Discrete Interpretations of Continuous Representations in Model-Theoretic Computational Phonology

Workshop on The Role of Representation in Computational Phonology

Scott Nelson

University of Illinois Urbana-Champaign

January 13, 2026

Download the slides!

Background

- There have been various proposals for the use of continuous representations within phonological theory.
- Here, I will focus on instances like subfeatural representations and gradient symbolic representations which both maintain phonological feature categories but change from discrete to continuous valuations of said features.

Lionnet (2017); Smolensky et al. (2014)

Background

- One reason for pursuing the continuous approach to phonological features is that it provides a representational account for exceptional behavior at both the phonemic and lexical level.
- These types of representational assumptions appear to be in conflict with the branch of computational phonology that explores the expressivity of phonological patterns.
- While this type of work can be given a gradient interpretation, it has been exclusively explored in terms of output structure and not input structure.

Heinz (2018); Chandlee and Heinz (2017)

Background

- The consequences of continuous *input* representations for this type of work have not yet been explored.
- Here, I show that there may be no conflict at all when viewing discrete category labels as something like equivalence classes over a continuous feature space.

Overview

- I discuss one way in which continuous phonological feature representations may be interpreted as discrete feature values under the purview of model-theoretic phonology.
- This approach assumes a threshold oriented mapping from continuous feature values to binary feature categories.
- Crucial to this approach is the idea that the threshold need not be the same for the positive and negative values.

Three types of patterns

The result of this assumption is that three types of patterns are predicted to emerge: a standard two-way split, a three-way split with elements that act as *both* positive and negative valued, and a three-way split with elements that act as *neither* positive nor negative valued.

Model-Theoretic Phonology

Model Signatures

A **model signature** $\mathcal{M}$ is a collection of symbols for the functions, relations, and constants that describe structures.

- This provides the ingredients for formally defining phonological representations.
- For example, the signature $\langle \triangleleft, \{R_\sigma \mid \sigma \in \Sigma\} \rangle$ gives us two types of operations.
  - The *ordering relation* $\triangleleft$ lets us linearly order individual elements.
  - The *labeling relations* R_σ allows us to give a label to an element where the labels are drawn from a finite set Σ.

Structures

An $\mathcal{M}$**-structure**, A, contains a set called the **domain**, as well as **denotations**.

- The domain is typically an initial sequence of the natural numbers: $\{1,2,\ldots\}$.
- A denotation takes the domain and says which elements of the domain satisfy which properties.

Structures

An $\mathcal{M}$**-structure**, A, contains a set called the **domain**, as well as **denotations**.

- Given a set of symbols $\Sigma = \{$æ,t,k$\}$, the structure for the word *cat* is $\langle \mathcal{D} = \{1, 2, 3\}; \lhd = \{(1, 2), (2, 3)\}; \{R_{\text{æ}} = 2, R_{\text{t}} = 3, R_{\text{k}} = 1\}\rangle$.
- We can also display this graphically:

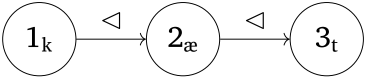

Logical Languages

A **logical language** in logic x is defined by combining the symbols of that logic with a specific model signature $\mathcal{S}$. Today I will use some variant of predicate (FO) logic.

- The model signature provides the tools to formalize representations.
- The logic provides a way to do inference over the representations that are built.

Determining static properties of a structure

The model-theoretic approach provides a way for identifying substructures with logic. This is useful when thinking about static phonotactic or morpheme structure constraints.

- Does structure *A* contain the substring *cad*?
- $\varphi := \exists x[a(x) \wedge \exists y[y \lhd x \wedge c(y)] \wedge \exists z[x \lhd z \wedge d(z)]]$

Interpretations

An **interpretation** of structure A in terms of structure B is a function denoted by a set of n formulas $\{\phi_i, \ldots, \phi_n\}$ where n is equal to the number of functions, relations, and constants in A's model signature.

- A formula $\phi_P(x) := Q(x)$ denotes that domain element x has property P in the output structure if it has property Q in the input structure.

Courcelle (1994); Engelfriet and Hoogeboom (2001)

Transductions

These types of logical interpretations can be used to define phonological input-output maps (transductions) such as **a** $\rightarrow$ **b/c_d**.

