultrafilter on X_i and so converges to some $x_i \in X_i$. Let x be the point in X such that $\pi_i(x) = x_i$ for each $i \in I$. Since $\mathcal{U} \downarrow x$, we conclude that X is compact. \square

Compactness is preserved by homeomorphisms. In fact, for any continuous surjection, if its domain is compact, then its codomain is also compact.

Proposition 2.90. *If X is a compact space, Y is a convergence space, and $f : X \rightarrow Y$ is a continuous surjection, then Y is compact.*

Proof. Let \mathcal{U} be an ultrafilter on X and let \mathcal{V} be an ultrafilter finer than $f^{-1}(\mathcal{U})$. Since f is a surjection, it follows that $\mathcal{U} = f(f^{-1}(\mathcal{U})) \subseteq f(\mathcal{V})$, and so $f(\mathcal{V}) = \mathcal{U}$; since X is compact, it follows that \mathcal{V} converges, and so $\mathcal{U} = f(\mathcal{V})$ converges. Therefore Y is compact. \square

Definition 2.87, of course, coincides with the usual definition when the convergence space is topological.

Proposition 2.91. *A topological convergence space X is compact if and only if every collection of open sets the union of which is X has a finite subcollection the union of which is also X .*

Proof. [Necessity]. We proceed by contraposition. If X is not compact, then there exists an ultrafilter \mathcal{U} that does not converge. Thus $U_x \notin \mathcal{U}$ for every $x \in X$, which implies that there exists a neighborhood U_x of x not included in \mathcal{U} . Since U_x is a neighborhood of x , it includes an open set V_x that contains x . Thus, for every $x \in X$, there exists an open set V_x that contains x but does not belong to \mathcal{U} , and so $X - V_x$ belongs to \mathcal{U} . Since $X = \bigcup_{x \in X} V_x$, it follows that there exists a finite subset S of X such that $X = \bigcup_{s \in S} V_s$. Thus $\emptyset = X - \bigcup_{s \in S} V_s = \bigcap_{s \in S} (X - V_s) \in \mathcal{U}$, which is absurd. Therefore X is compact.

[Sufficiency]. To the contrary, suppose that $X = \bigcup_{\alpha \in A} U_\alpha$ in which U_α is open for each $\alpha \in A$, but $X \supset \bigcup_{i \in I} U_i$ for every finite subset I of A . Thus, if I is a finite subset of A , then the set $A_I = X - \bigcup_{i \in I} U_i$ is nonempty, and so the set $\mathcal{F} = \{B : B \supseteq A_I \wedge I \text{ is a finite subset of } A\}$

is a filter. Let \mathcal{U} be an ultrafilter finer than \mathcal{F} . Since X is compact, it follows that \mathcal{U} converges to some point $x \in X$. By hypothesis, there exists an $\alpha \in A$ such that $x \in U_\alpha$. Since U_α is open, it must belong to \mathcal{U} ; but $X - U_\alpha$ belongs to \mathcal{F} and so also belongs to \mathcal{U} . Therefore, there exists a finite subset I of A such that $X = \bigcup_{i \in I} U_i$. \square

2.6.2 Kolmogorov Spaces

Two points of a topological space are *topologically distinguishable* if and only if there exists an open set containing exactly one of the sets. Topologists often require that every two points of a topological space are topologically distinguishable.³⁰ It is surprising then that we have not seen in the literature an attempt to extend distinguishability of points to convergence spaces.

Definition 2.92. Two points of a convergence space are *distinguishable* if and only if there exists a filter converging to exactly one of the points.

If there exists an open set containing exactly one of two points, then the two points are distinguishable; distinguishable points, however, might not have an open set containing exactly one of them.

Proposition 2.93. *Let X be a convergence space and let $x, y \in X$. If $\mathcal{N}_x \not\subseteq [x, y]$ or $\mathcal{N}_y \not\subseteq [x, y]$, then x and y are distinguishable.*

Proof. To the contrary, suppose that x and y are not distinguishable. Thus $\mathcal{F} \downarrow x$ if and only if $\mathcal{F} \downarrow y$; in particular $[y] \downarrow x$. If $[x, y] \not\subseteq \mathcal{N}_x$, then there exists $N \in \mathcal{N}_x$ that does not include $\{x, y\}$; but $N \in [y]$, and so $\{x, y\} \subseteq N$. Likewise $[x, y] \not\subseteq \mathcal{N}_x$ also implies a contradiction. Therefore x and y must be distinguishable. \square

If X is pretopological, then the converse of Proposition 2.93 also holds. Next, we verify that Definition 2.92 correctly generalizes topological distinguishability.

³⁰That is, they require the space to be *Kolmogorov* (T_0). Some even require the space to be Fréchet (T_1) or Hausdorff (T_2). See Munkres [30] for further discussion.

Theorem 2.94. *Two points x and y of a topological convergence space X are distinguishable if and only if there exists a open set containing exactly one of x or y .*

Proof. [Necessity]. By hypothesis, there exists an open set containing exactly one of x or y . If U contains x but not y , then $U \in \mathcal{N}_x - \mathcal{N}_y$, which implies that $\mathcal{N}_x \not\subseteq \mathcal{N}_y$, and so \mathcal{N}_y does not converge to x ; likewise, if U contains y but not x , then \mathcal{N}_x does not converge to y . Therefore, we conclude that x and y are distinguishable.

[Sufficiency]. We proceed by contraposition, that is, we suppose that an open set contains x if and only if it contains y . Thus every neighborhood of y is a neighborhood of x , and so $\mathcal{N}_y = \mathcal{N}_x$, from which it follows that a filter converges to x if and only if it converges to y . \square

Definition 2.95. A convergence space X is *Kolmogorov* if and only if each pair of distinct points in X are distinguishable.

Example 2.96. Each pair of points in the convergence space Y of Example 2.43 is distinguishable; thus Y is Kolmogorov. No pair of points in Y has an open set containing exactly one point of the pair.

Example 2.97. Every discrete space is Kolmogorov; every indiscrete space is not Kolmogorov. To see a non-trivial non-Kolmogorov space, consider the convergence space on $U = \{a, b, c\}$ defined by the equivalences

$$\mathcal{F} \downarrow a \text{ if and only if } U \in \mathcal{F}$$

$$\mathcal{F} \downarrow b \text{ if and only if } U \in \mathcal{F}$$

$$\mathcal{F} \downarrow c \text{ if and only if } \{a, c\} \in \mathcal{F}$$

Although a and b are both distinguishable from c , they are not distinguishable from each other; thus U is not Kolmogorov.³¹

³¹It is also not topological. See Chapter 6.

Next, we observe that Kolmogorov spaces are preserved by homeomorphisms.

Proposition 2.98. *Any convergence space homeomorphic to a Kolmogorov space is also Kolmogorov.*

Proof. Let X be a convergence space, let Y be a Kolmogorov space, let $f : X \rightarrow Y$ be a homeomorphism, and let a and b be distinct points of X . Since f is injective, the points $f(a)$ and $f(b)$ are distinct. By hypothesis, there exists a filter \mathcal{F} that converges to exactly one of $f(a)$ and $f(b)$. If \mathcal{F} converges to $f(a)$ but not to $f(b)$, then by the continuity of f^{-1} it follows that $f^{-1}(\mathcal{F})$ converges to a but not to b ; likewise, if \mathcal{F} converges to $f(b)$, then $f^{-1}(\mathcal{F})$ converges to b but not to a . Therefore X is Kolmogorov. \square

Products spaces are Kolmogorov if and only if each factor space is Kolmogorov; subspaces of Kolmogorov spaces need not be Kolmogorov.³²

Proposition 2.99. *If $(X_i)_{i \in I}$ is a collection of convergence spaces, then $\prod_{i \in I} X_i$ is Kolmogorov if and only if each X_i is Kolmogorov.*

Proof. [Necessity]. If x and y are distinct points of $\prod_{i \in I} X_i$, then there exists $j \in I$ such that $\pi_j(x) \neq \pi_j(y)$. Since X_j is Kolmogorov, it follows that there exists a filter \mathcal{F} converging to exactly one of $\pi_j(x)$ and $\pi_j(y)$. Suppose that \mathcal{F} converges to $\pi_j(x)$. Define the filter \mathcal{G} on $\prod_{i \in I} X_i$ by $\pi_i(\mathcal{G}) = \mathcal{F}$ if $i = j$ and $\pi_i(\mathcal{G}) = [\pi_i(x)]$ otherwise. Since $\pi_i(\mathcal{G}) \downarrow \pi_i(x)$ for each $i \in I$, it follows that $\mathcal{G} \downarrow x$; since $\pi_j(\mathcal{G}) \not\downarrow \pi_j(y)$, it follows that $\mathcal{G} \not\downarrow y$. Likewise, if \mathcal{F} converges to $\pi_j(y)$, then there exists a filter that converges to y but not to x . Therefore $\prod_{i \in I} X_i$ is Kolmogorov.

[Sufficiency]. To the contrary, if there exists $j \in I$ such that X_j is not Kolmogorov, then there exist distinct $\pi_j(x)$ and $\pi_j(y)$ in X_j such that $\mathcal{F} \downarrow \pi_j(x)$ if and only if $\mathcal{F} \downarrow \pi_j(y)$. Consider the point $z \in \prod_{i \in I} X_i$ defined by $\pi_i(z) = \pi_i(x)$ whenever $i \neq j$ and $\pi_j(z) = \pi_j(y)$.

³²See Example 3.2.

It follows that $x \neq z$, but

$$\mathcal{G} \downarrow z \Leftrightarrow (\forall i \in I)(\pi_i(\mathcal{G}) \downarrow \pi_i(z)) \Leftrightarrow (\forall i \in I)(\pi_i(\mathcal{G}) \downarrow \pi_i(x)) \Leftrightarrow \mathcal{G} \downarrow x,$$

in contradiction to the hypothesis that $\prod_{i \in I} X_i$ is Kolmogorov. Therefore each X_i is Kolmogorov. \square

Next, we verify that Definition 2.95 correctly generalizes Kolmogorov spaces from topology.

Theorem 2.100. *A topological convergence space X is Kolmogorov if and only if for every $x, y \in X$, there exists an open set containing exactly one of x or y .*

Proof. The desired result follows immediately from Theorem 2.94. \square

To show that a pretopological space is Kolmogorov, it is both necessary and sufficient to show that distinct points have distinct neighborhood filters; likewise, to show that a pseudotopological space is Kolmogorov, it is both necessary and sufficient to show that for each pair of distinct points, there exists an ultrafilter that converges to exactly one point of the pair.

Proposition 2.101. *A pretopological space X is Kolmogorov if and only if $\mathcal{N}_x \neq \mathcal{N}_y$ whenever x and y are distinct points of X .*

Proof. [Necessity]. We proceed by contraposition. If X is not Kolmogorov, then there exist distinct x and y in X such that $\mathcal{F} \downarrow x$ if and only $\mathcal{F} \downarrow y$. Since $\mathcal{N}_x \downarrow x$, it follows that $\mathcal{N}_x \downarrow y$, and so $\mathcal{N}_x \supseteq \mathcal{N}_y$; since $\mathcal{N}_y \downarrow y$, it follows that $\mathcal{N}_y \downarrow x$, and so $\mathcal{N}_y \supseteq \mathcal{N}_x$. Thus $\mathcal{N}_x = \mathcal{N}_y$.

[Sufficiency]. We proceed by contraposition. By hypothesis, there exists distinct points x and y of X such that $\mathcal{N}_x = \mathcal{N}_y$. If \mathcal{F} is a filter on X , then

$$\mathcal{F} \downarrow x \Leftrightarrow \mathcal{F} \supseteq \mathcal{N}_x \Leftrightarrow \mathcal{F} \supseteq \mathcal{N}_y \Leftrightarrow \mathcal{F} \downarrow y.$$

Thus x and y are not distinguishable, and therefore X is not Kolmogorov. \square

Proposition 2.102. *A pseudotopological space X is Kolmogorov if and only if for every pair of distinct points x and y of X , there exists an ultrafilter \mathcal{U} that converges to exactly one of x and y .*

Proof. [Necessity]. By hypothesis x and y are distinguishable. Therefore X is Kolmogorov.

[Sufficiency]. Let x and y be distinct points of X . By hypothesis, there exists a filter \mathcal{F} that converges to exactly one of x and y . If $\mathcal{F} \downarrow x$ but $\mathcal{F} \not\downarrow y$, then there exists an ultrafilter \mathcal{U} finer than \mathcal{F} that does not converge to y ; since $\mathcal{F} \downarrow x$, it follows that $\mathcal{U} \downarrow x$. Likewise, if $\mathcal{F} \downarrow y$ but $\mathcal{F} \not\downarrow x$, then there exists an ultrafilter that converges to y but not to x . \square

There is an interesting relationship between spaces of automorphisms and Kolmogorov spaces: $\text{Aut}(X)$ is discrete if X is finite and Kolmogorov; conversely, a topological space X is Kolmogorov if $\text{Aut}(X)$ is discrete.

Theorem 2.103. *If X is a finite Kolmogorov space, then $\text{Aut}(X)$ is discrete.*

Proof. Suppose that $[f]$ converges to g in $\text{Aut}(X)$. Let $x \in X$. Since X is finite, it follows that $\mathcal{N}_x = [N]$ for some $N \subseteq X$. Since $[f] \cdot \mathcal{N}_x$ converges to $g(x)$, it follows that

$$[g(N)] = g([N]) = g(\mathcal{N}_x) = \mathcal{N}_{g(x)} \subseteq [f] \cdot \mathcal{N}_x = [f] \cdot [N] = [f(N)],$$

and so $f(N) \subseteq g(N)$. Since f and g are automorphisms and N is finite, it follows that $f(N) = g(N)$; thus $\mathcal{N}_{f(x)} = f(\mathcal{N}_x) = f([N]) = [f(N)] = [g(N)] = \mathcal{N}_{g(x)}$, and so in view of Proposition 2.101 we see that $f(x) = g(x)$. Because $f(x) = g(x)$ for each $x \in X$, it must be that $f = g$. Therefore $\text{Aut}(X)$ is discrete. \square

We cannot guarantee the conclusion of Theorem 2.103 if we relax its hypothesis that the space is finite: we shall later see that, in view of Proposition 4.2, the space of automorphisms on an infinite Kolmogorov space need not be discrete. In seeking to strengthen Theorem

2.103, it would be natural to conjecture that the space of automorphisms of a compact Kolmogorov space is discrete. Later, we shall see a counterexample that refutes his conjecture.³³ Next we state a necessary condition on X when $\text{Aut}(X)$ is discrete.

Theorem 2.104. *Let X be a convergence space. If $\text{Aut}(X)$ is discrete, then X is Kolmogorov or not topological.*

Proof. We proceed by contraposition. By hypothesis, there exist distinct points a and b that are indistinguishable. If $f \in \text{Aut}(X)$, then the function $g : X \rightarrow X$ defined by

$$g(x) = \begin{cases} f(b), & x = a; \\ f(a), & x = b; \\ f(x), & \text{otherwise,} \end{cases}$$

also belongs to $\text{Aut}(X)$. If $\mathcal{F} \downarrow a$, then $\mathcal{F} \downarrow b$, and so $[f] \cdot \mathcal{F} \downarrow f(b) = g(a)$; likewise, if $\mathcal{F} \downarrow b$, then $[f] \cdot \mathcal{F} \downarrow g(b)$. Thus $[f] \downarrow g$; therefore X is not discrete. \square

Two corollaries follow immediately from Theorems 2.103 and 2.104.

Corollary 2.105. *If X is a topological space, then X is Kolmogorov whenever $\text{Aut}(X)$ is discrete.*

Proof. If X is a topological space and $\text{Aut}(X)$ is discrete, then by Theorem 2.104 it follows that X is Kolmogorov. \square

Corollary 2.106. *If X is a finite topological space, then $\text{Aut}(X)$ is discrete if and only if X is Kolmogorov.*

Proof. [Necessity]. This follows immediately from Theorem 2.103.

[Sufficiency]. By hypothesis X is pretopological; thus by Theorem 2.103, it follows that X is Kolmogorov. \square

³³See Example 3.19.

The next two examples, along with Proposition 4.2, show that all three cases of Theorem 2.104 are achievable.

Example 2.107. Equip \mathbb{N} with the pseudotopological structure defined by the conditions:

1. If $n \neq 2$, then $[n] \downarrow 0$.
2. If $n \neq 0$, then $[n]$ converges to n and $n - 1$.
3. Free filters do not converge.

Since $[n]$ converges to n but not to $n + 1$, it follows that \mathbb{N} is Kolmogorov.

This space, however, is not topological; in fact, it is not pretopological. To see this, observe that $\mathcal{N}_0 = [\mathbb{N} - \{2\}]$. But $\mathbb{N} - \{2\}$ belongs to the cofinite filter \mathcal{F} ; thus $\mathcal{N}_0 \subseteq \mathcal{F}$. If \mathcal{U} is a free filter, then $\mathcal{U} \supseteq \mathcal{F} \supseteq \mathcal{N}_0$; thus, if $\mathcal{N}_0 \downarrow 0$, then $\mathcal{U} \downarrow 0$, contrary to construction. Since \mathcal{N}_0 does not converge to 0, it follows that \mathbb{N} is not pretopological.

Finally, we note that $\text{Aut}(\mathbb{N})$ is discrete. Suppose that $f \in \text{Aut}(\mathbb{N})$. If $f(n) = 0$ for some $n \neq 0$, then $f(n - 1) = 0$, in contradiction to the hypothesis that f is a bijection; thus $f(0) = 0$. Likewise, if $n > 0$, then $f(n) - 1 = f(n - 1)$; thus $f(n) = n$, and so the only element of $\text{Aut}(\mathbb{N})$ is the identity function, which implies that \mathbb{N} is rigid. Therefore $\text{Aut}(\mathbb{N})$ is discrete.

Example 2.108. Modify the pseudotopological space \mathbb{N} of Example 2.107 with the pseudotopological structure defined by the conditions:

1. If $n \neq 2$, then $[n]$ converges to 0 and 1.
2. If $n = 2$, then $[n] \downarrow 2$.
3. If $n > 2$, then $[n]$ converges to n and $n - 1$.
4. Free filters do not converge.

By construction, a filter converges to 0 if and only if it converges to 1; thus \mathbb{N} is not Kolmogorov. By arguments similar to those given in Example 2.107, we see that \mathbb{N} is not topological; but \mathbb{N} is rigid, and therefore $\text{Aut}(\mathbb{N})$ is discrete.

2.6.3 Fréchet Spaces

Definition 2.109. A convergence space X is *Fréchet* if and only if $[x] \downarrow y$ implies $x = y$ for all $x, y \in X$.

Example 2.110. Every discrete space is Fréchet. The Euclidean line is also Fréchet. To see this, suppose that $[x] \downarrow y$ for some $x, y \in \mathbb{R}$. Since \mathbb{R} is topological, it follows that $\mathcal{U}_y \subseteq [x]$. For each $\varepsilon > 0$, define the set $U_\varepsilon = \{r : |r - y| < \varepsilon\}$. Thus, for every $\varepsilon > 0$, there exists $V \in [x]$ such $U_\varepsilon \supseteq V$; in particular $|x - y| < \varepsilon$. Since this inequality holds for every $\varepsilon > 0$, it follows that $x = y$.

Like Kolmogorov spaces, Fréchet spaces are also preserved by homeomorphisms.

Proposition 2.111. *Any convergence space homeomorphic to a Fréchet space is also Fréchet.*

Proof. Let X be a convergence space, let Y be a Fréchet space, let $f : X \rightarrow Y$ be a homeomorphism, and let $[a]$ converge to b in X . Since f is continuous, it follows that $[f(a)] \downarrow f(b)$; thus $f(a) = f(b)$, which implies that $a = b$. Therefore X is Fréchet. \square

A continuous convergence structure is Fréchet if its codomain space is also Fréchet.

Proposition 2.112. *If X is a convergence space and Y is a Fréchet space, then $\mathcal{C}(X, Y)$ is Fréchet.*

Proof. If $[f]$ converges to g in $\mathcal{C}(X, Y)$, then $[f] \cdot [x] = [f(x)]$ converges to $g(x)$ for each $x \in X$; thus $f(x) = g(x)$ for each $x \in X$, and so $f = g$. \square

Initial structures of Fréchet spaces need not be Fréchet unless they are constructed with respect to a separating³⁴ family of functions. Thus, subspaces and products of Fréchet spaces are also Fréchet.

Proposition 2.113. *Let $(X_i)_{i \in I}$ be a collection of convergence spaces, let $(f_i : X \rightarrow X_i)_{i \in I}$ be a separating family of functions, and let X have the initial convergence structure with respect to $(f_i)_{i \in I}$. If each $(X_i)_{i \in I}$ is Fréchet, then X is Fréchet.*

Proof. If $[x] \downarrow y$, then $f_i([x]) = [f_i(x)] \downarrow f_i(y)$ for each $i \in I$, and so $f_i(x) = f_i(y)$ for each $i \in I$; thus $x = y$. □

Before we verify that Definition 2.109 correctly generalizes Fréchet spaces from topology, we provide an equivalent condition for a space to be Fréchet.

Proposition 2.114. *A convergence space X is Fréchet if and only if every finite subset of X is closed.*

Proof. [Necessity]. Let x and y belong to X . Suppose that $[x] \downarrow y$. Since $\{x\} \in [x]$ it follows that $y \in \text{cl}(\{x\})$, and so by hypothesis $y \in \{x\}$. Therefore $x = y$.

[Sufficiency]. Let F be a finite subset of X . We proceed by induction on $|F|$. If $|F| = 1$, then $F = \{x\}$ for some $x \in X$. If $y \in \text{cl}(\{x\})$, then there exists a filter \mathcal{F} converging to y that contains $\{x\}$. But then $\mathcal{F} = [x]$, and so by hypothesis $x = y$. Thus $\text{cl}(\{x\}) = \{x\}$. Suppose that the hypothesis holds for finite subsets of cardinality n . Suppose that $|F| = n + 1$. If $x \in F$, then $|F - \{x\}| = n$, and so by the induction hypothesis $\text{cl}(F) = \text{cl}((F - \{x\}) \cup \{x\}) = \text{cl}(F - \{x\}) \cup \text{cl}(\{x\}) = (F - \{x\}) \cup \{x\} = F$. □

Proposition 2.115. *A topological convergence space X is Fréchet if and only if every finite subset of X is closed.*

Proof. The desired result follows immediately from Proposition 2.114. □

³⁴A family F of functions with domain X is *separating* if and only if for every $x, y \in X$, there exists $f \in F$ such that $f(x) \neq f(y)$.

