Finley, S. (2009). Directionality in vowel harmony. In H. Lehnert-LeHouillier and A.B Fine (Eds.), *University of Rochester Working Papers in the Language Sciences*, 5(1), 69-88.

Directionality in Vowel Harmony

Sara Finley[*]
University of Rochester

This paper presents an account of directionality in vowel harmony using Turbid Spreading (Finley, 2008, in preparation), a theory of vowel harmony for use with Optimality Theory (Prince and Smolensky, 1993/2004). In this theory, vowel features are represented on three levels of representation: an underlying form, a hidden abstract projection representation, and a phonetic surface form. These three levels capture the abstract nature of vowel harmony, which is induced by a directional SPREAD constraint. This SPREAD constraint captures both general directional spreading (e.g., SPREAD-R[αF]) and dominant-recessive spreading (e.g., SPREAD[+ATR]), which can be either directional or non-directional.

1. Introduction

In this paper, I extend previous work on Turbidity Theory (Goldrick, 1999, Goldrick, 2001) with a novel proposal for representations in vowel harmony. These representations can be used to account for various types of directional processes in vowel harmony, including feature-specific (dominant-recessive) non-directional harmony processes and directional feature-general spreading processes. The paper is structured into three parts. In the first part, I present the proposed theory of representations for Turbid Spreading, presenting the requirements on GEN that constrain the representations. In the second part, I present the constraints that induce the optimal representations for vowel harmony, including the constraints that derive both directional non-directional vowel harmony. Finally, I demonstrate the implementation of vowel harmony using the proposed constraints. This includes an account of non-participating (opaque and transparent) vowels.

2. Turbid Spreading

The Turbid Spreading theory for vowel harmony (Finley, 2008, Finley, in preparation) is based on the idea that spreading is an abstract process that takes place at a featural level. Vowel harmony is seen as a process where vowels share a particular feature value. Thus,

[*] Portions of this paper were presented at the 2008 and 2009 Linguistic Society of America Annual Meetings. I am very grateful to Colin Wilson, Paul Smolensky, Luigi Burzio, Jason Riggle and Eric Bakovic for helpful comments, discussions and suggestions. All errors are my own.

representations for vowels are concentrated to the value for the harmonic feature value (e.g., [HIGH] for height harmony, [ATR] for ATR harmony)[1]. The feature value for each segment is represented in terms of a triple: underlying form: projection form: surface form. All segments have a projection value (with the exception of some epenthetic vowels). The projection is interpreted differently from Goldrick's (1999, 2001) original formulation, in which each vowel feature has a pronunciation representation, and must also be licensed by a projection. In Turbid Spreading, all features have a projection for each feature value, which are each represented on a separate tier. This creates three levels of representation: the underlying form, the surface/phonetic form, and an intermediate, projection/phonological level, described in (1) below.

(1) Three Levels of Representation

/[+ATR]/	Underlying Form
[+ATR]	Projection/Phonological Level
[+ATR]	Surface/Phonetic Level

The underlying form is the standard formalization of the underlying representation: the input to the phonological optimization; it is found in the lexicon. The surface/phonetic form represents what is actually pronounced by the speaker, and is the representation subject to phonetic interpretation. The projection/phonological level is an intermediate representation that represents phonological processes, such as spreading. In Turbid Spreading, each feature value is licensed by a single projection, situated on the intermediate level of representation. There are three potential sources for the projection of a given vowel. The first is that of the underlying form of that particular segment, representing a phonologically unchanged (faithful) representation, depicted with a down arrow (🡣) in (2).

(2) Projection from the Underlying form

/[+ATR]/	Underlying Form
🡣	
[+ATR]	Projection/Phonological Level
[+ATR]	Surface/Phonetic Level

The second source of the projection value comes from a neighboring projection value, representing phonological spreading, depicted as a horizontal arrow (🡠, 🡢) in (3).

[1] While many harmony processes rely on interactions of multiple features, for demonstrative purposes, I assume that harmony takes place with a single harmony feature value.

(3) Projection from the Projection Level form

/[+ATR]/	Underlying Form
[+ATR] ← [+ATR]	Projection/Phonological Level
[+ATR]	Pronunciation/Surface Level

The surface/phonetic is also a potential source for the projection value, representing a phonetically induced change to the abstract representation, depicted as an up arrow (↑) in (4).