- $$\begin{aligned} \phi_a(x) &:= a(x) \wedge \neg\exists y,z[y \triangleleft x \triangleleft z \wedge c(y) \wedge d(z)] \\ \phi_b(x) &:= b(x) \vee (a(x) \wedge \exists y,z[y \triangleleft x \triangleleft z \wedge c(y) \wedge d(z)]) \\ \phi_c(x) &:= c(x) \\ \phi_d(x) &:= d(x) \end{aligned}$$
- 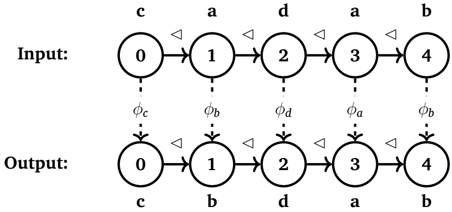

Strother-Garcia (2019); Chandlee and Jardine (2021); Nelson (2024); Heinz (Forthcoming)

Another way that interpretations have been used is to study the similarity between different types of proposed phonological representation schemes.

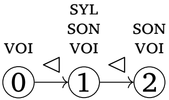

$$\phi_{\texttt{voi}}(x) \stackrel{\text{def}}{=} d(x) \vee n(x) \vee a(x)$$

$$\phi_{\texttt{son}}(x) \stackrel{\text{def}}{=} n(x) \vee a(x)$$

$$\phi_{\texttt{syl}}(x) \stackrel{\text{def}}{=} a(x)$$

$$\phi_d(x) \stackrel{\text{def}}{=} \texttt{voi}(x) \wedge \neg\texttt{son}(x) \wedge \neg\texttt{syl}(x)$$

$$\phi_n(x) \stackrel{\text{def}}{=} \texttt{voi}(x) \wedge \texttt{son}(x) \wedge \neg\texttt{syl}(x)$$

$$\phi_a(x) \stackrel{\text{def}}{=} \texttt{voi}(x) \wedge \texttt{son}(x) \wedge \texttt{syl}(x)$$

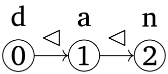

Strother-Garcia (2019); Oakden (2020); Jardine et al. (2021); Nelson (2022); Danis (2025); Nelson (2024)

A Proposal...

Continuous Feature Signatures

Suppose we adopt a signature for defining phonological representations that maintains a linear ordering relation, but replaces the labeling relations with a labeling functions which give each element a value for a feature between 0 and 1. We might also want to add a function for comparing real numbered values as well.

- $\mathcal{M} = \langle \lhd, \{f_\varphi \mid \varphi \in \mathcal{F}\}, \leq \rangle$
- $f_\varphi : \mathcal{D} \to \mathbb{R} \in [0, 1]$
- $\leq : (\mathbb{R}, \mathbb{R}) \to \{\texttt{TRUE}, \texttt{FALSE}\}$

An Aside About Discrete Feature Signatures

I am going to assume that the set of available labels are *valued* features as it gives us a better understanding of underspecification in logical approaches.

- $\{R_\sigma \mid \sigma \in \Sigma = \{+\texttt{voi}, -\texttt{voi}, +\texttt{son}, \ldots\}\}$

Nelson (2022)

Since $\leq$ is in the logical language for the continuous structures, this can be used for interpretation of discrete structures.

- Given thresholds i and j, we can write logical statements that interpret a continuous structure as a discrete structure:
  - $\phi_{-\varphi} := f_{\varphi}(x) \leq i$
  - $\phi_{+\varphi} := j \leq f_{\varphi}(x)$

Three Predictions

There are various ways to give semantic meaning to the (i, j) pair, but I'm interested in the three types of patterns which emerge.

- $i = j$
- $i < j$ (underspecification)
- $j < i$ (controversial???)

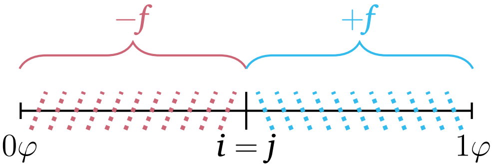

Three Predictions

There are various ways to give semantic meaning to the (i, j) pair, but I'm interested in the three types of patterns which emerge.

- $i = j$
- $i < j$ (underspecification)
- $j < i$ (controversial???)