The Fréchet condition, often assumed tacitly in topology, severely restricts finite convergence spaces: finite Fréchet spaces must be discrete.

Proposition 2.116. *If X is a finite Fréchet space, then X is discrete.*

Proof. If $x \in X$, then $X - \{x\}$ is finite, and so by Proposition 2.114 it follows that $X - \{x\}$ is closed, which implies that $\{x\}$ is open; thus $\mathcal{N}_x = [x]$. Therefore X is discrete. \square

Fréchet spaces are Kolmogorov spaces; Kolmogorov spaces, however, need not be Fréchet.

Proposition 2.117. *If a convergence space is Fréchet, then it is Kolmogorov.*

Proof. Let X be a Fréchet space. If x and y are distinct points of X , then $[x]$ does not converge to y . Therefore X is Kolmogorov. \square

Example 2.118. If \mathbb{N} is equipped with the pseudotopological structure defined by the conditions:

1. If $m, n \in \mathbb{N}$ and $n \neq 0$, then $[m] \downarrow n$ if and only if $m = n$.
2. If $m \in \mathbb{N}$, then $[m] \downarrow 0$ for $m = 0$ and all but finitely many m .
3. Free filters do not converge.

then \mathbb{N} is Kolmogorov but neither pretopological, compact, nor Fréchet. To see that \mathbb{N} is not pretopological, let $N = \{n \in \mathbb{N} : [n] \not\downarrow 0\}$. By construction, it follows that $\mathcal{N}_0 = [\mathbb{N} - N]$. Since N is finite, it follows that $\mathbb{N} - N$ belongs to the Fréchet filter \mathcal{F} ; thus $\mathcal{N}_0 \subseteq \mathcal{F}$. If \mathcal{U} is a free filter, then $\mathcal{U} \supseteq \mathcal{F} \supseteq \mathcal{N}_0$, and so \mathcal{N}_0 does not converge to 0.

2.6.4 Hausdorff Spaces

Definition 2.119. A convergence space X is *Hausdorff* if and only if no filter on X converges to two different points.

No indiscrete space is Hausdorff, but every discrete space is Hausdorff. Of course, the standard examples of Hausdorff topological spaces, such as the Euclidean line, are also Hausdorff convergence spaces.³⁵

Example 2.120. Consider \mathbb{N} equipped with the pseudotopological structure defined by the conditions

1. If $m, n \in \mathbb{N}$, then $[m] \downarrow n$ if and only if $m = n$.
2. Free filters converge to 0 only.

Since each filter converges to at most one point, it follows that \mathbb{N} is Hausdorff; since each ultrafilter converges, it follows that \mathbb{N} is compact.

Furthermore, this space is pretopological. Observe that $\mathcal{N}_0 = [0] \cap \mathcal{F}$, in which \mathcal{F} is the Fréchet filter, and $\mathcal{N}_n = [n]$ for nonzero $n \in \mathbb{N}$. To see that $\mathcal{N}_0 \downarrow 0$, note that if \mathcal{U} is an ultrafilter finer than \mathcal{N}_0 , then \mathcal{U} must be a free filter or $[0]$; in either case $\mathcal{U} \downarrow 0$.³⁶

Example 2.121. If \mathbb{N} is equipped with the pseudotopological structure defined by the conditions:

1. If $m, n \in \mathbb{N}$, then $[m] \downarrow n$ if and only if $m = n$.
2. If \mathcal{F} is a free filter, then $\mathcal{F} \downarrow m$ for each $m \in \mathbb{N}$.

then \mathbb{N} is compact and Fréchet but not Hausdorff.

Like Kolmogorov and Fréchet spaces, Hausdorff spaces are also preserved by homeomorphisms.

Proposition 2.122. *Any convergence space homeomorphic to a Hausdorff space is also Hausdorff.*

³⁵See Proposition 2.125.

³⁶In fact, this space is also topological. An example of a non-topological Hausdorff space is $\mathcal{C}(\mathbb{Q}, \mathbb{R})$, in which \mathbb{Q} inherits the subspace topology from \mathbb{R} . See Binz [6].

Proof. Let X be a convergence space, let Y be a Hausdorff space, let $f : X \rightarrow Y$ be a homeomorphism, and let \mathcal{F} converge to both a and b in X . Since f is continuous, it follows that $f(\mathcal{F})$ converges to both $f(a)$ and $f(b)$; thus $f(a) = f(b)$, which implies that $a = b$. Therefore X is Hausdorff. \square

A continuous convergence structure is Hausdorff if its codomain space is also Hausdorff.

Proposition 2.123. *If X is a convergence space and Y is a Hausdorff space, then $\mathcal{C}(X, Y)$ is Hausdorff.*

Proof. If \mathcal{F} converges to both f and g in $\mathcal{C}(X, Y)$, then $\mathcal{F} \cdot [x]$ converges to both $f(x)$ and $g(x)$ for each $x \in X$; thus $f(x) = g(x)$ for each $x \in X$, and so $f = g$. \square

As with initial structures of Fréchet spaces, initial structures of Hausdorff spaces need not be Hausdorff unless they are constructed with respect to a separating family of functions. Thus, subspaces and products of Hausdorff spaces are also Hausdorff.

Proposition 2.124. *Let $(X_i)_{i \in I}$ be a collection of convergence spaces, let $(f_i : X \rightarrow X_i)_{i \in I}$ be a separating family of functions, and let X have the initial convergence structure with respect to $(f_i)_{i \in I}$. If each $(X_i)_{i \in I}$ is Hausdorff, then X is Hausdorff.*

Proof. If \mathcal{F} converges to both x and y in X , then $f_i(\mathcal{F})$ converges to both $f_i(x)$ and $f_i(\mathcal{F}) \downarrow f_i(y)$ for each $i \in I$. Since each X_i is Hausdorff, it follows that $f_i(x) = f_i(y)$ for each $i \in I$; since $(f_i : X \rightarrow X_i)_{i \in I}$ is separating, it follows that $x = y$. Therefore X is Hausdorff. \square

Next, we verify that Definition 2.119 correctly generalizes Hausdorff spaces from topology.

Proposition 2.125. *A pretopological convergence space X is Hausdorff if and only if for every two points x and y of X there exists a neighborhood U of x and a neighborhood V of y such that $U \cap V = \emptyset$.*

Proof. [Necessity]. If X is not Hausdorff, then there exists a filter \mathcal{F} that converges to both x and y , which implies that \mathcal{F} contains both U and V , and so \mathcal{F} contains $U \cap V$; no filter, however, can contain the empty set. Therefore X is Hausdorff.

[Sufficiency]. If every neighborhood of x intersects every neighborhood of y , then the set $\mathcal{F} = \{U \cap V : U \in \mathcal{N}_x \wedge V \in \mathcal{N}_y\}$ is a filter finer than both \mathcal{N}_x and \mathcal{N}_y , which implies that \mathcal{F} converges to both x and y , in contradiction to the hypothesis that X is Hausdorff. \square

To show that a pseudotopological space is Hausdorff, it suffices to show that no ultrafilter converges to two distinct points.

Proposition 2.126. *A pseudotopological space X is Hausdorff if and only if no ultrafilter converges to two distinct points of X .*

Proof. [Necessity]. If \mathcal{F} converges to two distinct points x and y in X , then every ultrafilter finer than \mathcal{F} converges to both x and y , in contradiction to the hypothesis.

[Sufficiency]. This follows immediately by Definition 2.119. \square

Any continuous bijection from a compact pseudotopological space into a Hausdorff space is a homeomorphism; consequently, continuous bijections between compact pseudotopological Hausdorff spaces are homeomorphisms.

Proposition 2.127. *If X is a compact pseudotopological space, Y is a Hausdorff space, and $f : X \rightarrow Y$ is a continuous bijection, then f is a homeomorphism.*

Proof. Let \mathcal{F} converge to $f(x)$ in Y and let \mathcal{U} be an ultrafilter finer than $f^{-1}(\mathcal{F})$. Since X is compact, it follows that $\mathcal{U} \downarrow w$ for some $w \in X$; since f is continuous, it follows that $f(\mathcal{U}) \downarrow f(w)$. By hypothesis $f(\mathcal{U}) \downarrow f(x)$; but Y is Hausdorff, so $f(x) = f(w)$, which implies that $x = w$. Therefore $f^{-1}(\mathcal{F}) \downarrow x$. \square

Hausdorff spaces are Fréchet, and hence Kolmogorov; neither Kolmogorov nor Fréchet spaces need be Hausdorff.

Proposition 2.128. *If a convergence space is Hausdorff, then it is Fréchet.*

Proof. Let x and y belong to a Hausdorff space X . If $[x] \downarrow y$, then $x = y$, otherwise $[x]$ converges to x and y , in contradiction to the condition that X is Hausdorff. Therefore X is Fréchet. \square

We conclude this chapter with the observation that finite Hausdorff spaces must be discrete since they are finite Fréchet spaces.

Proposition 2.129. *If X is a finite Hausdorff space, then X is discrete.*

Proof. By hypothesis X is a finite Fréchet space; therefore X is discrete. □

Chapter 3

Reflexive Digraphs as Convergence Spaces

The category of convergence spaces, **CONV**, which is Cartesian closed, embeds as full subcategories both the category of topological spaces, **TOP**, and the category of reflexive digraphs, **ReRe**. The continuous functions in **TOP** are exactly the continuous functions in the embedded image of **TOP** in **CONV**; the continuous functions in the embedded image of **ReRe** are exactly the directed graph homomorphisms. Thus **CONV** unifies the two notions of directed graph homomorphism and continuous function. The essence of both notions is preservation of filter convergence, that is, a function f is continuous at a point x if and only if the image under f of each filter converging to x converges to $f(x)$. The Cartesian closure of **CONV** enables construction of hybrid spaces with discrete and continuous convergence structure.

A reflexive digraph's edge set induces a convergence structure on its vertex set. This observation provides a basis for extending concepts such as continuity and differentiability¹ from continua to discrete structures. Since the reflexive closure of a Cayley graph is a reflexive digraph, we can speak coherently of continuity and differentiability of functions on

¹See Chapter 4.

Cayley graphs; by identifying a group with one of its Cayley graphs, we then can speak of continuity and differentiability of functions between groups.

3.1 Reflexive Digraphs

Definition 3.1. Let (V, E) be a reflexive digraph. For each $v \in V$, the *graph neighborhood* of v is the set $\vec{v} = \{u \in V : (v, u) \in E\}$. The *reflexive digraph convergence structure* on V is defined by

$$\mathcal{F} \downarrow v \text{ if and only if } \vec{v} \in \mathcal{F}.$$

When no reasonable confusion is likely, we refer to a reflexive digraph (V, E) by V . Unless otherwise noted, we assume that all reflexive digraphs have the reflexive digraph convergence structure.

We can induce a convergence structure on the vertex set of a reflexive digraph in other ways. In [8], a filter \mathcal{F} converges to a vertex v of a reflexive digraph if and only if $\mathcal{F} = [u]$ for some vertex $u \in \vec{v}$; this structure requires the removal of the finite intersection property of convergence structures. In [31], a filter \mathcal{F} converges to a vertex v of a reflexive digraph if and only if there exists $u \in \bigcap_{F \in \mathcal{F}} F$ such that $v \in \vec{u}$; this structure reverses edges and prohibits non-principal filters from converging.

Example 3.2. Let P denote the pentacle, as shown in Figure 3.1, with the reflexive digraph convergence structure.² For each $p \in P$, the graph neighborhood of p is $\{q \in P : q \not\equiv p + 3 \pmod{5}\}$. Thus $\mathcal{F} \downarrow p$ if and only if $\mathcal{F} = [A]$ for some subset A of P that does not contain $p + 3 \pmod{5}$.

Note that $\text{Aut}(P) = \{(), (0\ 1\ 2\ 3\ 4), (0\ 2\ 4\ 1\ 3), (0\ 3\ 1\ 4\ 2), (0\ 4\ 3\ 2\ 1)\}$, which, as a group, is isomorphic to C_5 but, as a subspace of $\mathcal{C}(P, P)$, does not have the reflexive digraph convergence structure of C_5 , rather by Theorem 2.103 it is discrete.

²In the interest of clarity, we always omit reflexive loops from reflexive digraphs.




Figure 3.1: The Pentacle.

Also observe that the set $\{0, 1, 2\}$, as a subspace of P , is homeomorphic to the non-Kolmogorov space U of Example 2.97; thus U is not Kolmogorov although P is Kolmogorov. This shows that subspaces of Kolmogorov spaces need not be Kolmogorov.

As first observed in [8], continuity of functions on reflexive digraphs has an elegant formulation.

Proposition 3.3. *Let (V_1, E_1) and (V_2, E_2) be reflexive digraphs, and let V_1 and V_2 be the induced convergence spaces, respectively. A function $f : V_1 \rightarrow V_2$ is continuous if and only if $(f(v), f(u)) \in E_2$ whenever $(v, u) \in E_1$.*

Proof. [Necessity]. If $\mathcal{F} \downarrow v$, then $\vec{v} \in \mathcal{F}$. By hypothesis $f(\vec{v}) \subseteq \overrightarrow{f(v)}$. Since $f(\vec{v}) \in f(\mathcal{F})$, it follows that $\overrightarrow{f(v)} \in f(\mathcal{F})$. Therefore $f(\mathcal{F}) \downarrow f(v)$.

[Sufficiency]. If $(v, u) \in E_1$, then $u \in \vec{v}$, and so $\vec{v} \in [u]$, which implies that $[u] \downarrow v$. By hypothesis $[f(u)] \downarrow f(v)$; thus $\overrightarrow{f(v)} \in [f(u)]$, and so $f(u) \in \overrightarrow{f(v)}$. Therefore $(f(v), f(u)) \in E_2$. \square

In other words, a function is continuous on a reflexive digraph if and only if it is a graph homomorphism. Thus, the concepts of graph homomorphism and continuous function are, in fact, manifestations of the same concept.

In addition to Proposition 3.3, it is useful to have a formulation of local continuity.

Proposition 3.4. *Let (V_1, E_1) and (V_2, E_2) be reflexive digraphs, and let V_1 and V_2 be the induced convergence spaces, respectively. A function $f : V_1 \rightarrow V_2$ is continuous at $v \in V_1$ if and only if $f(\vec{v}) \subseteq \overrightarrow{f(v)}$.*

Proof. [Necessity]. If $\mathcal{F} \downarrow v$, then $\vec{v} \in \mathcal{F}$; thus $f(\vec{v}) \in f(\mathcal{F})$, and so by hypothesis $\overrightarrow{f(v)} \in f(\mathcal{F})$, which implies that $f(\mathcal{F}) \downarrow f(v)$. Therefore f is continuous at v .

[Sufficiency]. Since $[\vec{v}] \downarrow v$, by hypothesis it follows that $[f(\vec{v})] \downarrow f(v)$. Thus $\overrightarrow{f(v)} \in [f(\vec{v})]$, and therefore $f(\vec{v}) \subseteq \overrightarrow{f(v)}$. \square

In view of Proposition 3.4, we see that a function f on a reflexive digraph V is an automorphism if and only if $f(\vec{v}) = \overrightarrow{f(v)}$ for each $v \in V$.

If the edge relation of a reflexive digraph is a partial order, then the continuous functions on that reflexive digraph are precisely the nondecreasing monotone functions.

Proposition 3.5. *If (V, E) is a reflexive digraph and E is a partial order, then $f : V \rightarrow V$ is continuous if and only if $f(u) \leq f(v)$ whenever $u \leq v$, for each u and v in V .*

Proof. [Necessity]. If $w \in f(\vec{u})$, then there exists $v \in \vec{u}$ such that $f(v) = w$. By hypothesis $f(u) \leq f(v)$, which implies that $w = f(v) \in \overrightarrow{f(u)}$; thus $f(\vec{u}) \subseteq \overrightarrow{f(u)}$. Therefore $f \in \mathcal{C}(V, V)$.

[Sufficiency]. If $u \leq v$, then $v \in \vec{u}$, and so by hypothesis $f(v) \in f(\vec{u}) \subseteq \overrightarrow{f(u)}$; therefore $f(u) \leq f(v)$. \square

Next we determine the closure of a graph neighborhood: it is the set of all vertices the graph neighborhoods of which intersect the given graph neighborhood.

Proposition 3.6. *Let V be a reflexive digraph. If $u \in V$, then $cl(\vec{u}) = \{v \in V : \vec{u} \cap \vec{v} \neq \emptyset\}$.*

Proof. If $v \in cl(\vec{u})$, then there exists a filter \mathcal{F} that converges to v and contains \vec{u} ; thus $\vec{v} \in \mathcal{F}$, which implies that $\vec{u} \cap \vec{v} \neq \emptyset$. Conversely, if $\vec{u} \cap \vec{v} \neq \emptyset$ for some $v \in V$, then $[\vec{u} \cap \vec{v}]$ is a filter containing \vec{u} that converges to v ; thus $v \in cl(\vec{u})$. \square

Proposition 3.6 enables us to establish a necessary and sufficient condition for the graph neighborhood of a point to be closed.

Proposition 3.7. *Let V be a reflexive digraph. If $u \in V$, then \vec{u} is closed if and only if $\vec{u} \cap \vec{v} = \emptyset$ or $v \in \vec{u}$ for each $v \in V$.*

Proof. [Necessity]. If $v \in \text{cl}(\overrightarrow{u})$, then $\overrightarrow{u} \cap \overrightarrow{v} \neq \emptyset$, and so by hypothesis $v \in \overrightarrow{u}$; thus $\text{cl}(\overrightarrow{u}) = \overrightarrow{u}$.

[Sufficiency]. If $v \notin \overrightarrow{u}$, then by hypothesis $v \notin \text{cl}(\overrightarrow{u})$, and so $\overrightarrow{u} \cap \overrightarrow{v} = \emptyset$. \square

Now we attend to an equivalent condition for a graph neighborhood to be open: it must include the graph neighborhood of each of its members.

Proposition 3.8. *If V is a reflexive digraph, then U is an open subset of V if and only if $\overrightarrow{x} \subseteq U$ for each $x \in U$.*

Proof. [Necessity]. Let $x \in U$. If $\mathcal{F} \downarrow x$, then $\overrightarrow{x} \in \mathcal{F}$. By hypothesis $\overrightarrow{x} \subseteq U$, which implies that $U \in \mathcal{F}$. Thus U is an open subset of V .

[Sufficiency]. Let $x \in U$. If $y \in \overrightarrow{x}$, then $[y] \downarrow x$. By hypothesis $U \in [y]$; thus $y \in U$. Therefore $\overrightarrow{x} \subseteq U$. \square

The neighborhood filter of a vertex of a reflexive digraph is the principal filter generated by the graph neighborhood of that vertex.

Proposition 3.9. *If v belongs to a reflexive digraph V , then $\mathcal{N}_v = [\overrightarrow{v}]$.*

Proof. It suffices to show that $[\overrightarrow{v}] \subseteq \mathcal{N}_v$. If $N \in [\overrightarrow{v}]$, then $\overrightarrow{v} \subseteq N$. Since \overrightarrow{v} belongs to each filter converging to v , it follows that N belongs to each filter that converges to v ; thus $N \in \mathcal{N}_v$, as required. \square

Example 3.10. The standard topology on \mathbb{R} is not a reflexive digraph. To the contrary, suppose that the standard topology on \mathbb{R} is a reflexive digraph. By Proposition 3.9, it follows $\mathcal{N}_x = [\overrightarrow{x}]$ for every $x \in \mathbb{R}$. Thus $[\overrightarrow{x}] \subseteq \mathcal{N}_x$, which implies that there exists $\varepsilon > 0$ such that $(x - \varepsilon, x + \varepsilon) \subseteq \overrightarrow{x}$. Let $0 < \varepsilon' < \varepsilon$. Since $(x - \varepsilon', x + \varepsilon') \in \mathcal{N}_x \subseteq [\overrightarrow{x}]$, it follows that $(x - \varepsilon, x + \varepsilon) = \overrightarrow{x} \subseteq (x - \varepsilon', x + \varepsilon')$, which is absurd. Therefore, the standard topology on \mathbb{R} is not a reflexive digraph.

By Proposition 3.8, it follows that a reflexive digraph is topological if and only if it is transitive. Consequently, we have a canonical topology for preorders: if P is a preorder, then $\{U \subseteq P : (\forall v)(v \in U \rightarrow \vec{v} \subseteq U)\}$ is a topology on P .³

Proposition 3.11. *A reflexive digraph is topological if and only if it is transitive.*

Proof. Let V be a reflexive digraph.

[Necessity]. We must show that there exists a topology \mathcal{T} such that $\mathcal{F} \downarrow v$ if and only if $\mathcal{F} \subseteq \mathcal{U}_v$. Let $U \in \mathcal{T}$ if and only if $U \supseteq \vec{v}$ for every $v \in U$.

First, we verify that \mathcal{T} is a topology on V . Since $V \supseteq \vec{v}$ for every $v \in V$, it follows that $V \in \mathcal{T}$; likewise \emptyset has no elements, and so belongs to \mathcal{T} . Suppose that $U_\alpha \in \mathcal{T}$ for each α of some index set A . If $v \in \bigcup_{\alpha \in A} U_\alpha$, then $v \in U_\alpha$ for some $\alpha \in A$, and so $\vec{v} \subseteq U_\alpha \subseteq \bigcup_{\alpha \in A} U_\alpha$; thus $\bigcup_{\alpha \in A} U_\alpha \in \mathcal{T}$. Suppose that U and U' both belong to \mathcal{T} . If $v \in U \cap U'$, then v belongs to both U and U' , and so \vec{v} is a subset of both U and U' , which implies that $\vec{v} \subseteq U \cap U'$; thus $U \cap U' \in \mathcal{T}$.

Next, we show that $\mathcal{F} \downarrow v$ if and only if $\mathcal{F} \supseteq \mathcal{U}_v$. Suppose that $\mathcal{F} \supseteq \mathcal{U}_v$. Since V is transitive, it follows that \vec{v} is transitive, and so \vec{v} is open, which implies that $\vec{v} \supseteq \vec{x}$ for each $x \in \vec{v}$; thus $\vec{v} \in \mathcal{T}$, and in particular $\vec{v} \in \mathcal{U}_v \subseteq \mathcal{F}$, from which we conclude that $\mathcal{F} \downarrow v$. Conversely, suppose that $\mathcal{F} \downarrow v$. If $U \in \mathcal{U}_v$, then U includes a set $U' \in \mathcal{T}$ that contains v ; thus $\vec{v} \subseteq U' \subseteq U$. Since $\vec{v} \in \mathcal{F}$, it follows that for every $U \in \mathcal{U}_v$, there exists $F \in \mathcal{F}$, namely \vec{v} , such that $F \subseteq U$; thus $\mathcal{F} \supseteq \mathcal{U}_v$.