(4) Projection from the Pronunciation/Surface Level

/[+ATR]/	Underlying Form
[+ATR]	Projection/Phonological Level
↑	
[+ATR]	Pronunciation/Surface Level

A feature at the projection/phonological level must only have one projection, and all segments with an underlying representation must have a projection (epenthetic segments need not be projected). In (5), the medial vowel's [ATR] feature is projected by multiple segments, which is not produced by GEN.

(5) Banned Projection: Multiple Projections

* /[+ATR] [–ATR] [+ATR]/	Underlying Form
↓ ↓	
[+ATR] → [+ATR] ← [+ATR]	Projection/Phonological Level
[+ATR] [+ATR] [+ATR]	Pronunciation/Surface Level

The source of the projection must match the feature value at the projection level: both sides of the arrows in the pictorial representation must match. In (6), the source of the projection is [–ATR] but the projected value is [+ATR]. This type of representation is banned regardless of the source of the projection.

(6) Banned Projection: Mismatched Features

* /[+ATR]/	Underlying Form
[+ATR]	Projection/Phonological Level
↑	
[–ATR]	Pronunciation/Surface Level

Finally, the source of spreading must be faithful to its underlying form (i.e., it must be projected by its underlying feature). This restriction prevents representations of the type illustrated in (7) in which an underlyingly /–F/ segment spreads [+F] (e.g.,

/– – –/ ⇒ [+ + +]). The structure in (7) is banned because the underlying feature value of the initial vowel is [+ATR], but projects [–ATR] without spreading.

(7) Banned Representation

* / [–ATR]		[–ATR]/	Underlying Form
[+ATR]	→	[+ATR]	Projection/Phonological Level
↑			
[+ATR]		[+ATR]	Pronunciation/Surface Level

The structure in (8) represents proper spreading; the underlying feature of the initial vowel spreads at the projection level to the final vowel.[2] This example illustrates the possibility for a vowel to change its feature value without spreading. This occurs when the vowel is projected by its pronunciation or when the vowel pronounces a different feature value than its projection. The restriction on the source of spreading does not prevent 'abstract' analyses (Hyman, 2002); if an underlying value changes, it still may spread its underlying feature value.

(8) Turbid Spreading

/[+ATR]		[–ATR]/	Underlying Form
↓			
[+ATR]	→	[+ATR]	Projection/Phonological Level
[–ATR]		[+ATR]	Pronunciation/Surface Level

This completes the restrictions on GEN, which are summarized in (9) below.

(9) Restrictions on GEN

a. Feature values on both sides of projection arrow must match
b. All underlying feature values projected by one and only one element: underlying form, surface form, projection level
c. The initiator of spreading must be projected by its underlying feature value
d. Deletion occurs at the pronunciation level only
e. Epenthesis occurs at the pronunciation level or the projection level

These restrictions limit the number of candidates that can appear in a given optimization, and are essential for preventing an explosion of possible representations.

2.1 Representations of Spreading

I assume that spreading is an abstract process and occurs at the projection level. In rightwards spreading, the initial vowel spreads its projection (from the its underlying form) to the closest vowel to its right. This vowel may then spread its projection to the right, and onwards iteratively until spreading reaches the right edge of the word. In (10),

[2] Note that this does not preclude a case in which a vowel induces spreading after a change induced by a consonant. In this case, the source of the spreading would be the consonant.

left-to-right spreading occurs from the initial vowel, which is projected by its underlying [+ATR] feature value. This [+ATR] projection spreads to the second vowel. While the second vowel was underlyingly [–ATR], it is now projected as [+ATR]. This [+ATR] projection continues to spread to the right, and the final vowel projects differently from its underlyingly [–ATR] value. In this example, all vowels pronounce the same value of their projected values; all vowels pronounce [+ATR].

(10) Rightwards Vowel Harmony

/[+ATR]	[–ATR]	[–ATR]/	Underlying Form
↓			
[+ATR] →	[+ATR] →	[+ATR]	Projection/Phonological Level
[+ATR]	[+ATR]	[+ATR]	Pronunciation/Surface Level

Spreading need not change an underlying feature value; a projection may link to another projection regardless of whether it alters the feature value of the segment. For example, a [+ATR] segment may link to another [+ATR] projection, illustrated in (11). Spreading is considered 'vacuous' because the segments undergo spreading (by projecting their neighboring vowel) but do not change in feature value.