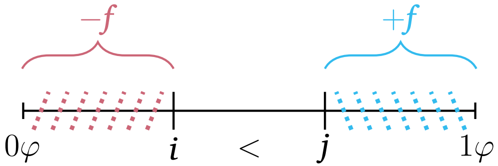

Three Predictions

There are various ways to give semantic meaning to the (i, j) pair, but I'm interested in the three types of patterns which emerge.

- $i = j$
- $i < j$ (underspecification)
- $j < i$ (controversial???)

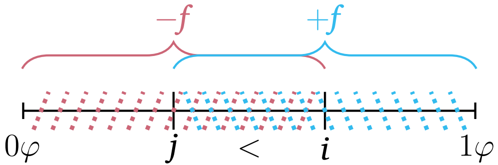

Empirical Phenomena

Rounding Harmony in Laal

- Simple rounding harmony in pronominal suffixes.
- "Doubly triggered" rounding harmony in number suffixes.

a.	/pɨ́-rù/	[púrù]	'catch her'
b.	/kɨ́r-ùn/	[kúrùn]	'place her'
c.	/də̀g-òn/	[dògòn]	'drag her'
d.	/də̀g-nŭ/	[dògnŭ]	'drag us'
e.	/ɓɨ̀r-ú/	[ɓùrú]	'fishhook-PL'
f.	/gɨ́n-ù/	[gɨ́nù]	'net-PL'
g.	/sə̀g-ó/	[sə̀gó]	'tree sp.-PL'
h.	/də́n-ú/	[də́nú]	'tree sp.-PL'

Lionnet (2017)

Rounding Harmony in Laal

- The “doubly triggered” harmony only occurs when the target is adjacent to a labial consonant and shares the same backness and height features with the trigger.
- Lionnet proposes that this type of harmony requires a continuous value of [`round`] *above some threshold* in order to occur.

Lionnet (2017)

Rounding Harmony in Laal

Suppose a predicate P indicates the triggering condition for simple harmony and a predicate Q indicates the triggering condition for the "double triggering" harmony.

- $\phi_u(x) := (P(x) \wedge \texttt{round}(x) \leq i) \vee (Q(x) \wedge j \leq \texttt{round}(x))$
- Discrete interpretation: simple harmony targets [-round] vowels and doubly triggered harmony targets [+round] vowels.
- The vowels which undergo the latter process should also undergo the former which is exactly what is the case (pp. 533–534).

Lionnet (2017)

Rounding Harmony in Laal

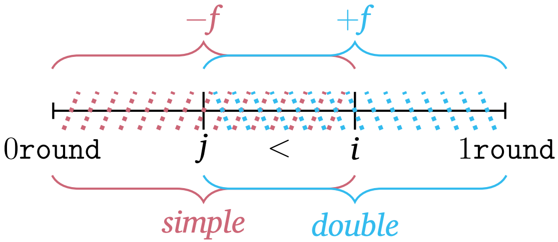

Lionnet (2017)

Rendaku Voicing in Japanese

In Japanese, the initial consonant of the second member of a compound is often voiced upon morphological concatenation. This process is referred to as rendaku and has long been known to be an exceptional process.

a.	/kuma + te/	[kumade]	‘rake’ (*bear* + *hand*)
b.	/yama + te/	[yamate]	‘mountainside’ (*mountain* + *hand*)
c.	/yama + tori/	[yamadori]	‘mountain bird’
d.	/niwa + tori/	[niwatori]	‘chicken’ (*garden* + *bird*)

Recent(ish) overview for generative phonology: Kawahara and Zamma (2016)

Rendaku Voicing in Japanese

- Rendaku as a morphological juncture feature of noun noun compounds.
- The voice feature can also be made continuous.
- Corpus study finds that there are some N1 and N2s which never voice as well as some N1 and N2s which always voice.
- There is a third set which sometimes voice and sometimes do not voice when put together.

Rosen (2016); Itô and Mester (1995)

One interpretation of these data is that the nouns which never trigger voicing are *below* some threshold ($-f$) while those that always trigger voicing are *above* some threshold ($+f$).