[Sufficiency]. Let \mathcal{T} be a topology that induces the convergence structure \downarrow . Suppose that $v \in \vec{u}$ and $w \in \vec{v}$. Since V is topological, it follows that $\vec{u} \in \mathcal{U}_u$, and so \vec{u} includes an open set U that contains u , which implies that $\vec{u} \subseteq U$, and so $U = \vec{u}$; thus \vec{u} is open, and so by Proposition 3.8 it follows that $\vec{v} \subseteq \vec{u}$; thus $w \in \vec{u}$. Therefore V is transitive. \square

Proposition 3.9 implies that each reflexive digraph is pretopological, and thus pseudotopological.

³Equivalently, if P is a preorder, then $\{\vec{p} : p \in P\}$ is a basis for a topology on P .

Proposition 3.12. *Every reflexive digraph is pretopological.*

Proof. Let V be a reflexive digraph. If $v \in V$, then by Proposition 3.9 it follows that $\vec{v} \in [\vec{v}] = \mathcal{N}_v$, and so $\mathcal{N}_v \downarrow v$. \square

Proposition 3.13. *Every reflexive digraph is pseudotopological.*

Proof. The desired result follows immediately from Propositions 2.79 and 3.12. \square

Next we give a partial characterization of $\mathcal{C}(X, Y)$ for reflexive digraphs X and Y .

Proposition 3.14. *If X and Y are reflexive digraphs and $f, g \in \mathcal{C}(X, Y)$, then $[f] \downarrow g$ if and only if $f(a) \in \overrightarrow{g(b)}$ whenever $a \in \vec{b}$ for each a and b in X .*

Proof. [Necessity]. If \mathcal{A} converges to x in X , then $\vec{x} \in \mathcal{A}$. By hypothesis, it follows that $f(\vec{x}) \subseteq \overrightarrow{g(x)}$. Thus $\overrightarrow{g(x)} \in [\overrightarrow{g(x)}] \subseteq [f(\vec{x})] \subseteq [f] \cdot [\vec{x}] \subseteq [f] \cdot \mathcal{A}$, and so $[f] \cdot \mathcal{A}$ converges to $g(x)$ in Y . Therefore $[f] \downarrow g$.

[Sufficiency]. If $a \in \vec{b}$, then $[a] \downarrow b$. By hypothesis, it follows that $[f(a)] = [f] \cdot [a]$ converges to $g(b)$; thus $f(a) \in \overrightarrow{g(b)}$. \square

Every finite convergence space can be represented by a reflexive digraph: if X is a finite convergence space, then for each $a \in X$, let \vec{a} be the set that generates \mathcal{N}_a . Subsequently, we treat, in the interest of clarity, all finite convergence spaces as reflexive digraphs.

The intuitive meaning of Proposition 3.14 in the case of functions between finite convergence space is that two continuous functions on are “close” if and only if their local behavior is approximately the same.

As with all finite convergence spaces, finite reflexive digraphs are compact; for finite reflexive digraphs to be compact it also suffices that the complement of some graph neighborhood is finite. Conversely, if an infinite reflexive digraph is compact, then one of its graph neighborhoods also must be infinite.

Proposition 3.15. *Let V be a reflexive digraph.*

1. If V is finite or $V - \vec{v}$ is finite for some $v \in V$, then V is compact.

2. If V is compact, then V is finite or \vec{v} is infinite for some $v \in V$

Proof.

1. If V is finite, then the only ultrafilters on V are point filters, which must converge. If V is infinite, then $V - \vec{v}$ is finite for some $v \in V$, which implies that \vec{v} belongs to the Fréchet filter, and so the Fréchet filter converges to v ; thus every free filter converges to v .


2. We proceed by contraposition: suppose that V is infinite but \vec{v} is finite for every $v \in V$. If \mathcal{U} is an ultrafilter that converges to v , then $\vec{v} \in \mathcal{U}$; but \vec{v} is finite, which implies that \mathcal{U} is a point filter. Thus free filters do not converge, and so V is not compact, in contradiction to the hypothesis. Therefore V is finite or \vec{v} is infinite for some $v \in V$.

□

Reflexive digraphs are rarely Hausdorff or even Fréchet; in fact, the only Fréchet, and hence Hausdorff, reflexive digraphs are discrete. On the other hand, reflexive digraphs are often Kolmogorov: for a reflexive digraph to be Kolmogorov, it is both necessary and sufficient that no two of its graph neighborhoods are identical.

Proposition 3.16. *A reflexive digraph V is Kolmogorov if and only if $\vec{u} \neq \vec{v}$ whenever u and v are distinct points of V .*

Proof. [Necessity]. If u and v are distinct points of V , then $\vec{u} \neq \vec{v}$. If $w \in \vec{u} - \vec{v}$, then $\vec{u} \in [w]$ but $\vec{v} \notin [w]$; thus $[w] \downarrow u$ but $[w] \not\downarrow v$. Likewise, if $w \in \vec{v} - \vec{u}$, then $[w] \downarrow v$ but $[w] \not\downarrow u$. In either case, there exists a filter that converges to exactly one of the points u and v . Therefore V is Kolmogorov.

Figure 3.2: A Partial Ordering on \mathbb{N} .

[Sufficiency]. We proceed by contraposition: suppose that there exist distinct points u and v of V such that $\vec{u} = \vec{v}$. If $\mathcal{F} \downarrow u$, then $\vec{v} = \vec{u} \in \mathcal{F}$, and so $\mathcal{F} \downarrow v$; likewise, if $\mathcal{F} \downarrow v$, then $\mathcal{F} \downarrow u$. Therefore V is not Kolmogorov. \square

Proposition 3.17. *A reflexive digraph is Fréchet if and only if it is discrete.*

Proof. Let V be a reflexive digraph.

[Necessity]. Since every discrete space is Fréchet, it follows that V is Fréchet.

[Sufficiency]. Let $v \in V$. If $u \in \vec{v}$, then $\vec{v} \in [u]$, and so $[u]$ converges to both u and v , and so by hypothesis $u = v$; thus $\vec{v} = \{v\}$, which implies that $\mathcal{N}_v = [v]$, and so $\mathcal{F} \downarrow v$ if and only if $\mathcal{F} = [v]$. Therefore V is discrete. \square

The next two examples preclude natural strengthenings of Theorem 2.103.

Example 3.18. Equip \mathbb{N} with a partial order, as represented by the reflexive digraph of Figure 3.2. Observe that \mathbb{N} is a topological Kolmogorov space in which each point has a finite neighborhood. Moreover, if $f : \mathbb{N} \rightarrow \mathbb{N}$ is a bijection, then $f \in \text{Aut}(X)$ if and only if $f(\infty) = \infty$; thus $\text{Aut}(X)$ is infinite.

Let $f \in \text{Aut}(\mathbb{N})$ and let

$$\mathcal{S} = \left\{ \bigcap_{i \in I} S_i : I \text{ is a finite subset of } \mathbb{N} \right\},$$

where $S_n = \{g \in \text{Aut}(\mathbb{N}) : g(n) = f(n)\}$. By Proposition 2.13, there exists a free filter \mathcal{U} on X that includes \mathcal{S} .

We claim that $\mathcal{U} \downarrow f$. It suffices to show that $\mathcal{U} \cdot [\vec{n}] \downarrow f(n)$ for each nonzero $n \in \mathbb{N}$. Since $S_n \in \mathcal{U}$, there exists $U \in \mathcal{U}$, namely S_n , such that $U \cdot \vec{n} = \overrightarrow{f(n)}$, which implies that




Figure 3.3: Another Partial Ordering on \mathbb{N} .

$\overrightarrow{f(n)} \in \mathcal{U} \cdot [\overrightarrow{n}]$. Therefore $\mathcal{U} \cdot [\overrightarrow{n}] \downarrow f(n)$, as desired.

Since there is a convergent nonprincipal ultrafilter on $\text{Aut}(\mathbb{N})$, it follows that $\text{Aut}(\mathbb{N})$ is not discrete. Therefore, the space of automorphisms on an infinite pretopological Kolmogorov space, in which each point has a finite neighborhood, need not be discrete.

Example 3.19. Now denote by \mathbb{N} the set of all non-negative integers together with an element ∞ that is strictly greater than every non-negative integer. Equip \mathbb{N} with the partial order depicted by the Hasse diagram of Figure 3.3. Observe that \mathbb{N} is a compact topological Kolmogorov space. By an argument exactly similar to that in Example 3.18, it follows that $\text{Aut}(\mathbb{N})$ is not discrete. Therefore, the space of automorphisms on a compact topological Kolmogorov space need not be discrete.

3.2 Cayley Graphs

Definition 3.20. Let Γ be a subset of a finite group G such that each element of G is a product of elements of Γ and no element of Γ is redundant.⁴ We call Γ a *generating set* for G and each element of Γ a *generator* of G . The *Cayley graph* for G generated by Γ is the reflexive digraph C such that the vertex set of C is G and the edge set of C is $\{(g, h) : g\gamma = h \wedge (\gamma = e \vee \gamma \in \Gamma)\}$.⁵

Under Definition 3.20, every Cayley graph is a reflexive digraph, and so the results of the previous section apply to Cayley graphs. The additional structure of Cayley graphs,

⁴An element $\gamma \in \Gamma$ is *redundant* if and only if it is a product of elements in $\Gamma - \{\gamma\}$.

⁵We assume a familiarity with the most basic facts of the theories of graphs and groups. For graph theory, see Merris [28]; group theory, Rotman [36].




Figure 3.4: A Cayley Graph of the Cyclic Group C_3 .

however, imply additional properties, which in this section we will investigate. First, we observe that left-multiplication is edge-preserving, and thus bicontinuous.

Proposition 3.21. *If v belongs to a Cayley graph C , then $\lambda x.vx$ is an automorphism on C .*

Proof. It is clear that $\lambda x.vx$ is a bijection. To see that $\lambda x.vx$ is bicontinuous, let $u \in C$. If $y \in (\lambda x.vx)(\vec{u})$, then there exists $w \in \vec{u}$ such that $vw = y$. Since $w \in \vec{u}$, there exists a generator g such that $ug = w$, which implies that $y = vug$, and so $y \in \vec{vu} = \overrightarrow{(\lambda x.vx)(u)}$. Since $(\lambda x.vx)(\vec{u}) \subseteq \overrightarrow{(\lambda x.vx)(u)}$, it follows that $\lambda x.vx$ is continuous at u . An exactly similar argument shows that $(\lambda x.vx)^{-1}$ is continuous at u . Because both $\lambda x.vx$ and its inverse are continuous at every point in V , we conclude that $\lambda x.vx$ is a bicontinuous, and therefore an automorphism. \square

In view of Proposition 3.21, we see that for any Cayley graph C , the set $\Lambda_C = \{\lambda x.vx : v \in C\}$ is a subset of $\mathcal{C}(C, C)$.⁶ Unless otherwise noted, we assume that Λ_C inherits the subspace structure from $\mathcal{C}(C, C)$.

Example 3.22. The Cayley graph for the cyclic group C_3 with presentation $\langle p : p^3 \rangle$, as depicted in see Figure 3.4, is homeomorphic to the convergence space Y of Example 2.43. To see this, identify e with 0, p with 1, and p^2 with 2. The continuous functions of this Cayley graph are the identity function, the constant functions, and the left-multiplications $\lambda x.px$ and $\lambda x.p^2x$; of these functions, only the identity function and the left-multiplications are automorphisms.

⁶In general, there might exist automorphisms on C that are not members of Λ_C .

Since no Cayley graph except C_1 is discrete, it follows that no Cayley graph is Fréchet or Hausdorff. Nearly every Cayley graph, however, is Kolmogorov: the only exception is the Cayley graph for C_2 .

Theorem 3.23. *If C is a non-Kolmogorov Cayley graph for a group G generated by Γ such that $|\Gamma| < 2$, then G is C_2 .*

Proof. By hypothesis, there exist distinct points a and b of C such that $\vec{a} = \vec{b}$. If $\Gamma = \{\gamma\}$ for some $\gamma \in G$ distinct from the identity element, then $a\gamma = b$ and $b\gamma = a$, which implies that $\gamma^2 = a^{-1}bb^{-1}a = e$; thus G is the cyclic group C_2 . \square

Theorem 3.24. *If C is a Cayley graph for a group G generated by Γ such that $|\Gamma| \geq 2$, then C is Kolmogorov.*

Proof. To the contrary, if C is not Kolmogorov, then there exist distinct points a and b of C such that $\vec{a} = \vec{b}$. Thus there exist generators γ_1 , γ_2 , and γ_3 such that γ_2 is distinct from γ_1 and γ_3 , $a\gamma_1 = b$, and $a\gamma_2 = b\gamma_3$, which implies that $\gamma_1\gamma_3 = a^{-1}a\gamma_1\gamma_3 = a^{-1}b\gamma_3 = a^{-1}a\gamma_2 = \gamma_2$, in contradiction to the condition that no element of Γ is redundant. Therefore C is Kolmogorov. \square

Chapter 4

Differential Calculus on Convergence Spaces

In [8], the concept of *differential* was extended from the Euclidean spaces of classical analysis to arbitrary convergence spaces. We restrict the definition of differential given in [8] and verify that this restriction also correctly generalizes the differentials of classical analysis. We then use this definition to construct theories of differential calculus on finite convergence spaces and sequences.

4.1 Generalized Differentials

Definition 4.1. Let X and Y be convergence spaces, let $\mathcal{D}(X, Y)$ be a subspace of $\mathcal{C}(X, Y)$, let $L \in \mathcal{D}(X, Y)$, let $f : X \rightarrow Y$ be a function, and let $a \in X$. Then L is a *differential* of f at a if and only if for every filter \mathcal{A} converging to a in X , there exists a filter \mathcal{L} converging to L in $\mathcal{D}(X, Y)$ such that for every element K in \mathcal{L} , there exists a set A in \mathcal{A} such that for every element x in A , there exists a function k in K such that $k(x) = f(x)$. The function f is *differentiable* at a if and only if f has a differential at a .

When we wish to refer to differentials in the context of elementary differential calculus

on, for example, Euclidean spaces, we will refer in the sequel to *classical differentials*.

The differentials are not fully determined by X and Y —for all (not necessarily distinct) convergence spaces we *choose* $\mathcal{D}(X, Y)$ to be a subspace of $\mathcal{C}(X, Y)$ so that the composition of differentials is always a differential.

In the case of functions on the Euclidean line, the usual choice for $\mathcal{D}(\mathbb{R}, \mathbb{R})$ is the space of linear functions on \mathbb{R} . Linearity, however, is not a necessary property of classical differentials: instead, we can choose—and, in this paper, we will choose—the space of affine functions for $\mathcal{D}(\mathbb{R}, \mathbb{R})$. It appears that the key property of the space of linear functions on \mathbb{R} is that it is homeomorphic to \mathbb{R} . Of course, the differentials of classical analysis are linear, not affine; our notion of differentials produces an affine approximation of a function at a point. This, however, is a reasonable trade-off: while we sustain only a marginal loss of simplicity in classical differential calculus, we gain a notion of differential applicable in arbitrary convergence spaces, not merely those convergence spaces possessing the required linear structure. Proposition 4.2 addresses this issue for affine functions.

Proposition 4.2. *Let f be a function on \mathbb{R} and let $a \in \mathbb{R}$. Define $\mathbb{A}_{f,a}$ to be the subspace of $\mathcal{C}(\mathbb{R}, \mathbb{R})$ consisting of all affine functions on \mathbb{R} that equal $f(a)$ at a . If \mathbb{R} is given the standard topology, then $\mathbb{A}_{f,a}$ is homeomorphic to \mathbb{R} .*

Proof. Define $\phi : \mathbb{R} \rightarrow \mathbb{A}_{f,a}$ by $\phi(p) = \lambda x.((x - a)p + f(a))$ for every $p \in \mathbb{R}$. It is clear that ϕ is a bijection.

Since \mathbb{R} is locally compact, by Corollary 1.5.17 of [4], the convergence structure of $\mathcal{C}(\mathbb{R}, \mathbb{R})$ is topological and is induced by the compact-open topology. We adapt a metric discussed in Willard [43] to obtain the metric

$$d(f, g) = \sum_{n=1}^{\infty} \frac{1}{2^n} \cdot \frac{d_n(f, g)}{1 + d_n(f, g)},$$

in which $d_n(f, g) = \sup\{|f(x) - g(x)| : x = [-n, n]\}$. Thus, if p and q belong to \mathbb{R} , then

$d_n(\phi(p), \phi(q)) = (n + |a|)|p - q|$, and so

$$d(\phi(p), \phi(q)) < \frac{|p - q|}{1 + |p - q|} \sum_{n=1}^{\infty} \frac{n + |a|}{2^n} = \frac{(2 + |a|)|p - q|}{1 + |p - q|} < 2 + |a|.$$

Now to see that ϕ is continuous, let p and ε belong to \mathbb{R} with $\varepsilon > 0$. If $\varepsilon < 2 + |a|$, then let $\delta = \varepsilon/(2 + |a| - \varepsilon)$; otherwise, let $\delta > 0$. In either case $d(\phi(p), \phi(q)) < \varepsilon$ whenever $|p - q| < \delta$. To see that ϕ^{-1} is continuous, let p and ε belong to \mathbb{R} with $\varepsilon > 0$. Let $\delta = ((1 + |a|)\varepsilon)/(2(1 + (1 + |a|)\varepsilon))$. If $d(\phi(p), \phi(q)) < \delta$, then

$$\frac{(1 + |a|)|p - q|}{2(1 + (1 + |a|)|p - q|)} < d(\phi(p), \phi(q)) < \frac{(1 + |a|)\varepsilon}{2(1 + (1 + |a|)\varepsilon)},$$

which implies that $|p - q| < \varepsilon$. □

Notice that this proof also shows that the space of linear functions is homeomorphic to \mathbb{R} : let f be linear and let $a = 0$.

Together with Proposition 4.2, the remainder of this section establishes that we have constructed a conservative extension of elementary differential calculus, that is, we have not changed the notion of differential in the context of elementary differential calculus.

The next two propositions specialize Definition 4.1 in the cases where X is pretopological and where $\mathcal{D}(X, Y)$ is pretopological.

Proposition 4.3. *Let X be a pretopological space, let Y be a convergence space, let $L \in \mathcal{D}(X, Y)$, let $f : X \rightarrow Y$ be a function, and let $a \in X$. If L is a differential of f at a , then for every neighborhood V of L , there exists a neighborhood U of a such that for every element x in U , there exists a function k in V such that $k(x) = f(x)$.*

Proof. Let V be a neighborhood of L . Since X is pretopological, it follows that $\mathcal{N}_a \downarrow a$. Because L is a differential of f at a , there exists a filter \mathcal{L} converging to L such that for every element K in \mathcal{L} , there exists a neighborhood U of a such that for every element x in U , there exists a function k in K such that $k(x) = f(x)$. By hypothesis $V \in \mathcal{L}$. Therefore, for

every neighborhood V of L , there exists a neighborhood U of a such that for every element x in U , there exists a function k in V such that $k(x) = f(x)$. \square

Proposition 4.4. *Let X and Y be convergence spaces, let $\mathcal{D}(X, Y)$ be a pretopological subspace of $\mathcal{C}(X, Y)$, let $L \in \mathcal{D}(X, Y)$, let $f : X \rightarrow Y$ be a function, and let $a \in X$. If for every neighborhood V of L , there exists a neighborhood U of a such that for every element x in U , there exists a function k in V such that $k(x) = f(x)$, then L is a differential of f at a*

Proof. If $\mathcal{A} \downarrow a$, then $\mathcal{A} \supseteq \mathcal{N}_a$; thus \mathcal{A} contains all of the neighborhoods of a . Since $\mathcal{D}(X, Y)$ is pretopological, it follows that $\mathcal{N}_L \downarrow L$. By hypothesis, for every $V \in \mathcal{N}_L$, there exists $A \in \mathcal{A}$ such that for every $x \in A$, there exists $k \in V$ such that $k(x) = f(x)$. Therefore L is a differential of f at a . \square

If both X and $\mathcal{D}(X, Y)$ are pretopological, we obtain an equivalence that simplifies Definition 4.1.

Proposition 4.5. *If X is a pretopological space, Y is a convergence space, $\mathcal{D}(X, Y)$ is a pretopological subspace of $\mathcal{C}(X, Y)$, $L \in \mathcal{D}(X, Y)$, $f : X \rightarrow Y$ is a function, and $a \in X$, then L is a differential of f at a if and only if for every neighborhood V of L , there exists a neighborhood U of a such that for every element x in U , there exists a function k in V such that $k(x) = f(x)$.*

Proof. The desired result follows immediately from Propositions 4.3 and 4.4. \square

In view of Proposition 4.2, we see that $\mathbb{A}_{f,a}$, the space of all affine functions on \mathbb{R} that equal $f(a)$ at a , is pretopological; thus Proposition 4.5 applies to the case of functions on \mathbb{R} .

Theorem 4.6. *If \mathbb{R} has the standard topology, $L \in \mathbb{A}_{f,a}$, f is a function on \mathbb{R} , and $a \in \mathbb{R}$, then L is a generalized differential of f at a if and only if L is a classical differential of f at a .*

Proof. [Necessity]. Let V be a neighborhood of $L(h) = (h - a)m + f(a)$. By Proposition 4.2, there exists $\varepsilon > 0$ such that N_ε is a neighborhood of m . By hypothesis, there exists $\delta > 0$

such that $|(f(x) - f(a))/(x - a) - m| < \varepsilon$ whenever $|x - a| < \delta$. Again by Proposition 4.2, the function $k(h) = (h - a)(f(x) - f(a))/(x - a) + f(a)$ belongs to V . Since $k(x) = f(x)$, it follows that L is a generalized differential of f at a .

[Sufficiency]. Let $\varepsilon > 0$. If N_ε is a neighborhood of m , then by Proposition 4.2 there exists an affine function $L(h) = (h - a)m + f(a)$ and a neighborhood V of L corresponding to N_ε . By hypothesis, there exists $\delta > 0$ such that for every x such that $0 < |x - a| < \delta$, there exists an affine function $k(h) = (h - a)m' + f(a)$ in V such that $k(x) = f(x)$. Since $k \in V$, it follows that $m' \in N_\varepsilon$. Because $k(x) = f(x)$, it follows that $m' = (f(x) - f(a))/(x - a)$. Thus

$$|(f(x) - f(a))/(x - a) - m| = |m' - m| < \varepsilon$$

whenever $0 < |x - a| < \delta$, which implies that $m = \lim_{x \rightarrow a} (f(x) - f(a))/(x - a) = f'(a)$, and so L is a classical differential of f at a . \square

We conclude this section with the observation that the so-called chain rule of classical analysis also holds for differentiable functions on convergence spaces.

Theorem 4.7. *Let X , Y , and Z be convergence spaces, let $a \in X$, let $f : Y \rightarrow Z$ be a function, and let $g : X \rightarrow Y$ be a function continuous at a . If L_f is a differential of f at $g(a)$ and L_g is a differential of g at a , then $L_f \circ L_g$ is a differential of $f \circ g$ at a .*

Proof. Let $\mathcal{A} \downarrow a$. Since L_g is a differential of g at a , there exists a filter \mathcal{K} converging to L_g such that for every element K in \mathcal{K} , there exists a set A in \mathcal{A} such that for every element x in A , there exists a function k in K such that $k(x) = g(x)$. Since g is continuous at a , it follows that $g(\mathcal{A}) \downarrow g(a)$; thus, by the hypothesis that L_f is a differential of f at $g(a)$, there exists a filter \mathcal{H} converging to L_f such that for every element H in \mathcal{H} , there exists a set B in $g(\mathcal{A})$ such that for every element x in B , there exists a function h in H such that $h(x) = f(x)$.

Now define the filter $\mathcal{L} = [\{h \circ k : h \in H \wedge k \in K\} : H \in \mathcal{H} \wedge K \in \mathcal{K}]$. If $\mathcal{F} \downarrow x$, then $\mathcal{H} \cdot (\mathcal{K} \cdot \mathcal{F}) \downarrow (L_f \circ L_g)(x)$. Since $\mathcal{L} \cdot \mathcal{F}$ includes $\mathcal{H} \cdot (\mathcal{K} \cdot \mathcal{F})$, it follows that $\mathcal{L} \downarrow L_f \circ L_g$. If

$L \in \mathcal{L}$, then $L \supseteq \{h \circ k : h \in H \wedge k \in K\}$ for some $H \in \mathcal{H}$ and $K \in \mathcal{K}$. Thus, there exists a set A in \mathcal{A} such that for every element x in A , there exists a function k in K such that $k(x) = g(x)$; likewise, there exists a set B in $g(\mathcal{A})$ such that for every element x in B , there exists a function h in H such that $h(x) = f(x)$. If $x \in A \cap g^{-1}(B)$, then there exist $k \in K$ and $h \in H$ such that $(h \circ k)(x) = h(k(x)) = h(g(x)) = f(g(x)) = (f \circ g)(x)$. Therefore, we conclude that $L_f \circ L_g$ is a differential of $f \circ g$ at a . \square

4.2 Differential Calculus on Finite Convergence Spaces

In this section, we first obtain an equivalent condition for differentiability of functions between finite convergence spaces without imposing any *a priori* restrictions on the space of differentials. We then specialize this condition for finite Kolmogorov spaces, finite groups, and Boolean hypercubes.

Theorem 4.8. *If X and Y are finite convergence spaces, $L \in \mathcal{D}(X, Y)$, $f : X \rightarrow Y$ is a function, and $a \in X$, then L is a differential of f at a if and only if for every $x \in \overrightarrow{a}$, there exists $k \in \overrightarrow{L}$ such that $k(x) = f(x)$.*

Proof. [Necessity]. Let V be a neighborhood of L . By hypothesis, there exists a neighborhood of a , namely \overrightarrow{a} , such that for every x in \overrightarrow{a} , there exists a function k in \overrightarrow{L} , and hence in V , such that $k(x) = f(x)$.

[Sufficiency]. If L is a differential of f at a , then by Proposition 4.5 for every neighborhood V of L , there exists a neighborhood U of a such that for every element x in U , there exists a function k in V such that $k(x) = f(x)$. Since \overrightarrow{L} is a neighborhood of L and any neighborhood of a includes \overrightarrow{a} , it follows that for every element x in \overrightarrow{a} , there exists a function k in \overrightarrow{L} such that $k(x) = f(x)$. \square

Corollary 4.9. *Let X and Y be finite convergence spaces, $L \in \mathcal{D}(X, Y)$, $f : X \rightarrow Y$ be a function, and $a \in X$. If L is a differential of f at a , then $f(x) \in \overrightarrow{L(a)}$ for every $x \in \overrightarrow{a}$.*

Proof. By Theorem 4.8, for every $x \in \vec{a}$, there exists $k \in \vec{L}$ such that $k(x) = f(x)$. In view of Proposition 3.14, we see that $f(x) = k(x) \in \overline{L(a)}$ for every element $x \in \vec{a}$. \square

The intuitive meaning of Corollary 4.9 is that a function between two finite convergence spaces that is differentiable at a point must locally behave as a differential. The converse of Corollary 4.9, however, does not hold in general. The question then arises: how close *should* a function be to a differential near a point for that function to be differentiable at that point. Consequently, we suggest a weakening of Definition 4.1.

Definition 4.10. Let X and Y be convergence spaces, let $\mathcal{D}(X, Y)$ be a subspace of $\mathcal{C}(X, Y)$, let $L \in \mathcal{D}(X, Y)$, let $f : X \rightarrow Y$ be a function, and let $a \in X$. Then L is a *weak differential* of f at a if and only if for every filter \mathcal{A} converging to a , there exists a filter \mathcal{L} converging to L such that for every element K in \mathcal{L} , there exists a set A in \mathcal{A} such that for every element x in A , there exists a function k in K such that $f(x) \in k(A)$. The function f is *weakly differentiable* at a if and only if f has a weak differential at a .

It is a simple verification that a function differentiable at a point is also weakly differentiable at that point; a function weakly differentiable at a point, however, need not be differentiable at that point. Moreover, if X is a finite convergence space and $\mathcal{D}(X, X) = \text{Aut}(X)$, then the converse of Corollary 4.9 holds *mutatis mutandis*. We adjourn to future research the exploration of this wider class of differentials.

4.2.1 Differential Calculus on Finite Kolmogorov Spaces

We presently restrict our attention to functions a finite Kolmogorov space. Given a convergence space X , a possible choice for $\mathcal{D}(X, X)$ is $\text{Aut}(X)$. Indeed, we assume throughout this subsection and the next that $\mathcal{D}(X, X)$ is the set of automorphisms on X . The discussion following the proof of Theorem 4.13 will clarify the motivation for this choice. We first obtain an equivalent condition for differentiability given a discrete space of differentials.

Proposition 4.11. *If X and Y are convergence spaces, $\mathcal{D}(X, Y)$ is a discrete subspace of $\mathcal{C}(X, Y)$, $L \in \mathcal{D}(X, Y)$, $f : X \rightarrow Y$ is a function, and $a \in X$, then L is a differential of f at a if and only if for every filter \mathcal{A} converging to a , there exists a set A in \mathcal{A} such that $L(x) = f(x)$ for every element x in A .*

Proof. [Necessity]. If \mathcal{A} converges to a , then by hypothesis there exists a set A in \mathcal{A} such that $L(x) = f(x)$ for every element x in A . Now $[L]$ is a filter converging to L such that L belongs to each element K of $[L]$ and $L(x) = f(x)$ for every element x in A . Therefore L is a differential of f at a .

[Sufficiency]. Let \mathcal{A} converge to a . Since $\mathcal{D}(X, Y)$ is discrete, the only filter that converges to L is $[L]$; since $\{L\} \in [L]$, there exists a set A in \mathcal{A} such that $L(x) = f(x)$ for every element x in A . □

Of course, we can reformulate Proposition 4.11 in terms of neighborhoods if we additionally require that X is pretopological.

Proposition 4.12. *If X is a pretopological space, Y is a convergence space, $\mathcal{D}(X, Y)$ is a discrete subspace of $\mathcal{C}(X, Y)$, $L \in \mathcal{D}(X, Y)$, $f : X \rightarrow Y$ is a function, and $a \in X$, then L is a differential of f at a if and only if there exists a neighborhood U of a such that $L(x) = f(x)$ for every element x in U .*

Proof. [Necessity]. If V is a neighborhood of L , then $L \in V$, and so by hypothesis there exists a neighborhood U of a such that for every element x in U , there exists a function k in V such that $k(x) = f(x)$. Thus by Proposition 4.5 it follows that L is a differential of f at a .

[Sufficiency]. By Proposition 4.5, for every neighborhood V of L , there exists a neighborhood U of a such that for every element x in U , there exists a function k in V such that $k(x) = f(x)$. Since $\mathcal{D}(X, Y)$ is discrete, it follows that $\{L\}$ is a neighborhood of L . Therefore there exists a neighborhood U of a such that $L(x) = f(x)$ for every element x in U . □

Theorem 4.13. *If X is a finite Kolmogorov space, $L \in \mathcal{D}(X, X)$, f is a function on X , and $a \in X$, then L is a differential of f at a if and only if $L(x) = f(x)$ for every $x \in \vec{a}$.*

Proof. [Necessity]. In view of Theorem 2.103 and Proposition 4.12, it suffices to show that there exists a neighborhood U of a such that $L(x) = f(x)$ for every $x \in U$. By hypothesis \vec{a} is such a neighborhood. Therefore L is a differential of f at a .

[Sufficiency]. By Theorem 2.103 and Proposition 4.12, it follows that there exists a neighborhood U of a such that $L(x) = f(x)$ for every $x \in U$. Since $\vec{a} \subseteq U$, we conclude that $L(x) = f(x)$ for every $x \in \vec{a}$. \square

In other words, Theorem 4.13 states that a function on a finite Kolmogorov space is differentiable at a point if and only if it is bicontinuous when restricted to the graph neighborhood of that point; this means that the only functions differentiable everywhere on a finite Kolmogorov space are the automorphisms on that space. This result supports the intuition that automorphisms are natural candidates for differentials. A classical differential of a function at a point is a linear approximation to the function that is identical to the function at that point. Although this approximation can be made arbitrarily close to the actual function on a sufficiently small neighborhood, it is not usually identical to the function (unless the function itself is linear on some neighborhood of the point). Since every point of a finite reflexive digraph, however, does have a smallest neighborhood, namely its graph neighborhood, we should expect that a differential of a function on a finite reflexive digraph at a point to coincide with that function on the graph neighborhood of that point.

Nevertheless, a possible objection to restricting $\mathcal{D}(X, X)$ to $\text{Aut}(X)$ is that this selection excludes from $\mathcal{D}(X, X)$ the constant functions on X , contrary to our experience with classical analysis on Euclidean space. Proposition 3.14, however, ensures that the results of this section still hold if we expand $\mathcal{D}(X, X)$ to include not only $\text{Aut}(X)$ but the constant functions on X as well. Intuitively, any discrete subspace of $\mathcal{C}(X, X)$ is an appropriate choice for $\mathcal{D}(X, X)$.¹

¹With that said, we argue that it is myopic to demand that generalized differentials behave *precisely* like

The next proposition establishes that a function on a finite Kolmogorov space that is differentiable at some point must also be continuous at that point.²

Proposition 4.14. *Let X be a finite Kolmogorov space, let f be a function on X , and let $a \in X$. If f is differentiable at a , then f is continuous at a .*

Proof. If $b \in f(\vec{a})$, then there exists $c \in \vec{a}$ such that $f(c) = b$. Since $c \in \vec{a}$, it follows that $\vec{a} \in [c]$, and so $[c] \downarrow a$. By hypothesis, there exists an automorphism L such that $L(x) = f(x)$ for every $x \in \vec{a}$. Thus $[b] = [f(c)] = [L(c)]$ converges to $L(a) = f(a)$, which implies that $\vec{f(a)} \in [b]$, and so $b \in \vec{f(a)}$, from which it follows that $f(\vec{a}) \subseteq \vec{f(a)}$. Therefore, by Proposition 3.4 we conclude that f is continuous at a . \square

If we know the differential of each function at a particular point of a homogeneous finite Kolmogorov space, then the homogeneity of the space provides us with a means to compute, on the basis of this knowledge alone, the differential of each function at every point of the space.

Proposition 4.15. *If X is a finite Kolmogorov space, $L \in \mathcal{D}(X, X)$, f is a function on X , a and b belong to X , and ϕ is automorphism that maps b to a , then L is a differential of f at a if and only if $L \circ \phi$ is a differential of $f \circ \phi$ at b .*

Proof. [Necessity]. Since ϕ^{-1} is an automorphism, it belongs to $\mathcal{D}(X, X)$; in particular ϕ^{-1} is a differential of itself at a . Thus, by the chain rule $L = L \circ \phi \circ \phi^{-1}$ is a differential of $f \circ \phi \circ \phi^{-1} = f$ at a .

[Sufficiency]. Since ϕ is an automorphism, it belongs to $\mathcal{D}(X, X)$; in particular ϕ is a differential of itself at b . Thus, by the chain rule $L \circ \phi$ is a differential of $f \circ \phi$ at b . \square

classical differentials. In the sequel, we witness several unexpected characteristics of generalized differentials, even in the absence of *a priori* restrictions—for instance, Example 4.52 establishes that a function may be differentiable at a point at which it is not continuous.

²A function on an arbitrary convergence space that is differentiable at some point need not be continuous at that point. See Example 4.52.

4.2.2 Differential Calculus on Finite Groups

Now we apply the results of Chapter 3 and Section 4.2.1 to develop differential calculus on finite groups. To this end, we identify a group with one of its Cayley graphs, and thereby speak of differentiability of functions between groups.³

In view of Theorems 3.23 and 3.24, we see that if C is a Cayley graph distinct from C_1 or C_2 , then C is Kolmogorov, and so Theorem 4.13 and its consequences hold; otherwise $\mathcal{D}(C, C)$ is indiscrete, and so every member of $\mathcal{D}(C, C)$ is a differential of every function on C at every point in C .

Theorem 4.16. *If C is either C_1 or C_2 , $L \in \mathcal{D}(C, C)$, f is a function on C , and $a \in C$, then L is a differential of f at a .*

Proof. Since C is indiscrete, by Proposition 2.86 it follows that $\mathcal{D}(C, C)$ is indiscrete. Thus, if $L \in \mathcal{D}(C, C)$, then the only neighborhood of L is $\mathcal{D}(C, C)$. Now X is a neighborhood of a such that for every x in X , there exists k in $\mathcal{D}(C, C)$, namely f , such that $k(x) = f(x)$. Therefore, by Proposition 4.5 we conclude that L is a differential of f at a . \square

Theorem 4.17. *If C is a Cayley graph distinct from C_1 and C_2 , $L \in \mathcal{D}(C, C)$, f is a function on C , and $a \in C$, then L is a differential of f at a if and only if $L(x) = f(x)$ for every $x \in \vec{a}$.*

Proof. By Theorems 3.23 and 3.24, it follows that C is a finite Kolmogorov space. In view of Proposition 3.9, we see that $\mathcal{N}_a = [\vec{a}]$, and so by Theorem 4.13 we conclude that L is a differential of f at a if and only if $L(x) = f(x)$ for every $x \in \vec{a}$. \square

In view of the homogeneity of Cayley graphs, we observe a specialization of Proposition 4.15: we need only know the differential of each function at e (or, in fact, any other element) to determine the differential of each function at each point.

³Identification of a group with one of its *non-redundant* Cayley graphs is analogous to selection of a basis for a vector space.




Figure 4.1: A Cayley graph of the Symmetric Group S_3 .


Proposition 4.18. *If C is a Cayley graph, $L \in \mathcal{D}(C, C)$, f is a function on C , and $a \in C$, then L is a differential of f at a if and only if $L \circ (\lambda x.ax)$ is a differential of $f \circ (\lambda x.ax)$ at e .*

Proof. The desired result follows immediately from Proposition 4.15. \square

We conclude this section with three examples.

Example 4.19. Consider the Cayley graph of the symmetric group S_3 , as represented in Figure 4.1. Observe that $\text{Aut}(S_3)$ is Λ_{S_3} . Which functions on S_3 have a differential at e ? Exactly those functions that are bicontinuous when restricted to $\vec{e} = \{e, r, t\}$. In other words L is a differential of f at e if and only if $f|_{\vec{e}} = L|_{\vec{e}}$. For example, the function $(e \ r \ r^2 \ t \ tr^2 \ tr \ e)$ is differentiable at e : its differential at e is $\lambda x.rx$. The function $\lambda x.xr$, however, is not differentiable at e : it is not identical with any of the candidate differentials on \vec{e} . If we determined all functions that have a differential at e , then by using Proposition 4.18 we could determine all functions differentiable somewhere on S_3 ; since there are $6^4 = 1296$ functions differentiable at e , we leave these computations as an exercise for the reader.⁴

⁴Note that the symmetric group S_3 is the automorphism group of S^3 , the tertiary Cartesian product of the Sierpiński space (that is, the set $\{0, 1\}$ equipped with the convergence structure defined by the equivalences $\mathcal{F} \downarrow 0$ if and only if $\{0, 1\} \in \mathcal{F}$ and $\mathcal{F} \downarrow 1$ if and only if $\{1\} \in \mathcal{F}$). In this role, however, the space S_3 is not the reflexive digraph of Figure 4.1, but rather a 6-point discrete space, as can be deduced from Theorem 4.13 or verified directly by computation.

Figure 4.2: A Cayley Graph of \mathbb{Z} .

Example 4.20. Suppose that we modify Definition 3.20 to include infinite groups. Identify \mathbb{Z} with the Cayley graph for \mathbb{Z} generated by $\{1\}$, as shown in Figure 4.2. Note that $\vec{n} = \{n, n + 1\}$ for each $n \in \mathbb{Z}$. Observe that $\text{Aut}(\mathbb{Z}) = \Lambda_{\mathbb{Z}}$.

If $\mathcal{F} \downarrow \lambda x.x + k$, then $\mathcal{F} \cdot [0, 1] \downarrow k$, which implies that $[k, k + 1] \subseteq \mathcal{F} \cdot [0, 1]$, and so there exists $F \in \mathcal{F}$ and $A \supseteq \{0, 1\}$ such that $F \cdot A \subseteq \{k, k + 1\}$. If $\lambda x.x + h \in F$, then

$$\{h, h + 1\} = (\lambda x.x + h)(\{0, 1\}) \subseteq (\lambda x.x + h)(A) \subseteq F \cdot A \subseteq \{k, k + 1\},$$

which implies that $h = k$; thus $\text{Aut}(\mathbb{Z})$ is discrete. Therefore, in view of Proposition 4.12, we see that a function f on \mathbb{Z} is differentiable at $n \in \mathbb{Z}$ if and only if $f(n + 1) = f(n) + 1$; its differential at n is $\lambda x.x + f(n) - n$.

Example 4.21. Finally, we consider differentials of functions between two distinct spaces. Consider \mathbb{Z} as in Example 4.21 and equip $\{0, 1\}$ with the discrete convergence structure. Observe that $\mathcal{C}(\mathbb{Z}, \{0, 1\})$ is discrete. Since $\{0, 1\}$ is discrete, the only members of $\mathcal{C}(\mathbb{Z}, \{0, 1\})$ are the constant functions $\bar{0}$ and $\bar{1}$. Moreover, as a convergence space $\mathcal{C}(\mathbb{Z}, \{0, 1\})$ is discrete. If we choose $\mathcal{D}(\mathbb{Z}, \{0, 1\})$ to be $\mathcal{C}(\mathbb{Z}, \{0, 1\})$, then, in view of Proposition 4.12, we see that L is a differential of f at n if and only if $L(n) = f(n)$ and $L(n + 1) = f(n + 1)$. This means that f is differentiable at n if and only if $f(n + 1) = f(n)$.

4.2.3 Differential Calculus on Boolean Hypercubes

We now consider, as an application of Theorem 4.8, differential calculus on Boolean hypercubes.




Figure 4.3: The 3-Boolean Hypercube.

Definition 4.22. Denote by B the set $\{0, 1\}$. For every positive integer n , the n -*Boolean hypercube* is the Cartesian product B^n equipped with the reflexive digraph convergence structure defined by the condition: $b \in \vec{a}$ if and only if there exists at most one $1 \leq i \leq n$ such that $\pi_i(b) \neq \pi_i(a)$.

Example 4.23. Figure 4.3 depicts the 3-Boolean hypercube.

According to Rudeanu [37], construction of a Boolean differential calculus was first attempted by Daniell [15]. Recently, Bacsó and Lábos [3] suggest that Boolean differentials⁵ should be linear and satisfy the Leibniz identity; no proposed Boolean differential, they claim, meets both criteria. By Theorem 4.7, however, we know that generalized differentials always satisfy the chain rule, and hence the Leibniz identity if multiplication and addition are defined on the codomain. Thus, if we restrict the space of differentials to linear functions, then we satisfy the criteria of [3].

In conformity with [3], throughout this section we restrict the space of differentials to linear functions. We also choose as a (vector space) basis for each Boolean hypercube its standard basis.

⁵To be precise, total differentials, in contrast to partial differentials. This, however, is presently moot since generalized differentials are total differentials.

Proposition 4.24. *A linear function $L : B^m \rightarrow B^n$ is continuous if and only if $L(\vec{0}_m) \subseteq \vec{0}_n$.*

Proof. [Necessity]. Let $b \in B^m$. If $b \neq y \in L(\vec{b})$, then there exists $x \in \vec{b}$ such that $L(x) = y$. By hypothesis it follows that

$$y - L(b) = L(x) - L(b) = L(x - b) \in L(\vec{0}_m) \subseteq \vec{0}_n.$$

Thus $y \in \vec{0}_n$, which implies that $L(\vec{b}) \subseteq \vec{0}_n$, and so L is continuous by Proposition 3.4.

[Sufficiency]. The desired result follows immediately from Proposition 3.4. \square

In view of Proposition 4.24, we see that a linear function $L : B^m \rightarrow B^n$ is continuous if and only if it is represented by an $n \times m$ matrix, the columns of which are elements of $\vec{0}_n$. Together with Proposition 3.14, this observation reveals the following characterization of $\mathcal{C}(B^m, B^n)$ restricted to linear functions.

Proposition 4.25. *If L and K are linear functions in $\mathcal{C}(B^m, B^n)$, then $K \in \vec{L}$ if and only if $K = L$ or all nonzero columns of K and L are identical.*

Proof. [Necessity]. Let $x \in B^m$. If $y \in \vec{x}$, then $x - y = b$ for some $b \in \vec{0}_m$. By hypothesis $L(x) - K(y) = (L - K)(x) - K(b) \in \vec{0}_n$, which implies that $K(y) \in \vec{L}(x)$; thus $K \downarrow [L]$ by Proposition 3.14. Therefore $K \in \vec{L}$.