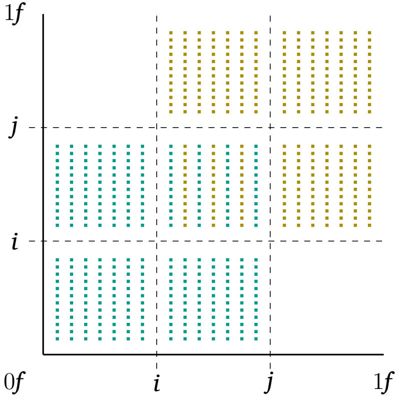

Japanese Rendaku Voicing

Variable behavior arises in compounds where both nouns are in between the thresholds (underspecified). The variation therefore becomes externally conditioned rather than internally conditioned.

- $\phi_{-\texttt{voice}}(x) := R(x,y) \land \texttt{voice}(x) \leq i \lor \texttt{voice}(y) \leq i$
- $\phi_{+\texttt{voice}}(x) := R(x,y) \land j \leq \texttt{voice}(x) \lor j \leq \texttt{voice}(y)$

Japanese Rendaku Voicing

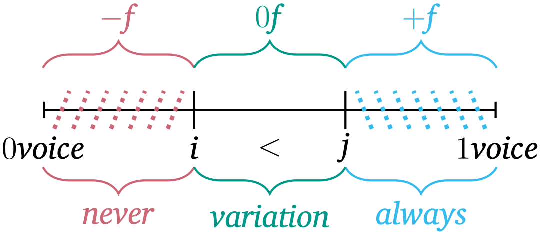

Concluding Thoughts

Conclusion

- I have sketched out how continuous feature spaces may be given a discrete interpretation using model theory.
- It suggests computational analyses of expressivity can be viewed as operating over equivalence class labels.
- Further issues such as *what features should represent* and *how are features phonetically interpreted* are important and should be thought of as next steps in the goal of relating discrete and continuous theories of phonological features.

THANK YOU!

Bibliography I

Chandlee, J. and Heinz, J. (2017). Computational phonology. In Oxford Research Encyclopedia of Linguistics.

Chandlee, J. and Jardine, A. (2021). Computational universals in linguistic theory: Using recursive programs for phonological analysis. Language, 97(3):485–519.

Courcelle, B. (1994). Monadic second-order definable graph transductions: a survey. Theoretical Computer Science, 126(1):53–75.

Danis, N. (2025). Logical transductions are not sufficient for notational equivalence. In Proceedings of the Annual Meetings on Phonology.

Engelfriet, J. and Hoogeboom, H. J. (2001). MSO definable string transductions and two-way finite-state transducers. ACM Transactions on Computational Logic (TOCL), 2(2):216–254.

Heinz, J. (2018). The computational nature of phonological generalizations. In Hyman, L. and Plank, F., editors, Phonological Typology, Phonetics and Phonology, chapter 5, pages 126–195. De Gruyter Mouton.

Heinz, J. (Forthcoming). Doing Computational Phonology. Oxford University Press.

Itô, J. and Mester, A. (1995). The core-periphery structure of the lexicon and constraints on reranking. University of Massachusetts Occasional Working Papers in Linguistics (UMOP).

Jardine, A., Danis, N., and Iacoponi, L. (2021). A formal investigation of q-theory in comparison to autosegmental representations. Linguistic Inquiry, 52(2):333–358.

Kawahara, S. and Zamma, H. (2016). Generative treatments of rendaku and related issues. In Sequential Voicing in Japanese, pages 13–34. John Benjamins Publishing Company.

Lionnet, F. (2017). A theory of subfeatural representations: The case of rounding harmony in laal. Phonology, 34(3):523–564.

Nelson, S. (2022). A model theoretic perspective on phonological feature systems. In Proceedings of the Society for Computation in Linguistics 2022, pages 1–10.

Nelson, S. (2024). The Computational Structure of Phonological and Phonetic Knowledge. PhD thesis, Stony Brook University.

Oakden, C. (2020). Notational equivalence in tonal geometry. Phonology, 37(2):257–296.

Rosen, E. (2016). Predicting the unpredictable: Capturing the apparent semi-regularity of rendaku voicing in japanese through harmonic grammar. In Proceedings of BLS, volume 42, pages 235–249. Berkeley Linguistic Society Berkeley, CA.

Smolensky, P., Goldrick, M., and Mathis, D. (2014). Optimization and quantization in gradient symbol systems: A framework for integrating the continuous and the discrete in cognition. Cognitive science, 38(6):1102–1138.

Strother-Garcia, K. (2019). Using model theory in phonology: a novel characterization of syllable structure and syllabification. PhD thesis, University of Delaware.