[Sufficiency]. We proceed by contraposition. Thus, there exist distinct $i, j \in \vec{0}_m - \{0_m\}$ such that $L(i) \neq L(j)$. Since K is continuous, it follows that $K(i) \in \vec{0}_n$. But $L(i+j) - K(i) \notin \vec{0}_n$, which implies that $K(i) \notin \vec{L}(i+j)$, and so $[K] \not\downarrow L$ by Proposition 3.14. Therefore $K \notin \vec{L}$. \square

Theorem 4.8 and Proposition 4.25 provide us with an equivalent condition for differentiability on Boolean hypercubes.

Theorem 4.26. *If $L \in \mathcal{D}(B^m, B^n)$, $f : B^m \rightarrow B^n$ is a function, and $b \in B^m$, then L is a differential of f at b if and only if*

1. $L(x) = f(x)$ for every $x \in \vec{b}$, or
2. $f(\vec{b}) \subseteq L(B^m) = \{0, \beta\}$ for some $\beta \in \vec{0}_n$ and $f(0_m) = 0_n$ if $b \in \vec{0}_m$.

Proof. [Necessity]. If the first case holds, then by Theorem 4.8 it immediately follows that L is a differential of f at b . If the second case holds, then each nonzero column of L is β . Thus, if $f(x) \neq L(x)$, then $L + \beta f(x)^T$ belongs to \vec{L} and equals $f(x)$ at x , and so L is a differential of f at b by Theorem 4.8.

[Sufficiency]. By Theorem 4.8, for every $x \in \vec{b}$ there exists $K \in \vec{L}$ such that $K(x) = f(x)$. If the second case does not hold, then by hypothesis there does not exist $\beta \in \vec{0}_n$ such that $L(B^m) = \{0, \beta\}$; in other words, the nonzero columns of L are not identical, which by Proposition 4.25 implies that $K = L$. Therefore $L(x) = f(x)$ for every $x \in \vec{b}$. \square

A corollary to Theorem 4.26 is that every function $f : B^m \rightarrow B$ is differentiable at every point not in $\vec{0}_m$ and at every point in $\vec{0}_m$ if $f(0_m) = 0$.

Example 4.27. The function $f : B^2 \rightarrow B^3$ defined by $f(p, q) = (p, (1+p)(1+q), q)$ is differentiable only at $(1, 1)$; its differential at $(1, 1)$ is $\lambda(p, q) \cdot (p, 0, q)$. The function $g : B^3 \rightarrow B^2$ defined by $g(p, q, r) = ((1+q)(1+p+pr), (1+r)q)$ is differentiable at $(1, 0, 1)$; one of its differentials at $(1, 0, 1)$ is $\lambda(p, q, r) \cdot (q+r, 0)$. By Theorem 4.7, it follows that $(g \circ f)(p, q) = (q, (1+p)(1+q))$ is differentiable at $(1, 1)$; one of its differentials at $(1, 1)$ is $\lambda(p, q) \cdot (q, 0)$.

4.3 Differential Calculus on Labeled Graphs

Definition 4.28. Let (V, E) be a reflexive digraph and let L be a set. Associate with each edge in E exactly one element of L . We call each element of L a *label*. If $(u, v) \in E$ is associated with $x \in L$, we write $x(u) = v$. We say that the triple (V, E, L) is a *labeled graph*. When no reasonable confusion is likely, we refer to a labeled graph (V, E, L) by V and denote

L by $\mathcal{L}(V)$. If V and W are two labeled graphs, then a function $\lambda : \mathcal{L}(V) \rightarrow \mathcal{L}(W)$ is called a *labeling function*.

Given reflexive digraphs V and W , there are often several labeling functions $\lambda : \mathcal{L}(V) \rightarrow \mathcal{L}(W)$; we require, however, the choice of precisely one labeling function (although the choice itself may vary according to the intended model or application). The labeling function provides us with a means of restricting the space of differentials $\mathcal{D}(V, W)$. Since differentials are continuous, they preserve edges. To strengthen this property of differentials, we require that differentials also preserve the labeling function: if V and W are labeled graphs, $\lambda : \mathcal{L}(V) \rightarrow \mathcal{L}(W)$ is a labeling function, and $d \in \mathcal{D}(V, W)$, then $d(x(v)) = (\lambda(x))(d(v))$ for every $x \in \mathcal{L}(V)$ and $v \in V$. The next proposition shows that this property constrains not only the choice of differentials but the choice of labeling function as well.

Proposition 4.29. *Let U , V , and W be labeled graphs. If $\lambda : \mathcal{L}(U) \rightarrow \mathcal{L}(V)$, $\mu : \mathcal{L}(V) \rightarrow \mathcal{L}(W)$, and $\nu : \mathcal{L}(U) \rightarrow \mathcal{L}(W)$ are labeling functions, then $\nu = \mu \circ \lambda$.*

Proof. Let $h \in \mathcal{D}(U, V)$ and $k \in \mathcal{D}(V, W)$. By hypothesis, we have that

$$\begin{aligned} (\nu(x))((k \circ h)(u)) &= (k \circ h)(x(u)) = k(h(x(u))) \\ &= k((\lambda(x))(h(u))) = (\mu(\lambda(x)))(k(h(u))) = ((\mu \circ \lambda)(x))((k \circ h)(u)) \end{aligned}$$

for every $x \in \mathcal{L}(U)$ and $u \in U$, from which the desired result follows. \square

Note that we have not used the chain rule in deriving this property of labeling functions but only that spaces of differentials are chosen so that the composition of differentials is always a differential. Also observe that if $U = V = W$, then $\lambda = \mu = \nu$, and so λ is idempotent.

Examples of labeled graphs include Cayley graphs and deterministic finite automata. Indeed, if V is a Cayley graph, then the labels are the generators of V together with the

identity element e . Thus every label $x \in \mathcal{L}(V)$ corresponds to a generator (or e), and so we may write $x(v) = vx$ for each $v \in V$.

Proposition 4.30. *If V and W are Cayley graphs and $\lambda : \mathcal{L}(V) \rightarrow \mathcal{L}(W)$ is a labeling function, then $d(vx) = d(v)\lambda(x)$ for every $x \in \mathcal{L}(V)$ and $v \in V$.*

Proof. Since $d(x(v)) = (\lambda(x))(d(v))$ and $x(v) = vx$, it follows that $d(vx) = d(x(v)) = (\lambda(x))(d(v)) = d(v)\lambda(x)$. \square

Note that Proposition 4.30 implies that labeling functions on Cayley graphs preserve identity elements and orders of elements: if V and W are Cayley graphs, $\lambda : \mathcal{L}(V) \rightarrow \mathcal{L}(W)$, and $v \in V$ has order n , then $\lambda(e_V) = e_W$ and $\lambda(v)$ has order n .

In Section 4.2.2, we restricted $\mathcal{D}(V, V)$ for a Cayley graph V to the set of automorphisms on V . As observed in Chapter 3, among these automorphisms are the left-multiplications on V . Independent of this restriction on $\mathcal{D}(V, V)$, if $\mathcal{D}(V, V)$ contains *any* left-multiplications on V , then the only labeling function on $\mathcal{L}(V)$ is the identity function.

Proposition 4.31. *If V is a Cayley graph and $\mathcal{D}(V, V) \cap \Lambda_V \neq \emptyset$, then the only labeling function on $\mathcal{L}(V)$ is the identity function on $\mathcal{L}(V)$.*

Proof. Let $\mu : \mathcal{L}(V) \rightarrow \mathcal{L}(V)$ be a labeling function. By hypothesis, there exists $a \in V$ such that $\lambda x.ax \in \mathcal{D}(V, V)$. Thus, if $y \in \mathcal{L}(V)$ and $v \in V$, then by Proposition 4.30 it follows that $avy = (\lambda x.ax)(vy) = (\lambda x.ax)(v)\mu(y) = av\mu(y)$, which implies that $\mu(y) = y$. Therefore μ is the identity function on $\mathcal{L}(V)$. \square

In view of Proposition 4.31, we see that for any Cayley graph V , if $\mathcal{D}(V, V)$ contains any left-multiplication on V , then the only possible labeling function on $\mathcal{L}(V)$ is the identity function on $\mathcal{L}(V)$. This result forces $\mathcal{D}(V, V)$ to contain only automorphisms on V .

Proposition 4.32. *If V is a Cayley graph and $\mathcal{D}(V, V) \cap \Lambda_V \neq \emptyset$, then $\mathcal{D}(V, V)$ is a subset of the set of automorphisms on V .*

Proof. Let $d \in \mathcal{D}(V, V)$ and let $v \in V$. Since $\mathcal{D}(V, V) \subseteq \mathcal{C}(V, V)$, it follows that $d(\vec{v}) \subseteq \overrightarrow{d(v)}$. If $w \in \overrightarrow{d(v)}$, then there exists a generator g such that $w = d(v)g$. By Propositions 4.30 and 4.31, we infer that $w = d(vg)$. Since $vg \in \vec{v}$, it follows that $w \in d(\vec{v})$; thus $d(\vec{v}) \supseteq \overrightarrow{d(v)}$, and therefore d is an automorphism. \square

It is possible to strengthen Proposition 4.32: if V is a Cayley graph and the chosen labeling function is the identity function on $\mathcal{L}(V)$, then $\mathcal{D}(V, V)$ is a subset of the set of automorphisms on V .

We conclude this section by revisiting the Cayley graph of the symmetric group S_3 , as in Example 4.19.

Example 4.33. Consider the Cayley graph of the symmetric group S_3 , as in Example 4.19. Let $\lambda : \mathcal{L}(S_3) \rightarrow \mathcal{L}(S_3)$ be a labeling function. Since λ preserves orders of elements, it follows that $\lambda(r)$ is either e or r and $\lambda(t)$ is either e or t . If λ is constant, then $\mathcal{D}(S_3, S_3)$ contains only constant functions. As noted above, if λ is the identity function, then $\mathcal{D}(S_3, S_3)$ contains only automorphisms; since the set of automorphisms on S_3 is precisely Λ_{S_3} , it follows that $\mathcal{D}(S_3, S_3)$ contains only left-multiplications. If $\lambda(r) = e$ and $\lambda(t) = t$, then each function $d \in \mathcal{D}(S_3, S_3)$ is of the form

$$d(v) = \begin{cases} k, & v = r^n; \\ kt, & \text{otherwise,} \end{cases}$$

for some constant k . It is easy to verify that $\mathcal{D}(S_3, S_3)$ is discrete. An exactly similar argument holds if $\lambda(r) = r$ and $\lambda(t) = e$.

4.4 Differential Calculus on the Cantor Tree

In this section, we apply Proposition 4.12 to determine an equivalent condition for differentiability on infinite binary trees.




Figure 4.4: The Cantor Tree.

Definition 4.34. Denote by \mathbb{T} the *Cantor tree*, that is, the infinite full binary tree, equipped with a partial order, as represented by the Hasse diagram of Figure 4.4. For each $a \in \mathbb{T}$, the *depth* of a , denoted by $\text{depth}(a)$, is the length of the path in the tree from the root to a .

In view of Proposition 3.5, we see that depth is an invariant under automorphisms on \mathbb{T} ; in particular, each automorphism on the Cantor tree has the root as a fixed point. Consequently, the space of automorphisms on \mathbb{T} is discrete.

Proposition 4.35. *If $f \in \text{Aut}(\mathbb{T})$, then $\text{depth}(f(a)) = \text{depth}(a)$ for each $a \in \mathbb{T}$.*

Proof. Let $a \in \mathbb{T}$. We proceed by strong induction on $\text{depth}(a)$.

First suppose that $\text{depth}(a) = 0$. If $\text{depth}(a) < \text{depth}(f(a))$, then $a < f(a)$, and so $f^{-1}(a) < f^{-1}(f(a)) = a$, which is absurd; thus $\text{depth}(f(a)) = \text{depth}(a)$.

Now suppose that $\text{depth}(f(b)) = \text{depth}(b)$ for all b such that $\text{depth}(b) < \text{depth}(a)$. Since f is a bijection, the induction hypothesis precludes the case in which $\text{depth}(f(a)) < \text{depth}(a)$. If $\text{depth}(a) < \text{depth}(f(a))$, then there exists $c \in \mathbb{T}$ such that $\text{depth}(c) = \text{depth}(a)$ and $c < f(a)$. Since f is a bijection, there exists $c' \in \mathbb{T}$ such that $f(c') = c$; thus $c' = f^{-1}(c) < f^{-1}(f(a)) = a$, which implies that $\text{depth}(c') < \text{depth}(a)$, and so by the induction hypothesis

$\text{depth}(c) = \text{depth}(f(c')) = \text{depth}(c') < \text{depth}(a)$, which is absurd. Therefore $\text{depth}(f(a)) = \text{depth}(a)$, as desired. \square

Proposition 4.36. *$\text{Aut}(\mathbb{T})$ is discrete.*

Proof. Suppose that \mathcal{F} converges to f in $\text{Aut}(\mathbb{T})$. If $a \in \mathbb{T}$, then $\mathcal{F} \cdot [a]$ converges to $f(a)$ in \mathbb{T} , and so $\overrightarrow{f(a)} \in \mathcal{F} \cdot [a]$, which implies that there exists $F \in \mathcal{F}$ such that $F \cdot \{a\} \subseteq \overrightarrow{f(a)}$; thus, there exists $F \in \mathcal{F}$ such that $f(a) \leq g(a)$ for every $g \in F$. By Proposition 4.35, it follows that $\text{depth}(a) = \text{depth}(f(a)) \leq \text{depth}(g(a)) = \text{depth}(a)$, which implies that $f(a) = g(a)$; thus $g = f$, from which we infer that $\mathcal{F} = [f]$. Therefore $\text{Aut}(\mathbb{T})$ is discrete. \square

Propositions 4.12 and 4.36 imply an equivalent condition for differentiability on \mathbb{T} .

Theorem 4.37. *If $L \in \mathcal{D}(\mathbb{T}, \mathbb{T}) = \text{Aut}(\mathbb{T})$, $f : \mathbb{T} \rightarrow \mathbb{T}$, and $a \in \mathbb{T}$, then L is a differential of f at a if and only if $L(b) = f(b)$ for every $b \geq a$.*

Proof. [Necessity]. Since $\overrightarrow{a} = \{b : b \geq a\}$, it follows that $L(b) = f(b)$ for every $b \in \overrightarrow{a}$. Therefore, by Proposition 4.12 we conclude that L is a differential of f at a .

[Sufficiency]. By Proposition 4.36, the space $\mathcal{D}(\mathbb{T}, \mathbb{T})$ is discrete. Thus, by Proposition 4.12, there exists a neighborhood U of a such that $L(x) = f(x)$ for every element x in U . Since $U \supseteq \overrightarrow{a} = \{b : b \geq a\}$, it follows that $L(b) = f(b)$ for every $b \geq a$. \square

4.5 Differential Calculus on Sequences

In this section, we use Definition 4.1 to develop differential calculi for real and binary sequences, that is, functions from the natural numbers to either the Euclidean line or a two-point set.

By \mathbb{N} we will denote the set of all non-negative integers together with an element ∞ that is strictly greater than every non-negative integer.⁶ Prior to any construction of a differential calculus for functions from \mathbb{N} , we must choose an appropriate convergence structure for \mathbb{N}

⁶As in Example 3.19, but not equipped with the same convergence structure.

from among the infinitely many convergence structures available. In classical analysis, a convergent sequence is a function from \mathbb{N} to \mathbb{R} that is continuous at ∞ . Thus, an ideal convergence structure for \mathbb{N} is one in which the convergent real sequences are precisely the real-valued functions continuous at ∞ .

An ostensibly natural choice is to equip \mathbb{N} with the subspace convergence structure induced by \mathbb{R} . Consider, in this case, the neighborhood filter \mathcal{N}_x of any point $x \in \mathbb{N}$. Since \mathbb{N} is pretopological, it follows that $\mathcal{N}_x \downarrow x$, which implies that $\iota(\mathcal{N}_x)$ converges to $\iota(x)$, and so $\mathcal{N}_{\iota(x)} \subseteq \iota(\mathcal{N}_x)$. Since $M = \{r \in \mathbb{R} : |r - \iota(x)| < 1\}$ is a neighborhood of $\iota(x)$, it follows that $M \in \mathcal{N}_{\iota(x)}$; thus, there exists $N \in \mathcal{N}_x$ such that $\iota(N) \subseteq M$. If $y \in N$, then $1 > |\iota(y) - \iota(x)| = |y - x|$; since x and y are both natural numbers, it follows that $x = y$. Thus $N = \{x\}$, which implies that $\mathcal{N}_x = [x]$. Therefore \mathbb{N} is discrete, and so every function from \mathbb{N} —not just the convergent sequences—is continuous.

Alternatively, we equip \mathbb{N} with the pretopology, the neighborhood filters of which are defined by

$$\mathcal{N}_n = \begin{cases} [m : m \geq n], & n \neq \infty; \\ \{F \subseteq \mathbb{N} : \infty \in F \wedge \mathbb{N} - F \text{ is finite}\}, & \text{otherwise.} \end{cases}$$

Throughout this section, we assume that \mathbb{N} is equipped with this convergence structure. Proposition 4.38 justifies this choice of convergence structure for \mathbb{N} : a function from \mathbb{N} into a pretopological space is continuous at ∞ if and only if it is a convergent sequence of that pretopological space.⁷

Proposition 4.38. *If X is a pretopological space and $f : \mathbb{N} \rightarrow X$ is a function, then $x \in \lim_{n \rightarrow \infty} f(n)$ if and only if for every neighborhood V of x , there exists $M \in \mathbb{N}$ such that $f(m) \in V$ for every $\infty > m \geq M$.*

Proof. [Necessity]. If V is a neighborhood of $f_{\infty \mapsto x}(\infty) = x$, then by hypothesis, there exists $M \in \mathbb{N}$ such that $f(m) \in V$ for every $\infty > m \geq M$. Let $U = \{m : m \geq M\}$.

⁷Cf. Definition 1.6.1 of [4].

Since U is cofinite, it belongs to the Fréchet filter; thus U is a neighborhood of ∞ such that $f_{\infty \rightarrow x}(U) \subseteq V$. Therefore, by Proposition 2.70, we conclude that $x \in \lim_{n \rightarrow \infty} f(n)$.

[Sufficiency]. Let V be a neighborhood of x . By Proposition 2.70, there exists a neighborhood U of ∞ such that $f_{\infty \rightarrow l}(U) \subseteq V$. Let $M = \max(\mathbb{N} - U) - 1$. Then $f(m) \in V$ for every $\infty > m \geq M$ \square

It is useful to formulate the convergence structure of \mathbb{N} in terms of ultrafilters.

Proposition 4.39. *If $a, b \in \mathbb{N}$, then $[a] \downarrow b$ if and only if $a \geq b$; free filters on \mathbb{N} converge to each point of \mathbb{N} .*

Proof. If $[a] \downarrow b$, then $\mathcal{N}_b \subseteq [a]$, and so for every $N \in \mathcal{N}_b$, there exists $A \in [a]$ such that $A \subseteq N$. In particular $a \in \{m : b \leq m\}$, which implies that $b \leq a$. Conversely, suppose that $b \leq a$. If $N \in \mathcal{N}_b$, then $\{m : n \leq m\} \subseteq N$ for some $n \leq b$. Since $b \leq a$, it follows that $n \leq a$, and so $a \in N$, which implies that $N \in [a]$. Thus $\mathcal{N}_b \subseteq [a]$, and so $[a] \downarrow b$.

Since every neighborhood of every point is cofinite, every neighborhood of every point belongs to the Fréchet filter, and thus to every free filter. Therefore, every free filter includes the neighborhood filter of each point, and so every free filter converges to every point. \square

4.5.1 Differentials of Real Sequences

In this section, we determine an equivalent condition for the differentiability of a real sequence. We first establish that the continuous functions \mathbb{N} to \mathbb{R} are precisely the constant functions. From this observation it follows that $\mathcal{C}(\mathbb{N}, \mathbb{R})$ is homeomorphic to \mathbb{R} ; thus $\mathcal{C}(\mathbb{N}, \mathbb{R})$ is a pretopological space.

Proposition 4.40. *If $f : \mathbb{N} \rightarrow \mathbb{R}$ is continuous at n , then for every $\varepsilon > 0$, there exists $M \in \mathbb{N}$ such that $|f(m) - f(n)| < \varepsilon$ for every $m \geq M$.*

Proof. Let $\varepsilon > 0$ and let $V = \{x : |x - f(n)| < \varepsilon\}$. Since V is a neighborhood of $f(n)$, by Proposition 2.66 there exists a neighborhood U of n such that $f(U) \subseteq V$. Since $U \supseteq \{m :$

$m \geq M\}$ for some $M \in \mathbb{N}$, it follows that there exists $M \in \mathbb{N}$ such that $|f(m) - f(n)| < \varepsilon$ for every $m \geq M$. \square

Proposition 4.41. *A function $f : \mathbb{N} \rightarrow \mathbb{R}$ is continuous if and only if it is constant.*

Proof. [Necessity]. Let m and n be distinct elements of \mathbb{N} . Let $\varepsilon > 0$. By Proposition 4.40, there exist M and N such that $|f(m') - f(m)| < \varepsilon/2$ and $|f(n') - f(n)| < \varepsilon/2$ for every $m' \geq M$ and $n' \geq N$. Let $K = \max(M, N)$. If $k \geq K$, then

$$\begin{aligned} |f(n) - f(m)| &= |(f(n) - f(k)) + (f(k) - f(m))| \\ &\leq |f(k) - f(n)| + |f(k) - f(m)| \\ &< \varepsilon. \end{aligned}$$

Since this argument holds for every $\varepsilon > 0$, it follows that $f(m) = f(n)$. Therefore f is constant.

[Sufficiency]. Constant functions are always continuous. \square

Proposition 4.42. $\mathcal{C}(\mathbb{N}, \mathbb{R})$ is homeomorphic to \mathbb{R} .

Proof. The desired result follows from Proposition 4.41 and the proof of Proposition 2.58. \square

Since the only members of $\mathcal{C}(\mathbb{N}, \mathbb{R})$ are the constant functions, which are distinguished only by comparison of their ranges, we choose $\mathcal{D}(\mathbb{N}, \mathbb{R})$ to be the entire space $\mathcal{C}(\mathbb{N}, \mathbb{R})$. Consequently, a real sequence is differentiable at some point n distinct from ∞ if and only if it is constant at all points greater than or equal to n .

Theorem 4.43. *If $f : \mathbb{N} \rightarrow \mathbb{R}$ is a function, then f is differentiable at n if and only if f is constant for every $m \geq n$.*

Proof. [Necessity]. By hypothesis, there exists $r \in \mathbb{R}$ such that $f(m) = r$ for every $m \geq n$. If U is a neighborhood of n , then U includes $\{m : m \geq n\}$. If V is a neighborhood of the

constant function $\bar{r} : \mathbb{N} \rightarrow \mathbb{R}$. Since \bar{r} belongs to each of its neighborhoods, it follows that \bar{r} is a differential of f at n .

[Sufficiency]. By hypothesis, there exists $r \in \mathbb{R}$ such that the constant function $\bar{r} : \mathbb{N} \rightarrow \mathbb{R}$ is a differential of f at n . If $\varepsilon > 0$, then $V = \{\bar{p} \in \mathcal{D}(\mathbb{N}, \mathbb{R}) : |p - r| < \varepsilon\}$ is a neighborhood of \bar{r} . Thus, there exists a neighborhood U of n such that for every $m \in U$, there exists $\bar{p} \in V$ such that $f(m) = p$. Since U includes $\{m : m \geq n\}$, it follows that $|f(m) - r| = |p - r| < \varepsilon$ for every $m \geq n$. Since this argument holds for every $\varepsilon > 0$, we conclude that $f(m) = r$ for every $m \geq n$. \square

Example 4.44. In particular, every real sequence is differentiable at ∞ . Thus, if (x_n) converges to x , then \bar{x} is a differential of (x_n) .

4.5.2 Differentials of Binary Sequences

There are three non-homeomorphic convergence structures on the set $\{0, 1\}$: the discrete, indiscrete, and Sierpiński convergence structures. We will analyze $\mathcal{D}(\mathbb{N}, \{0, 1\})$ for each of these convergence structures on $\{0, 1\}$. The next two theorems briefly dispose of the discrete and indiscrete cases.

Theorem 4.45. *If $\{0, 1\}$ is equipped with the discrete topology, $f : \mathbb{N} \rightarrow \{0, 1\}$ is a function, and $n \in \mathbb{N}$, then f is differentiable at n if and only if there exists an $M \leq n$ such that f is constant for every $m \geq M$.*

Proof. First, we show that $\mathcal{C}(\mathbb{N}, \{0, 1\})$ is discrete. If g is a continuous function, then $[g(m)] \downarrow g(0)$ for each $m \in \mathbb{N}$, and so $g(m) = g(0)$ for each $m \in \mathbb{N}$; thus g is constant. If $[\bar{0}]$ converges to $\bar{1}$, then $[0] = [\bar{0}] \cdot [0]$ converges to $\bar{1}(0) = 1$, in contradiction to the hypothesis that $\{0, 1\}$ is discrete; likewise $[\bar{1}]$ does not converge to $\bar{0}$. Thus $\mathcal{C}(\mathbb{N}, \{0, 1\})$ is discrete, and so by Proposition 2.85, it follows that $\mathcal{D}(\mathbb{N}, \{0, 1\})$ is discrete. By Proposition 4.12, it follows that f is differentiable at n if and only if there exists a neighborhood U of n such that f is constant for every element m in U . Since U is a neighborhood of n if and only if there exists

$M \leq n$ such that $\{m : m \geq M\} \subseteq U$, we conclude that f is differentiable at n if and only if there exists an $M \leq n$ such that f is constant for every $m \geq M$. \square

Theorem 4.46. *If $\{0, 1\}$ is equipped with the indiscrete topology, $f : \mathbb{N} \rightarrow \{0, 1\}$ is a function, and $n \in \mathbb{N}$, then f is differentiable at n if and only if there exists $M \leq n$ and $L \in \mathcal{D}(\mathbb{N}, \{0, 1\})$ such that $f(m) = L(m)$ for every $m \geq M$.*

Proof. By Proposition 2.86, the space $\mathcal{C}(\mathbb{N}, \{0, 1\})$ is indiscrete, and so by Proposition 2.85 the space $\mathcal{D}(\mathbb{N}, \{0, 1\})$ is indiscrete. Thus, by Propositions 2.84 and 4.5, it follows that f is differentiable at n if and only if there exists a neighborhood U of n such that for every element m in U , there exists a function $L \in \mathcal{D}(\mathbb{N}, \{0, 1\})$ such that $f(m) = L(m)$; equivalently, there exists $M \leq n$ and $L \in \mathcal{D}(\mathbb{N}, \{0, 1\})$ such that $f(m) = L(m)$ for every $m \geq M$. \square

We now attend to $\mathcal{D}(\mathbb{N}, S)$, in which S is the set $\{0, 1\}$ equipped with the Sierpiński convergence structure. Define the function $\mathbf{0} : \mathbb{N} \rightarrow S$ by $\mathbf{0}(n) = 0$ for every $n \in \mathbb{N}$; define for each $n \in \mathbb{N}$ the function $\mathbf{2}^{-n} : \mathbb{N} \rightarrow S$ by

$$\mathbf{2}^{-n}(m) = \begin{cases} 0, & m < n; \\ 1, & \text{otherwise.} \end{cases}$$

The function $\mathbf{0}$ and all functions of the form $\mathbf{2}^{-n}$ are continuous; in fact, these are the only functions from \mathbb{N} to S that are continuous.

Proposition 4.47. *A function $f : \mathbb{N} \rightarrow S$ is continuous if and only if it is either $\mathbf{0}$ or $\mathbf{2}^{-n}$ for some $n \in \mathbb{N}$.*

Proof. [Necessity]. Since constant functions are continuous, it suffices to consider $\mathbf{2}^{-n}$ for some $n \in \mathbb{N}$ distinct from 0. In view of Theorem 2.76, it suffices to consider the ultrafilters on \mathbb{N} . Let $m, k \in \mathbb{N}$ and suppose that $[m] \downarrow k$. If $m < n$, then $k < n$; thus $\mathbf{2}^{-n}([m]) = [0]$, which converges to $\mathbf{2}^{-n}(k) = 0$. If $m \geq n$, then $\mathbf{2}^{-n}([m]) = [1]$ which converges to $\mathbf{2}^{-n}(k)$. Finally, if \mathcal{U} is a free filter, then $\mathbf{2}^{-n}(\mathcal{U})$ must be $[1]$, which converges to both 0 and 1.

[Sufficiency]. Suppose that f is not constant. Let n be the smallest number for which f is 1. Since f is continuous, it follows that $f([m]) \downarrow f(n)$ for all $m \geq n$. Thus $f(m) = 1$ for all $m \geq n$. Therefore $f = \mathbf{2}^{-n}$. \square

Now that we know precisely which functions belong to $\mathcal{C}(\mathbb{N}, S)$, we can determine the convergence structure of $\mathcal{C}(\mathbb{N}, S)$. Since S is pseudotopological, it follows that $\mathcal{C}(\mathbb{N}, S)$ is also pseudotopological; in fact $\mathcal{C}(\mathbb{N}, S)$ is pretopological.

Proposition 4.48. *If \mathcal{U} is an ultrafilter on $\mathcal{C}(\mathbb{N}, S)$, then*

1. $\mathcal{U} \downarrow \mathbf{0}$.
2. $\mathcal{U} \downarrow \mathbf{2}^{-\infty}$ if and only if $\mathcal{U} \neq [\mathbf{0}]$.
3. If $n < \infty$, then $\mathcal{U} \downarrow \mathbf{2}^{-n}$ and if and only if $\mathcal{U} = [\mathbf{2}^{-m}]$ for some $m \leq n$.

Proof. Let \mathcal{A} converge to a in \mathbb{N} .

Since $\mathcal{U} \cdot \mathcal{A}$ converges to $0 = \mathbf{0}(a)$, it follows that $\mathcal{U} \downarrow \mathbf{0}$.

Suppose that $\mathcal{U} \neq [\mathbf{0}]$. If $\mathcal{A} \neq [\infty]$, then for each $A \in \mathcal{A}$, the set $\{1\} = \{\mathbf{2}^{-k} : k \geq \min(A)\} \cdot A$ belongs to $\mathcal{U} \cdot \mathcal{A}$, and so $\mathcal{U} \cdot \mathcal{A} = [1]$, which converges to $\mathbf{2}^{-n}(a)$; if $\mathcal{A} = [\infty]$, then the set $\{1\} = (\mathcal{C}(\mathbb{N}, S) - \{\mathbf{0}\}) \cdot \{\infty\}$ belongs to $\mathcal{U} \cdot [\infty]$, and so $\mathcal{U} \cdot \mathcal{A} = [1]$, which converges to $\mathbf{2}^{-n}(a)$. Conversely, if $\mathcal{U} \downarrow \mathbf{2}^{-\infty}$ but $\mathcal{U} = [\mathbf{0}]$, then $[0] = \mathcal{U} \cdot [\infty]$ converges to 1, which is absurd; thus $\mathcal{U} \neq [\mathbf{0}]$.

Let $n < \infty$. Suppose that $\mathcal{U} = [\mathbf{2}^{-m}]$ for some $m \leq n$. If $\mathbf{2}^{-n}(a) = 1$, then $n \leq a$, which implies that $\mathbf{2}^{-m}(a) = 1$, and so $\mathcal{U} \cdot \mathcal{A}$ converges to $\mathbf{2}^{-m}(a) = \mathbf{2}^{-n}(a)$; otherwise $\mathbf{2}^{-n}(a) = 0$, and so $\mathcal{U} \cdot \mathcal{A}$ converges to $\mathbf{2}^{-n}(a)$. Thus $\mathcal{U} \downarrow \mathbf{2}^{-n}$. Conversely, if $m > n$, then $[0] = [\mathbf{2}^{-m}] \cdot [n]$ converges to $\mathbf{2}^{-n}(n) = 1$, which is absurd; thus $[\mathbf{2}^{-m}]$ does not converge to $\mathbf{2}^{-n}$. If \mathcal{U} is a free filter, then $\{0\} = \{\mathbf{2}^{-m} : m > n\} \cdot \{n\} \in \mathcal{U} \cdot [n]$, which implies that $[0] = \mathcal{U} \cdot [n]$ converges to $\mathbf{2}^{-n}(n) = 1$, which is absurd. Therefore \mathcal{U} does not converge to $\mathbf{2}^{-n}$. \square

Proposition 4.49. *Let $f \in \mathcal{C}(\mathbb{N}, S)$.*

1. If $f = \mathbf{0}$, then $\mathcal{N}_f = [\mathcal{C}(\mathbb{N}, S)]$.
2. If $f = \mathbf{2}^{-\infty}$, then $\mathcal{N}_f = [\mathcal{C}(\mathbb{N}, S) - \{\mathbf{0}\}]$.
3. If $f = \mathbf{2}^{-n}$ and $n < \infty$, then $\mathcal{N}_f = [\mathbf{2}^{-m} : m \leq n]$.

Proof. Since every ultrafilter converges to $\mathbf{0}$, it follows that every filter converges to $\mathbf{0}$, and so $\mathcal{N}_{\mathbf{0}} = \mathcal{C}(\mathbb{N}, S)$.

Now suppose that $f = \mathbf{2}^{-\infty}$. If $\mathcal{C}(\mathbb{N}, S) - \{\mathbf{0}\}$ is not a neighborhood of f , then $f \in \text{cl}(\{\mathbf{0}\})$, which implies that $[\mathbf{0}] \downarrow f$, in contradiction to Proposition 4.48. If A is a neighborhood of f but $A \not\supseteq \mathcal{C}(\mathbb{N}, S) - \{\mathbf{0}\}$, then there exists $g \in (\mathcal{C}(\mathbb{N}, S) - \{\mathbf{0}\}) - A$. Since $g \neq \mathbf{0}$, it follows that $[g] \downarrow f$, and so $A \in [g]$, which implies that $g \in A$, which is absurd. Thus $\mathcal{N}_f = [\mathcal{C}(\mathbb{N}, S) - \{\mathbf{0}\}]$.

Finally, suppose that $f = \mathbf{2}^{-n}$ and $n < \infty$. If $\{\mathbf{2}^{-m} : m \leq n\}$ is not a neighborhood of f , then there exists a filter \mathcal{F} that converges to f and contains $\mathcal{C}(\mathbb{N}, S) - \{\mathbf{2}^{-m} : m \leq n\}$. If \mathcal{U} is an ultrafilter finer than \mathcal{F} , then $\mathcal{U} \downarrow f$, and so $\mathcal{U} = [\mathbf{2}^{-k}]$ for some $k \leq n$; thus $\mathbf{2}^{-k} \in \mathcal{C}(\mathbb{N}, S) - \{\mathbf{2}^{-m} : m \leq n\}$, which is absurd. If A is a neighborhood of f but $A \not\supseteq \{\mathbf{2}^{-m} : m \leq n\}$, then there exists $g \in \{\mathbf{2}^{-m} : m \leq n\} - A$, but then $[g] \downarrow f$, and so $A \in [g]$, which implies that $g \in A$, which is absurd; thus $\mathcal{N}_f = [\mathbf{2}^{-m} : m \leq n]$. \square

Proposition 4.50. *The convergence space $\mathcal{C}(\mathbb{N}, S)$ is pretopological.*

Proof. Since every ultrafilter converges to $\mathbf{0}$, it follows that every filter converges to $\mathbf{0}$; in particular $\mathcal{N}_{\mathbf{0}}$ converges to $\mathbf{0}$. If \mathcal{U} is an ultrafilter finer than $\mathcal{N}_{\mathbf{2}^{-\infty}}$, then by Proposition 4.49 it follows that $\mathcal{U} \neq [\mathbf{0}]$, and so $\mathcal{U} \downarrow \mathbf{2}^{-\infty}$; thus $\mathcal{N}_{\mathbf{2}^{-\infty}} \downarrow \mathbf{2}^{-\infty}$. If $f = \mathbf{2}^{-n}$ for some $n < \infty$ and \mathcal{U} is an ultrafilter finer than $\mathcal{N}_{\mathbf{2}^{-n}}$, then $\mathcal{U} = [\mathbf{2}^{-m}]$ for some $m \leq n$, and so $\mathcal{U} \downarrow \mathbf{2}^{-n}$; thus $\mathcal{N}_{\mathbf{2}^{-n}} \downarrow \mathbf{2}^{-n}$. Since the neighborhood filter of each point in $\mathcal{C}(\mathbb{N}, S)$ converges to that point, we conclude that $\mathcal{C}(\mathbb{N}, S)$ is pretopological. \square

In view of Proposition 4.50, we see that \mathbb{N} and $\mathcal{D}(\mathbb{N}, S)$ satisfy Proposition 4.5. Without imposing any constraints on $\mathcal{D}(\mathbb{N}, S)$, we determine whether a member of $\mathcal{C}(\mathbb{N}, S)$ is a differential of a given function at a particular point.

Theorem 4.51. *Let $f : \mathbb{N} \rightarrow S$ be a function and let $n \in \mathbb{N}$.*

1. *The constant function $\mathbf{0}$ is a differential of f at n .*
2. *The function $\mathbf{2}^{-\infty}$ is a differential of f at n if and only if $f \neq \mathbf{0}$.*
3. *If $j \neq \infty$, then the function $\mathbf{2}^{-j}$ is a differential of f at n if and only if $f(m) = 1$ for every $m \geq \max(n, j)$.*

Proof. By Proposition 4.5, we have that $\mathbf{0}$ is a differential of f at n if and only if for every neighborhood U of n , there exists $g \in \mathcal{C}(\mathbb{N}, S)$ such that $g(m) = f(m)$ for every $m \in U$. Since U is a neighborhood of n if and only if there exists $M \leq n$ such that $\{m : m \geq M\} \subseteq U$, it follows that $\mathbf{0}$ is a differential of f at n if and only if there exists $M \leq n$ and $g \in \mathcal{C}(\mathbb{N}, S)$ such that $g(m) = f(m)$ for every $m \geq M$. Since there always exists such a function g , it follows that $\mathbf{0}$ is a differential of f at n .

Likewise $\mathbf{2}^{-\infty}$ is a differential of f at n if and only if there exists $M \leq n$ and $g \in \mathcal{C}(\mathbb{N}, S) - \{\mathbf{0}\}$ such that $g(m) = f(m)$ for every $m \geq M$. Since such a function exists if and only if $f \neq \mathbf{0}$, the function $\mathbf{2}^{-\infty}$ is a differential of f at n if and only if $f \neq \mathbf{0}$.

If $j \neq \infty$, then $\mathbf{2}^{-j}$ is a differential of f at n if and only if there exists $M \leq n$ and $k \leq j$ such that $2^{-k}(m) = f(m)$ for every $m \geq M$. Let $M = \min(n, j)$. If $f(m) = 1$ for every $m \geq \max(n, j)$, then $\mathbf{2}^{-j}(m) = 1 = f(m)$ for every $m \geq M$; if $f(m) = 0$ for some $m \geq \max(n, j)$, then $\mathbf{2}^{-j}(m) \neq f(m)$ for every $m \geq M$. Thus the function $\mathbf{2}^{-j}$ is a differential of f at n if and only if $f(m) = 1$ for every $m \geq \max(n, j)$. \square

In view of Theorem 4.51, it is reasonable to delete $\mathbf{0}$ and $\mathbf{2}^{-\infty}$ from $\mathcal{D}(\mathbb{N}, S)$ regardless of any other constraints we wish to impose on $\mathcal{D}(\mathbb{N}, S)$.

Example 4.52. A binary sequence differentiable at a point need not be continuous at that same point. Consider the binary sequence $f : \mathbb{N} \rightarrow S$ defined by

$$f(m) = \begin{cases} 0, & m = 0 \text{ or } 2; \\ 1, & \text{otherwise.} \end{cases}$$

By Theorem 4.51, it follows that 2^{-3} is a differential of f at 1. But since f is 0 at 1, it is not continuous at 1.

Chapter 5

Mereology on Topological and Convergence Spaces

In 1981, Clarke [14] proposed an axiomatization of mereology—the study of the relationship between *part* and *whole*—with one primitive, *connection*: two regions are connected if they intersect. Some ten years later Randell, et al., [35] weakened this definition of connection by requiring only that the topological closures of the two regions intersect. Although Casati and Varzi [12] discuss several approaches to mereology, their “favored strategy” uses this definition of connection; in turn, they define parthood in terms of connection: one thing is a part of another if and only if anything connected to the former is also connected to the latter. In contrast, Guarino, et al., [20] used this definition of connection but chose, for reasons unclear, not to define parthood in terms of connection—an uneconomical choice since it increases the number of primitives in the theory.

Although those cited above have examined thoroughly the relational logic implied by the mereotopological axioms, they have neglected the topological consequences of these same axioms. In this chapter, we remedy this oversight in two ways. First, we determine the topological structures determined by the mereotopological axioms of [35]; in particular, we show that the definition of Randell, et al., is equivalent to the condition that a topological

space is discrete. Second, we generalize this result to the Cartesian closed category of convergence spaces, of which the category of topological spaces, is a full subcategory.

Conventional treatment of mereology begins by defining the binary relations connection, as in [35], and parthood, with the condition that parthood is a partial order. We take an ostensibly different, but in fact equivalent, approach: rather than explicitly require parthood to be a partial order, we define connection to be an extensional relation.¹

Definition 5.1. Let (X, \mathcal{T}) be a topological space. Let $|$ be a extensional binary relation on $2^X - \{\emptyset\}$ such that $A|B$ if and only if the topological closures of A and B intersect; formally, we write

$$(\forall A)(\forall B)(A|B \Leftrightarrow \text{cl}(A) \cap \text{cl}(B) \neq \emptyset). \quad (5.1)$$

Two subsets A and B of X are *connected* if and only if $A|B$. We refer to $|$ as the *connection* relation.

Proposition 5.2. *The connection relation is reflexive and symmetric.*

Proof. Suppose that A and B are nonempty subsets of a topological space X .

[Reflexivity]. Since $\text{cl}(A) \cap \text{cl}(A) = \text{cl}(A) \neq \emptyset$, it follows that $A|A$.

[Symmetry]. If $A|B$, then $\text{cl}(B) \cap \text{cl}(A) = \text{cl}(A) \cap \text{cl}(B) \neq \emptyset$, and thus $B|A$. \square

Connection is not an equivalence relation: it is not transitive. To see this, consider the following counterexample. Suppose that X has the discrete topology, and that $\emptyset \neq A \subset B \subseteq X$. Since every set of a discrete space is closed, it follows that $\text{cl}(B) \cap \text{cl}(X - A) = B \cap (X - A) \neq \emptyset$; thus $B|(X - A)$. Likewise $A|B$ because $\text{cl}(A) \cap \text{cl}(B) = A \cap B \neq \emptyset$. But it is not the case that $A|(X - A)$; otherwise $A \cap (X - A) = \text{cl}(A) \cap \text{cl}(X - A) \neq \emptyset$, which is absurd.

Definition 5.3. Let (X, \mathcal{T}) be a topological space. Let \leq be a relation on $2^X - \{\emptyset\}$ such that

$$(\forall A)(\forall B)(A \leq B \Leftrightarrow (\forall C)(C|A \rightarrow C|B)).$$

¹A relation R on a set S is *extensional* if and only if $a = b$ whenever $Rca \leftrightarrow Rcb$ for every a, b , and $c \in S$.

A subset A of X is a *part* of a subset B of X if and only if $A \leq B$. We refer to \leq as the *parthood* relation.

Proposition 5.4. *The parthood relation is a partial order.*

Proof. Suppose that A , B , and C are nonempty subsets of a topological space X .

[Reflexivity]. Since $(\forall C)(C|A \rightarrow C|A)$, it follows that $A \leq A$.

[Transitivity]. Let $A \leq B$ and $B \leq C$. If $D|A$, then by definition of parthood $D|B$, and so $D|C$. Therefore $A \leq C$.

[Antisymmetry]. Let $A \leq B$ and $B \leq A$. By definition of parthood, it follows that $(\forall C)(C|A \leftrightarrow C|B)$. Thus, by hypothesis, we have $A = B$. \square

Notice that connection need not be extensional for parthood to be reflexive and transitive; connection, however, must be extensional for parthood to be antisymmetric. To see this, suppose that the parthood relation is antisymmetric, but do not require that connection is extensional. Let A and B be subsets of a topological space X . Suppose that $(\forall C)(C|A \leftrightarrow C|B)$. By definition of parthood, it follows that $A \leq B$ and $B \leq A$. Thus, by hypothesis, we have $A = B$. Therefore, if we require parthood to be a partial order, then we must require connection to be extensional.

While it is conceivable that additional topological constraints could guarantee parthood to be antisymmetric, no separation axiom will suffice. To see this, suppose that connection is defined as above but without the requirement that it is extensional. Consider the standard topology on \mathbb{R} . Let $A = \{1/n : n \in \mathbb{Z}^+\}$ and let $B = A \cup \{0\}$. Then $\text{cl}(A) = B = \text{cl}(B)$, and so both $A \leq B$ and $B \leq A$. But $A \neq B$. Since \mathbb{R} is a T_6 space, we conclude that no separation axiom implies antisymmetry of parthood.

Definition 5.5. A *mereology* is a triple (X, \mathcal{T}, \leq) in which (X, \mathcal{T}) is a topological space and \leq is a parthood relation.² When no reasonable confusion is likely, we refer to (X, \mathcal{T}, \leq) by X .

²In view of Definition 5.1 and Proposition 5.4, the parthood relation must be a partial order. Moreover, the topological space (X, \mathcal{T}) completely determines (X, \mathcal{T}, \leq) .

We now ask two questions. Given a mereology (X, \mathcal{T}, \leq) , what is the structure of \mathcal{T} ? Conversely, which topological spaces are candidates for mereologies? To motivate an answer, we study a simple example.

Example 5.6. Consider $X = \{a, b, c\}$ and let (X, \mathcal{T}, \leq) be a mereology. Which of the nine nonhomeomorphic topologies on X can \mathcal{T} be?

If $\mathcal{T} = \{\emptyset, \{a\}, X\}$, then $\text{cl}(\{b\}) = \{b, c\} = \text{cl}(\{c\})$, and so for any $A \subseteq X$ both $A|\{b\}$ and $A|\{c\}$; thus $\{b\} \leq \{c\}$ and $\{c\} \leq \{b\}$, which implies that $\{b\} = \{c\}$, which is absurd.

Of the eight remaining topologies, we can apply minor variations of this *reductio ad absurdum* to all but the discrete topology, which is, as is easily verified, a mereology.

Not only can we generalize the argument given in Example 5.6 to show that every mereology is discrete, but we can also show that every discrete topological space is a mereology.

Theorem 5.7. *A topological space is a mereology if and only if it is discrete.*

Proof. Let X be a topological space. It suffices to show that the connection relation is extensional if and only if X is discrete.

[Necessity]. Let X have the discrete topology. Suppose that A and B are subsets of X and that $C|A$ if and only if $C|B$ for every subset C of X . If $x \in A$, then $\{x\} \subseteq A$. Since X is discrete, it follows that $\text{cl}(\{x\}) \cap \text{cl}(A) = \{x\} \cap A \neq \emptyset$; thus $\{x\}|A$, and so by hypothesis $\{x\}|B$, which implies that $\{x\} \cap B = \text{cl}(\{x\}) \cap \text{cl}(B) \neq \emptyset$, from which we infer that $x \in B$. An exactly similar argument shows that $x \in A$ whenever $x \in B$. Therefore $A = B$.

[Sufficiency]. Let the connection relation be extensional. Without loss of generality, suppose that A is a nonempty subset of X . Since $\text{cl}(\text{cl}(A)) \cap \text{cl}(A) = \text{cl}(A) \neq \emptyset$, it follows that $\text{cl}(A)|A$ and $A|\text{cl}(A)$, and so by hypothesis $\text{cl}(A) = A$. Thus every subset of X is closed, that is, as a topological space X is discrete. \square

In view of Theorem 5.7, we can simplify Definition 5.1 by replacing (5.1) with

$$(\forall A)(\forall B)(A|B \Leftrightarrow A \cap B \neq \emptyset).$$

As a consequence, we obtain a corollary to Theorem 5.7.

Corollary 5.8. *If A and B are nonempty subsets of a mereology X , then $A \leq B$ if and only if $A \subseteq B$.*

Proof. Let A and B be nonempty subsets of a topological space X .

[Necessity]. Suppose that $A \subseteq B$. If $C|A$ for some subset C of X , then $C \cap B \supseteq C \cap A \neq \emptyset$.

Thus $C|B$, and therefore $A \leq B$.

[Sufficiency]. Suppose that $A \leq B$. If $x \in A$, then $\{x\} \subseteq A$, and so $\{x\} \cap A \neq \emptyset$, which implies that $\{x\}|A$; thus $\{x\}|B$, from which it follows that $\{x\} \cap B \neq \emptyset$, and so $x \in B$.

Therefore, we conclude that $A \subseteq B$. \square

Corollary 5.8 shows that the parthood relation is equivalent to the subset relation, and thereby effectively demonstrates the equivalence of mereology and set theory. We can obtain a stronger result than this, namely the equivalence of general extensional mereology and set theory. To proceed, we present several definitions and intermediate results.

Definition 5.9. Let X be a mereology. Let $<$ be a binary relation on $2^X - \{\emptyset\}$ such that

$$(\forall A)(\forall B)(A < B \Leftrightarrow (A \leq B \wedge A \neq B)).$$

A subset A of X is a *proper part* of a subset B of X if and only if $A < B$. We refer to $<$ as the *proper parthood* relation.

Proposition 5.10. *The proper parthood relation is asymmetric and transitive.*

Proof. Let A , B , and C be subsets of a mereology X .

[Asymmetry]. To the contrary, suppose that $A < B$ and $B < A$. By definition of proper parthood, it follows that $A \leq B$, $B \leq A$, and $A \neq B$. By antisymmetry of parthood, however, we see that $A = B$, in contradiction to the former result. Therefore, we conclude that $B \not< A$.

[Transitivity]. Suppose that $A < B$ and $B < C$. By definition of proper parthood, it follows that $A \leq B$, $B \leq C$, $A \neq B$, and $B \neq C$. By transitivity of parthood, we see that $A \leq C$. If $A = C$, then $B \leq A$, and so antisymmetry of parthood implies that $A = B$, in contradiction to the result that $A \neq B$. Thus $A \neq C$, and therefore $A < C$. \square

Definition 5.11. Let X be a mereology. Let \odot be a binary relation on $2^X - \{\emptyset\}$ such that

$$(\forall A)(\forall B)(A \odot B \Leftrightarrow (\exists C)(C \leq A \wedge C \leq B)).$$

Two subsets A and B of X *overlap* if and only if $A \odot B$. We refer to \odot as the *overlap* relation.

While reflexive and symmetric, the overlap relation is not transitive. In view of Theorem 5.7, it follows that two sets overlap if and only if they intersect.

Definition 5.12. Let A and B be subsets of a mereology X . The *weak supplementation principle* asserts that if A is a proper part of B then there is a part of B not overlapping A . Formally, we write

$$(\forall A)(\forall B)(A < B \rightarrow (\exists C)(C \leq B \wedge \neg(C \odot A))).$$

The *strong supplementation principle* asserts that unless B is a part of A there is a part of B not overlapping A . Formally, we write

$$(\forall A)(\forall B)(\neg(B \leq A) \rightarrow (\exists C)(C \leq B \wedge \neg(C \odot A))).$$

Proposition 5.13. *The strong supplementation principle implies the weak supplementation principle.*

Proof. Suppose that the strong supplementation principle holds in a mereology X . Let A and B be subsets of X such that $A < B$. By definition of proper parthood, it follows that

$A \leq B$ and $A \neq B$. By the antisymmetry of parthood $\neg(B \leq A)$. By hypothesis, therefore, we conclude that there exists a subset C of X such that $C \leq B$ and $\neg(C \odot A)$. \square

Observe that the proof of Proposition 5.13 does not require connection to be extensional. If, on the other hand, connection is extensional, then the strong supplementation principle will hold in every mereology.

Proposition 5.14. *The strong supplementation principle holds in every mereology.*

Proof. Let X be a mereology and suppose that $\neg(B \leq A)$ for some nonempty subsets A and B of X . By Corollary 5.8, it follows that $\neg(\text{cl}(B) \subseteq \text{cl}(A))$, which implies that there exists an $x \in X$ such that $x \in \text{cl}(B)$ and $x \notin \text{cl}(A)$. By the former conjunct we have $\text{cl}(\{x\}) \subseteq \text{cl}(B)$, and so by Corollary 5.8 again we infer that $\{x\} \leq B$; by the latter conjunct and Theorem 5.7, we have $\text{cl}(\{x\}) \cap \text{cl}(A) = \{x\} \cap \text{cl}(A) = \emptyset$, and so $\neg(\{x\} | A)$, which implies that $\neg(\{x\} \odot A)$. \square

Definition 5.15. The *general sum* of all sets satisfying a predicate ϕ is the set

$$\Sigma\phi = \iota A \forall B (A \odot B \leftrightarrow \exists C (\phi C \wedge C \odot B)). \quad (5.2)$$

Proposition 5.16. *Let X be a mereology. If $\Sigma\phi$ exists, then it is the supremum of $\Phi = \{A \subseteq X : \phi A\}$.*

Proof. First, we establish that $U = \Sigma\phi$ is an upper bound of Φ . By (5.2), we have that

$$U \odot C \leftrightarrow \exists B (\phi B \wedge B \odot C). \quad (5.3)$$

Suppose that ϕS holds. An instance of the contrapositive of the strong supplementation principle is

$$(\forall A)(A \leq S \rightarrow A \odot U) \rightarrow S \leq U. \quad (5.4)$$

If $A \leq S$, then $A \odot S$, and so by (5.3) it follows that $A \odot U$; thus by (5.4) we infer that $S \leq U$.

Now we show that U is the least upper bound of Φ . Let V be any other upper bound of Φ . An instance of the strong supplementation principle is

$$\neg(U \leq V) \rightarrow (\exists A)(A \leq U \wedge \neg A \odot V).$$

Thus, if $\neg(U \leq V)$, then there exists a A such that $A \leq U$ and $\neg(A \odot V)$. The former conjunct implies that $A \odot U$, and so by (5.3) there exists B such that ϕB holds and $B \odot A$. Since $B \odot A$, there exists E such that $E \leq B$ and $E \leq A$; because V is a least upper bound of Φ , it follows that $B \leq V$, and so $E \leq V$ and $E \leq A$, which implies that $A \odot V$, in contradiction to $\neg(A \odot V)$. Therefore, we conclude that $U \leq V$. \square

Finally, we define a general extensional mereology (sometimes referred to as **GEM** in the literature). If connection is extensional, then any mereology is a general extensional mereology.

Definition 5.17. A *general extensional mereology* is a mereology X in which the weak supplementation principle holds and $\Sigma\phi$ exists for every relation ϕ on X .

Theorem 5.18. *Every mereology is a general extensional mereology.*

Proof. Let X be a mereology. By Theorem 5.7, it follows that X is a discrete topological space. Since the supremum of $\{A \subseteq X : \phi A\}$ is $\bigcup_{\phi A} A$, which always exists, it suffices to show that $\bigcup_{\phi A} A$ satisfies Definition 5.15.

If $\bigcup_{\phi A} A$ and B overlap, then there exists $C \subseteq X$ that is a subset of both $\bigcup_{\phi A} A$ and B . Now C intersects some A for which ϕA holds. Since $A \cap C$ is a subset of both A and B , it follows that $A \cap C$ is a part of both A and B , and so A and B overlap.

Conversely, suppose that $B \subseteq X$ and B overlaps with some $C \subseteq X$ such that ϕC . Since B intersects C , it follows that B intersects $\bigcup_{\phi A} A$, and so B and $\bigcup_{\phi A} A$ overlap. \square

Theorem 5.7 together with Corollary 5.8 and Theorem 5.18 show that general extensional mereology reduces to set theory. To prevent this reduction, we might relax the definition of connection; it is not obvious, however, in which way we should do this: for we cannot remove extensionality from connection without simultaneously removing antisymmetry from parthood. Alternatively, we observe that the category of topological spaces is a full subcategory of the category of convergence spaces; subsequently, we can amend the definition of mereology to be over convergence spaces rather than topological spaces. Since topological spaces are “coarser” structures than convergence spaces, the “finer” structure of the latter might prevent reduction of general extensional mereology to set theory not unlike the way that increasing microscopic resolution allows the observation of formerly indistinguishable features. Theorem 5.19, however, precludes this possibility.

Theorem 5.19. *A convergence space is a mereology if and only if it is discrete.*

Proof. Let X be a convergence space. It suffices to show that the connection relation is extensional if and only if X is discrete. The argument is exactly similar to the proof of Theorem 5.7. □

If we generalize the definition of connection from topological spaces to convergence spaces, then all previous results of this chapter hold *mutatis mutandis*.³

We have investigated mereology⁴ on topological and convergence spaces. Theorem 5.7 establishes that every mereology, as defined by Definition 5.5, is a discrete topological space; likewise, Theorem 5.19 shows that every mereology defined on a convergence space must be a discrete convergence space. The latter result allows to generalize Theorem 5.18 to convergence spaces: thus, every mereology defined on a convergence space is a general extensional

³There is at least one other approach to this generalization of connection from topological spaces to convergence spaces. Instead of using limit spaces, we could have used general convergence spaces and, in place of Theorem 5.19, prove that a convergence space is a mereology if and only if it is postdiscrete and pretopological (see [8]).

⁴In particular, we investigated mereology with a connection relation, often referred to as *mereotopology*. Since we also extended the definition of connection to convergence spaces, we were reluctant to use that term, while hoping that context could resolve any ambiguities.

mereology, that is, a mereology in which the weak supplementation principle holds and the general sum exists for each predicate of the mereology.

A topological space is simply a subcollection of the power set of the underlying set; a discrete topological space, then, is the entire power set. Thus, our work demonstrates that general extensional mereology is indistinguishable from set theory—not a surprising result when one considers that general extensional mereology was developed in pursuit of an alternative to set theory. We attribute this reduction (of general extensional mereology to set theory) to the extensionality imposed on connection (or, equivalently and more conventionally, the antisymmetry imposed on parthood): Theorem 5.19 shows that the reduction is not a topological consequence since it also occurs in the category of convergence spaces, of which the category of topological spaces is a full subcategory.

One possible objection to our work is that it assumes the elements of a mereology are sets. While we concede this point, we also answer that the assumption arises not from any peculiarity of our work but from the mereotopological axioms in [35]. In particular, the axioms require that the topological closures of two connected regions intersect; but the only objects that have topological closures are indeed sets.

Chapter 6

A Universal Homogeneous Pretopological Space

In 1954, Shimrat [39] showed that every topological space can be embedded in a homogeneous topological space. Uspenskii [41] later provided an alternative construction also establishing this result. Since the category of topological spaces is a full subcategory of the category of convergence spaces, it is reasonable to ask whether we can generalize the work of Shimrat or Uspenskii to convergence spaces. Indeed, Blair, et al., [8] recently generalized Shimrat's construction to convergence spaces. In Section 6.1, we adapt to convergence spaces Uspenskii's technique of embedding an arbitrary topological space into a homogeneous space.

The existence of universal topological spaces is also well-known: for example, in 1956, Mrówka [29] provided several examples of universal spaces. By contrast, universal convergence spaces are largely unexplored. In Section 6.2, we use the technique developed in Section 6.1 to construct a universal homogeneous pretopological space—that is, a homogeneous pretopological space with the property that any other pretopological space embeds in some product of it—and we determine that this space is neither compact nor separable.

6.1 Uspenskiĭ Spaces

In this section, we construct a homogeneous convergence space, which we call an Uspenskiĭ space. We show that an Uspenskiĭ space preserves several properties—in particular, the separation properties—of the underlying space. First, we present some relevant terminology.

Definition 6.1. Let X be a convergence space; let A be an infinite set with cardinality at least $|X|$. Give X^A the product convergence structure. The *Uspenskiĭ space on X with respect to A* is the subspace

$$U = \{u \in X^A : (\forall x)(x \in X \Rightarrow |u^{-1}(x)| = |A|)\}.$$

Example 6.2. Let us examine a simple example of an Uspenskiĭ space. Define a convergence structure on $X = \{0, 1, 2\}$ by the equivalences

$$\mathcal{F} \downarrow 0 \text{ if and only if } \{0, 1\} \in \mathcal{F}$$

$$\mathcal{F} \downarrow 1 \text{ if and only if } \{1, 2\} \in \mathcal{F}$$

$$\mathcal{F} \downarrow 2 \text{ if and only if } \{0, 2\} \in \mathcal{F}$$

The Uspenskiĭ space U of X with respect to \mathbb{N} is the set of all tuples that have countably many 0's, 1's, and 2's. Observe that U is not closed in $X^{\mathbb{N}}$. To see this, consider the point filter $\mathcal{P} = [(1, 2, 0, 1, 2, 0, \dots)]$. Since $\{(1, 2, 0, 1, 2, 0, \dots)\} \subseteq U$, it follows that $U \in \mathcal{P}$. Now $\pi_{3n-2}(\mathcal{P}) = [1]$, $\pi_{3n-1}(\mathcal{P}) = [2]$, and $\pi_{3n}(\mathcal{P}) = [0]$ for every $n \geq 1$. Since $[1] \downarrow 1$, $[2] \downarrow 0$, and $[0] \downarrow 1$, it follows that \mathcal{P} converges to $(1, 0, 1, 1, 0, 1, \dots)$, which does not belong to U . Thus $(1, 0, 1, 1, 0, 1, \dots)$ belongs to $\text{cl}(U) - U$. Therefore U is not closed.

We wish to show that convergence space can be embedded into a homogeneous convergence space. We shall see that Uspenskiĭ spaces are homogeneous and admit the desired embedding. First, we prove that any Uspenskiĭ space is homogeneous. To do so, we must show that given any two elements of an Uspenskiĭ space, there exists an automorphism that

maps one of the elements to the other.

Lemma 6.3. *Let U be the Uspenskiĭ space on X with respect to A . For every two elements u and v of U , there exists a bijection $\sigma : A \rightarrow A$ such that $u = v \circ \sigma$.*

Proof. If u and v belong to U , then $|u^{-1}(x)| = |v^{-1}(x)|$ for each $x \in X$, and so there exists a bijection $\sigma_x : u^{-1}(x) \rightarrow v^{-1}(x)$ for each $x \in X$. The function $\sigma : A \rightarrow A$, defined by $\sigma(a) = \sigma_x(a)$ if and only if $a \in u^{-1}(x)$, is a bijection such that $\sigma(u^{-1}(x)) = v^{-1}(x)$ for each $x \in X$. Thus $\sigma^{-1}(v^{-1}(x)) = u^{-1}(x)$ for each $x \in X$. If $(v \circ \sigma)(a) = y$, then $v(\sigma(a)) = y$, and so $\sigma(a) \in v^{-1}(y)$, which implies that $\sigma(a) \in \sigma(u^{-1}(y))$. Hence there exists $b \in u^{-1}(y)$ such that $\sigma(b) = \sigma(a)$. Since σ is a bijection, it follows that $b = a$, and so $a \in u^{-1}(y)$, which implies that $u(a) = y$. Thus $v \circ \sigma = u$. \square

Lemma 6.4. *Let U be the Uspenskiĭ space on X with respect to A . For every two elements u and v of U , there exists an automorphism τ on U such that $\tau(v) = u$.*

Proof. If u and v belong to U , then by Lemma 6.3 there exists a bijection $\sigma : A \rightarrow A$ such that $u = v \circ \sigma$. Define the function $\tau : U \rightarrow U$ by $\tau(w) = w \circ \sigma$. Note that $\tau(v) = u$. We will show that τ is an automorphism on U .

To see that τ is bijective, note that if $\tau(w_1) = \tau(w_2)$, then $w_1 \circ \sigma = w_2 \circ \sigma$, and so $w_1 = w_1 \circ \sigma \circ \sigma^{-1} = w_2 \circ \sigma \circ \sigma^{-1} = w_2$, which establishes that τ is injective. Moreover, if $w \in U$, then $\tau(w \circ \sigma^{-1}) = w$, which implies that τ is surjective.

Finally, we show that τ is bicontinuous. Let $\iota : U \rightarrow X^A$ be the inclusion map. For each projection $\pi_a : \prod_{a \in A} X \rightarrow X$ both $(\pi_a \circ \iota \circ \tau)(w) = (\pi_{\sigma(a)} \circ \iota)(w)$ and $(\pi_a \circ \iota \circ \tau)(\mathcal{F}) = (\pi_{\sigma(a)} \circ \iota)(\mathcal{F})$. Thus, if $\mathcal{F} \downarrow w$, then $(\pi_a \circ \iota)(\mathcal{F}) \downarrow (\pi_a \circ \iota)(w)$ for each $a \in A$, and so $(\pi_{\sigma^{-1}(a)} \circ \iota \circ \tau)(\mathcal{F}) \downarrow (\pi_{\sigma^{-1}(a)} \circ \iota \circ \tau)(w)$ for each $a \in A$, which implies that $\tau(\mathcal{F}) \downarrow \tau(w)$. Conversely, if $\tau(\mathcal{F}) \downarrow \tau(w)$, then $(\pi_a \circ \iota \circ \tau)(\mathcal{F}) \downarrow (\pi_a \circ \iota \circ \tau)(w)$ for each $a \in A$, and so $(\pi_{\sigma(a)} \circ \iota)(\mathcal{F}) \downarrow (\pi_{\sigma(a)} \circ \iota)(w)$ for each $a \in A$, which implies that $\mathcal{F} \downarrow w$. \square

Lemma 6.5. *Every Uspenskiĭ space is homogeneous.*

Proof. From Lemmas 6.3 and 6.4 the desired result immediately follows. \square

Since an Uspenskii space is simply a subspace of a product space, its construction preserves any separation of the underlying space. Because the convergence structure of an Uspenskii space is initial, an Uspenskii space is topological, pretopological, or pseudotopological whenever the underlying space is also topological, pretopological, or pseudotopological.

If X and Y are homeomorphic spaces and A is a set with cardinality at least the cardinality of X , then the Uspenskii spaces of X and Y with respect to A are homeomorphic.

Proposition 6.6. *If X and Y are homeomorphic convergence spaces, then there exist homeomorphic Uspenskii spaces U_X and U_Y of X and Y , respectively.*

Proof. Let A be some set such that $|A| \geq \max(|X|, |Y|)$. Let U_X and U_Y be the Uspenskii spaces of X and Y with respect to A , let $\pi_a : \prod_{a \in A} X \rightarrow X$ and $\rho_a : \prod_{a \in A} Y \rightarrow Y$ be projections, and let $\iota : U_X \rightarrow \prod_{a \in A} X$ and $\kappa : U_Y \rightarrow \prod_{a \in A} Y$ be inclusion maps. Define $\hat{f} : U_X \rightarrow U_Y$ by $\hat{f}(x_1, x_2, \dots) = (f(x_1), f(x_2), \dots)$. It is clear that \hat{f} is bijective.

To see that \hat{f} is continuous, suppose that $\mathcal{F} \downarrow (x_1, x_2, \dots)$. By construction $(\pi_a \circ \iota)(\mathcal{F}) \downarrow x_a$ for each $a \in A$, and so by hypothesis $f((\pi_a \circ \iota)(\mathcal{F})) \downarrow f(x_a)$ for each $a \in A$. Since $f \circ (\pi_a \circ \iota) = (\rho_a \circ \kappa) \circ \hat{f}$, it follows that $(\rho_a \circ \kappa)(\hat{f}(\mathcal{F})) \downarrow f(x_a)$ for each $a \in A$; equivalently $(\rho_a \circ \kappa)(\hat{f}(\mathcal{F})) \downarrow (\rho_a \circ \kappa)(\hat{f}(x_1, x_2, \dots))$ for each $a \in A$; thus $\hat{f}(\mathcal{F}) \downarrow \hat{f}(x_1, x_2, \dots)$.

To see that \hat{f}^{-1} is continuous, it suffices to show that $\mathcal{F} \downarrow (x_1, x_2, \dots)$ whenever $\hat{f}(\mathcal{F}) \downarrow \hat{f}(x_1, x_2, \dots)$. If $\hat{f}(\mathcal{F}) \downarrow \hat{f}(x_1, x_2, \dots)$, then $(\rho_a \circ \kappa)(\hat{f}(\mathcal{F})) \downarrow (\rho_a \circ \kappa)(\hat{f}(x_1, x_2, \dots))$ for each $a \in A$, which implies that $f((\pi_a \circ \iota)(\mathcal{F})) \downarrow f(x_a)$ for each $a \in A$, and so by hypothesis $(\pi_a \circ \iota)(\mathcal{F}) \downarrow x_a$ for each $a \in A$; thus $\mathcal{F} \downarrow (x_1, x_2, \dots)$. \square

We conclude this section with its main result, namely that every convergence can be embedded into its Uspenskii space, and hence into a homogeneous convergence space.

Theorem 6.7. *Every convergence space can be embedded into its Uspenskii space, and hence into a homogeneous convergence space.*

Proof. Let U be the Uspenskii space on a convergence space X with respect to A . Lemma 6.5 guarantees that U is homogeneous. Now we claim that the function $e : X \rightarrow U$, defined by $e(x) = (x, u_1, u_2, \dots)$ for some constant $u = (u_1, u_2, \dots) \in U$, is an embedding. Since $x = u_i$ for infinitely many i , it follows that $e(x) \in U$. If $e(x_1) = e(x_2)$, then $(x_1, u_1, u_2, \dots) = (x_2, u_1, u_2, \dots)$, and so $x_1 = x_2$; thus e is injective. Also, note that if $\mathcal{A} \in \Phi(X)$, then $(\pi_0 \circ e)(\mathcal{A}) = \mathcal{A}$ and $(\pi_a \circ e)(\mathcal{A}) = [u_a]$ for each $a \in A$. Thus, if $\mathcal{A} \downarrow x$, then $(\pi_0 \circ e)(\mathcal{A}) \downarrow x$ and $(\pi_a \circ e)(\mathcal{A}) \downarrow u_a$ for each $a \in A$, which implies that $e(\mathcal{A}) \downarrow e(x)$, and so e is continuous. Conversely, if $e(\mathcal{A}) \downarrow e(x)$, then $\mathcal{A} = (\pi_0 \circ e)(\mathcal{A})$ also converges to x ; thus e^{-1} is also continuous. \square

6.2 Construction and Analysis

In this section, we construct a universal space, and then homogenize it with the technique developed in the previous section.

Definition 6.8. A convergence space U is *universal* if and only if for every convergence space X , there exists an index set I such that $\prod_I U$ embeds X .

A universal convergence space cannot have any properties preserved under initial constructions. Thus, it cannot be separated nor can it pretopological. It is the latter constraint that is the most severe: since every finite space is pretopological, a universal convergence space—if one exists—cannot be finite. This condition is prohibitive. Nevertheless, we can construct a finite universal pretopological space; any Uspenskii space of it will be a universal homogeneous pretopological space.

Lemma 6.9. *If U is the reflexive digraph of Figure 6.1, then U is a universal pretopological space.*

Proof. First observe that U is not topological since it is not transitive. Now let X be a pretopological space. For each $x \in X$ and each neighborhood N_x of x , define the function




Figure 6.1: A Universal Pretopological Space.

$f_{N_x} : X \rightarrow U$ by

$$f_{N_x}(y) = \begin{cases} a, & y \in N_x - \{x\}; \\ c, & y = x; \\ b, & \text{otherwise.} \end{cases}$$

Each f_{N_x} is continuous. To see this, let $y \in X$ and let V be a neighborhood of $f_{N_x}(y)$. It suffices to show that there exists a neighborhood of y for which its image under f_{N_x} is a subset of V . If $y \in N_x - \{x\}$ or $y \notin N_x$, then $V = U$; thus $f_{N_x}(W) \subseteq V$ for any neighborhood W of y . If $y = x$, then $f_{N_x}(y) = c$, and so V is either $\{a, c\}$ or U ; thus N_x is a neighborhood of y such that $f_{N_x}(N_x) = \{a, c\} \subseteq V$.

Define the function $F : X \rightarrow \prod_{\mathcal{C}(X,U)} U$ by $\pi_f(F(x)) = f(x)$ for each $f \in \mathcal{C}(X, U)$. We will show that F is an embedding.

To see that F is injective, suppose that $F(x) = F(y)$ for some $x, y \in X$. Thus $f(x) = \pi_f(F(x)) = \pi_f(F(y)) = f(y)$ for each $f \in \mathcal{C}(X, U)$; in particular $f_{N_x}(x) = f_{N_x}(y)$ for any neighborhood N_x of x , and so $f_{N_x}(y) = c$, which implies that $x = y$.

To see that F is continuous, suppose that $x \in X$. It suffices to show that $F(\mathcal{N}_x) \downarrow F(x)$. By construction $f(\mathcal{N}_x) \downarrow f(x)$ for each $f \in \mathcal{C}(X, U)$; thus $\pi_f(F(\mathcal{N}_x)) \downarrow \pi_f(F(x))$ for each $f \in \mathcal{C}(X, U)$, from which the desired result follows.

To see that F^{-1} is continuous on $F(X)$, suppose that $u \in F(X)$ and V is a neighborhood of $F^{-1}(u)$. It suffices to show that there exists a neighborhood W of u such that $W \subseteq F(V)$. Since $\pi_{f_V}(u) = \pi_{f_V}(F(F^{-1}(u))) = f_V(F^{-1}(u)) = c$, the only neighborhoods of $\pi_{f_V}(u)$ are $\{a, c\}$ and U . Since π_{f_V} is continuous, there exists a neighborhood W of u such that

$\pi_{f_V}(W) \subseteq \{a, c\}$. If $w \in W$, then there exists $x \in X$ such that $F(x) = w$, from which it follows that $f_V(x) = \pi_{f_V}(F(x)) = \pi_{f_V}(w)$, and so $f_V(x) \in \{a, c\}$, which implies that $x \in V$, from which we infer that $w = F(x) \in F(V)$; thus $W \subseteq F(V)$. \square

Although we independently arrived at Lemma 6.9, we subsequently discovered that Bourdaud [10], and more recently Herrlich, et al., [22], anticipated this result by some 40 years. Nevertheless, it is apparently absent from the literature that any pretopological space can be embedded in a homogeneous pretopological space.¹

Theorem 6.10. *Any Uspenskiĭ space of U is a universal homogeneous pretopological space.*

Proof. If X is a pretopological space, then there exists an embedding $F : X \rightarrow \prod_I U$ for some index set I . Let H be an Uspenskiĭ space of U . Since H is homogeneous and pretopological, it follows that $\prod_I H$ is homogeneous and pretopological. By Proposition 2.56, there exists an embedding $G : \prod_I U \rightarrow \prod_I H$. Since $G \circ F : X \rightarrow \prod_I H$ is an embedding, we conclude that H is a universal pretopological homogeneous space. \square

We will denote by V the Uspenskiĭ space of the universal pretopological space U with respect to \mathbb{N} . To determine precisely the convergence structure of V it suffices to determine the neighborhoods of points in V .

Proposition 6.11. *For each $p \in V$, the neighborhood filter of p is $[M]$, in which $M = \prod_{i \in \mathbb{N}} M_i$ and*

$$M_i = \begin{cases} \{a, c\}, & \pi_i(p) = c; \\ U, & \text{otherwise.} \end{cases}$$

Proof. If $\pi_i(p) = c$, then $\pi_i([M]) = [\pi_i(M)] = [M_i] = [a, c]$, which converges to c ; otherwise $\pi_i([M]) = [\pi_i(M)] = [U]$, which converges to a and b . Thus $[M]$ converges to p , and so $\mathcal{N}_p \subseteq [M]$. For the reverse inclusion, suppose that $N \supseteq M$ and \mathcal{F} converges to p in V . Since

¹In [8], however, it is claimed that any convergence space can be embedded in a homogeneous convergence space.

$M_i \in \pi_i(\mathcal{F})$ for each $i \in \mathbb{N}$, it follows that for every $i \in \mathbb{N}$, there exists $F \in \mathcal{F}$ such that $\pi_i(F) \subseteq M_i$; thus $F \subseteq M \subseteq N$, and so $N \in \mathcal{F}$, which implies that $M \in \mathcal{N}_p$. \square

The final propositions of this chapter establish that the space V is neither compact nor Kolmogorov.

Proposition 6.12. *The space V is not compact.*

Proof. Let $p \in V$ and let \mathcal{F} denote the Fréchet filter. By construction of V , there exists $j \in \mathbb{N}$ such that $\pi_j(p) = c$. Define $N = \prod_{i \in \mathbb{N}} N_i$ and

$$N_i = \begin{cases} \{a, c\}, & i = j; \\ U, & \text{otherwise.} \end{cases}$$

By Proposition 6.11, it follows that N is a neighborhood of p . Since $V - N$ is infinite, it follows that $N \notin \mathcal{F}$, which implies that \mathcal{F} does not converge to p ; thus, there exists an ultrafilter finer than \mathcal{F} that does not converge to p . Therefore V is not compact. \square

Proposition 6.13. *The space V is not Kolmogorov.*

Proof. If p is a point of V , then q defined by

$$\pi_i(q) = \begin{cases} b, & \pi_i(p) = a; \\ a, & \pi_i(p) = b; \\ c, & \pi_i(p) = c. \end{cases}$$

for each $i \in \mathbb{N}$ is also a point of V distinct from p . By Proposition 6.11, it follows that $\mathcal{N}_p = \mathcal{N}_q$, and so by Theorem 2.101 we conclude that V is not Kolmogorov. \square

Chapter 7

Conclusion and Future Work

We investigated several problems in the theory of convergence spaces: generalization of Kolmogorov separation from topological spaces to convergence spaces, representation of reflexive digraphs as convergence spaces, construction of differential calculi on convergence spaces, mereology on convergence spaces, and construction of a universal homogeneous pretopological space. First, we generalized Kolmogorov separation from topological spaces to convergence spaces; we then studied properties of Kolmogorov spaces. Second, we developed a theory of reflexive digraphs as convergence spaces, which we then specialized to Cayley graphs. Third, we conservatively extended the concept of differential from the spaces of classical analysis to arbitrary convergence spaces; we then used this extension to obtain differential calculi for finite convergence spaces, finite Kolmogorov spaces, finite groups, Boolean hypercubes, labeled graphs, the Cantor tree, and real and binary sequences. Fourth, we showed that a standard axiomatization of mereology is equivalent to the condition that a topological space is discrete, and consequently, any model of general extensional mereology is indistinguishable from a model of set theory; we then generalized these results to the cartesian closed category of convergence spaces. Finally, we showed that every convergence space can be embedded into a homogeneous convergence space; we then used this result to construct a universal homogeneous pretopological space.

Above¹ we addressed the implications of our work on mereology; here we conclude by attending to several interesting questions arising from our work on differential calculus. First, we obtained a conservative extension of the concept of *differential* from classical analysis to arbitrary convergence spaces; yet we have not developed a conservative extension of *derivative*. Although it is clear what the concept of *derivative* must involve—the derivative of $f : X \rightarrow Y$ is the function $D_f : X \rightarrow \mathcal{D}(X, Y)$ defined by for each $a \in X$ the differential of f at a is $D_f(a)$ —it is a concept rife with difficulties in the context of arbitrary convergence spaces: not only is it computationally daunting but it is not guaranteed to be a function. Second, the observation that identification of a group with one of its non-redundant Cayley graphs is analogous to selection of a basis for a vector space leads to the question of how “change of Cayley graph” affects differentiation—and more fundamentally, continuity—in groups. Third, while it is evident that practical application of our work on differential calculus on convergence spaces is impeded, in general, by the computational complexity inherent in the continuous convergence structure, there is potential to mechanize—in the spirit of classical analysis—specific differential calculi such as that on Boolean hypercubes, which is essentially a matter of solving matrix equations. Fourth, our work shows that differentiability is fundamentally a structural—as opposed to quantitative—property. Can we abstract the structural essence of other analytic concepts² and thereby obtain conservative extensions of those concepts to arbitrary convergence spaces?

These questions, together with the results presented here, are part of our larger project to develop analysis on convergence spaces and related structures. Our project is both iconoclastic and radical.³ It is iconoclastic insofar that it effectively repeals the bankrupt dichotomization of mathematics into the discrete and continuous. From Graham, et al., [18], we hear that “[d]espite all the parallels between continuous and discrete math, some continuous notions have no discrete analog.” Lovász [25] perpetuates this myth by asserting

¹See Chapter 5.

²For example, curvature.

³As was our project to consider seriously the topological consequences of the mereotopological axioms of [35].

the existence of “a rather clear dividing line between discrete and continuous mathematics.” Such unprincipled statements gloss over the fundamental nature of continua. Our project, however, is a radical investigation of continua—radical, in the sense of going to the root, and the root, of course, is mathematical structure. The structure of continuity *is* the structure of homomorphism: preservation of filter convergence. Thus we obtain a unification, a *synthesis*, of ostensibly disparate concepts: the continuous and the discrete. From this correct understanding of continuity we obtain a correct understanding of differentiability, which in turn forms the foundation of discrete analysis.

Bibliography

- [1] Kazuyuki Aihara and Hideyuki Suzuki. Theory of hybrid dynamical systems and its applications to biological and medical systems. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 368(1930):4893–4914, 2010.
- [2] Andrej Bauer, Lars Birkedal, and Dana S. Scott. Equiological spaces. *Theoret. Comput. Sci.*, 315(1):35–59, 2004.
- [3] Fülöp Bacsó and Elemér Lábos. Boolean-Lie algebras and the Leibniz rule. *J. Phys. A*, 39(22):6871–6876, 2006.
- [4] R. Beattie and H.-P. Butzmann. *Convergence structures and applications to functional analysis*. Kluwer Academic Publishers, Dordrecht, 2002.
- [5] E. Binz and H. H. Keller. Funktionenräume in der Kategorie der Limesräume. *Ann. Acad. Sci. Fenn. Ser. A I No.*, 383:21, 1966.
- [6] Ernst Binz. Recent results in the functional analytic investigations of convergence spaces. In *General topology and its relations to modern analysis and algebra, III (Proc. Third Prague Topological Sympos., 1971)*, pages 67–72. Academia, Prague, 1972.
- [7] Ernst Binz. *Continuous convergence on $C(X)$* . Springer-Verlag, Berlin, 1975. Lectures Notes in Mathematics, Vol. 469.
- [8] Howard A. Blair, David W. Jakel, Robert J. Irwin, and Angel Rivera. Elementary differential calculus on discrete and hybrid structures. In *Logical foundations of computer*

- science*, volume 4514 of *Lecture Notes in Comput. Sci.*, pages 41–53. Springer, Berlin, 2007.
- [9] Nicolas Bourbaki. *Elements of mathematics. General topology. Part 1*. Hermann, Paris, 1966.
- [10] Gérard Bourdaud. Espaces d’Antoine et semi-espaces d’Antoine. *Cahiers Topologie Géom. Différentielle*, 16(2):107–133, 1975.
- [11] Henri Cartan. Théorie des filtres. *C. R. Acad. Sci., Paris*, 205:595–598, 1937.
- [12] Roberto Casati and Achille C. Varzi. *Parts and places: the structures of spatial representation*. MIT Press, Cambridge, 1999.
- [13] G. Choquet. Convergences. *Ann. Univ. Grenoble. Sect. Sci. Math. Phys. (N.S.)*, 23:57–112, 1948.
- [14] Bowman L. Clarke. A calculus of individuals based on “connection”. *Notre Dame J. Formal Logic*, 22(3):204–218, 1981.
- [15] P. J. Daniell. The modular difference of classes. *Bull. Amer. Math. Soc.*, 23(10):446–450, 1917.
- [16] D. Dikranjan and W. Tholen. *Categorical structure of closure operators*, volume 346 of *Mathematics and its Applications*. Kluwer Academic Publishers Group, Dordrecht, 1995. With applications to topology, algebra and discrete mathematics.
- [17] H. R. Fischer. Limesräume. *Math. Ann.*, 137:269–303, 1959.
- [18] Ronald L. Graham, Donald E. Knuth, and Oren Patashnik. *Concrete mathematics*. Addison-Wesley Publishing Company, Advanced Book Program, Reading, MA, 1989. A foundation for computer science.

- [19] Nicola Guarino. Formal ontology and information systems. In *Formal Ontology in Information Systems: Proceedings of the First International Conference (FOIS'98), June 6-8, Trento, Italy*, Frontiers in artificial intelligence and applications. IOS Press, 1998.
- [20] Nicola Guarino, Massimiliano Carrara, and Pierdaniele Giaretta. Formalizing ontological commitments. In *Proceedings of the twelfth national conference on Artificial intelligence (vol. 1)*, AAAI '94, pages 560–567, Menlo Park, CA, USA, 1994. American Association for Artificial Intelligence.
- [21] Reinhold Heckmann. A non-topological view of deposes as convergence spaces. *Theoret. Comput. Sci.*, 305(1-3):159–186, 2003. Topology in computer science (Schloß Dagstuhl, 2000).
- [22] H. Herrlich, E. Lowen-Colebunders, and F. Schwarz. Improving Top: PrTop and PsTop. In *Category theory at work (Bremen, 1990)*, volume 18 of *Res. Exp. Math.*, pages 21–34. Heldermann, Berlin, 1991.
- [23] John L. Kelley. *General topology*. D. Van Nostrand Company, Inc., Toronto-New York-London, 1955.
- [24] Hans-Joachim Kowalsky. Limesräume und Komplettierung. *Math. Nachr.*, 12:301–340, 1954.
- [25] László Lovász. Discrete and continuous: two sides of the same? *Geom. Funct. Anal.*, (Special Volume, Part I):359–382, 2000. GAFA 2000 (Tel Aviv, 1999).
- [26] Saunders Mac Lane. *Categories for the working mathematician*, volume 5 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, second edition, 1998.
- [27] Maurice Mashaal. *Bourbaki*. American Mathematical Society, Providence, RI, 2006. A secret society of mathematicians, Translated from the 2002 French original by Anna Pierrehumbert.

- [28] Russell Merris. *Combinatorics*. Wiley-Interscience Series in Discrete Mathematics and Optimization. Wiley-Interscience [John Wiley & Sons], New York, second edition, 2003.
- [29] S. Mrówka. On universal spaces. *Bull. Acad. Polon. Sci. Cl. III.*, 4:479–481, 1956.
- [30] James R. Munkres. *Topology: a first course*. Prentice-Hall Inc., Englewood Cliffs, N.J., 1975.
- [31] Frédéric Mynard. An introduction to convergence spaces. 2005.
- [32] L. D. Nel. Topological universes and smooth Gel'fand-Naïmark duality. In *Mathematical applications of category theory (Denver, Col., 1983)*, volume 30 of *Contemp. Math.*, pages 244–276. Amer. Math. Soc., Providence, RI, 1984.
- [33] Daniel R. Patten. Mereology on topological and convergence spaces. *Notre Dame J. Form. Log.*, 54(1):21–31, 2013.
- [34] Gerhard Preuß. Semiuniform convergence spaces and filter spaces. In *Beyond topology*, volume 486 of *Contemp. Math.*, pages 333–373. Amer. Math. Soc., Providence, RI, 2009.
- [35] David A. Randell, Zhan Cui, and Anthony G. Cohn. A spatial logic based on regions and connection. In *Principles of Knowledge Representation and Reasoning*, pages 165–176, 1992.
- [36] Joseph J. Rotman. *An introduction to the theory of groups*, volume 148 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, fourth edition, 1995.
- [37] Sergiu Rudeanu. What can be expected from a Boolean derivative? *An. Ştiinţ. Univ. "Ovidius" Constanţa Ser. Mat.*, 17(1):177–186, 2009.
- [38] Walter Rudin. *Principles of mathematical analysis*. McGraw-Hill Book Co., New York, third edition, 1976. International Series in Pure and Applied Mathematics.

- [39] M. Shimrat. Embedding in homogeneous spaces. *Quart. J. Math., Oxford Ser. (2)*, 5:304–311, 1954.
- [40] Patrick Suppes. *Axiomatic set theory*. Dover Publications Inc., New York, 1972. Unabridged and corrected republication of the 1960 original with a new preface and a new section (8.4).
- [41] Vladimir V. Uspenskiĭ. For any X , the product $X \times Y$ is homogeneous for some Y . *Proc. Amer. Math. Soc.*, 87(1):187–188, 1983.
- [42] Arjan van der Schaft and Hans Schumacher. *An introduction to hybrid dynamical systems*, volume 251 of *Lecture Notes in Control and Information Sciences*. Springer-Verlag London, Ltd., London, 2000.
- [43] Stephen Willard. *General topology*. Dover Publications Inc., Mineola, NY, 2004. Reprint of the 1970 original [Addison-Wesley, Reading, MA; MR0264581].

Biographical Data

Name of Author: Daniel R. Patten

Place of Birth: Utica, New York

Date of Birth: November 28, 1979

Graduate and Undergraduate Schools Attended:

- Utica College of Syracuse University, Utica NY
- Syracuse University, Syracuse NY

Degrees Awarded:

- Bachelor of Arts, *magna cum laude*, Mathematics, Utica College of Syracuse University, 5/01.
- Master of Science, Mathematics, Syracuse University, 5/04.
- Master of Science, Computer Science, Syracuse University, 12/13.