(11) Rightwards 'Vacuous' Spreading

/[+ATR]	[+ATR]	[+ATR]/	Underlying Form
↓			
[+ATR] →	[+ATR] →	[+ATR]	Projection/Phonological Level
[+ATR]	[+ATR]	[+ATR]	Phonological Level

In addition to leftwards/rightwards spreading, there are cases of vowel harmony that appear to be primarily bi-directional, particularly dominant-recessive vowel harmony. In these cases, a dominant feature value (e.g., [+ATR]) spreads to both the left and right. In other words, if a [+ATR] vowel is present in the input, all vowels will become [+ATR]. In these dominant-recessive languages, spreading is often bi-directional such that both stems and suffixes undergo harmony. In these cases, the source of harmony is the vowel with the underlyingly dominant feature value, and that feature value spreads to the left and right from the source.

(12) Bi-Directional Dominant Harmony

/[–ATR]	[+ATR]	[–ATR]/	Underlying Form
	↓		
[+ATR] ←	[+ATR] →	[+ATR]	Projection/Phonological Level
[+ATR]	[+ATR]	[+ATR]	Phonological Level

In (12), the underlyingly [+ATR] vowel spreads outwards bi-directionally. In this case, spreading is non-directional from the dominant feature-value.

2.2 Non-Participating Vowels

There are two ways for a segment to fail to participate in harmony. The first is at the projection level: the segment fails to undergo spreading at the projection level, producing an opaque vowel. The second is at the surface level: the vowel may undergo spreading at the projection level, but fail to pronounce the projected feature value, producing a transparent vowel, as in (13).

(13) Transparent Vowels

/[+ATR]	[–ATR]	[–ATR]/	Underlying Form
↓			
[+ATR] →	[+ATR] →	[+ATR]	Projection/Phonological Level
[+ATR]	[–ATR]	[–ATR]	Pronunciation/Surface Level

In (13), the projection spreads from left to right through the non-participating vowel. All three vowels project [+ATR], but the medial, non-participating vowel pronounces [–ATR], while the other vowels pronounce [+ATR].

When the features of the projection level match the features on the pronunciation level, the added level of representation merely serves to mark the direction and source of spreading. In the case of transparent vowels, the added level of representation provides a means for accounting for the long-distance, phonologically opaque case of 'skipping' vowels in harmony. Transparent vowels have created complications for accounts of spreading, typically requiring additional derivations or representations (Anderson, 1980, Bakovic and Wilson, 2000, Booij, 1984, Clements, 1977, O'Keefe, 2007, Pulleyblank, 1996, Pulleyblank, 2002, Smolensky, 2006). The added level of representation in Turbid Spreading allows for a seamless account of transparent vowels. The representation for opaque vowels in (14) is much simpler, as the features on the pronunciation and projection match. However, these vowels not undergo complete spreading.

(14) Opaque Vowels

/[+ATR]	[–ATR]	[+ATR]/	Underlying Form
↓	↓		
[+ATR]	[–ATR] →	[–ATR]	Projection/Phonological Level
[+ATR]	[–ATR]	[–ATR]	Pronunciation/Surface Level

In (14), the projection from the initial target vowel does not spread to the medial, non-undergoing vowel, but the projection of this non-undergoer is able to continue to spread rightwards, thereby behaving as an opaque vowel.

3. Constraints on Vowel Harmony

The constraints that induce harmony are: SPREAD, ID[F], RECIPROCITY and Featural Markedness. These are defined and discussed below.

3.1 SPREAD[αF]

The general harmony-inducing constraint SPREAD is defined in (15) and (16). These constraints require the projection representation for each vowel to originate from a neighboring vowel (visually represented by a horizontal arrow in the specified direction for spreading), with the source of spreading projected by its underlying form. For example, for the sequence /+ –/ to undergo spreading from left to right, the initial vowel /+/ must be projected by its underlying form, and the final vowel /–/ must have a [+] projection licensed by the initial vowel. The SPREAD-R constraint does not target the initial vowel because there is no vowel to spread to it. The converse is true for SPREAD-L; the final vowel cannot undergo leftward spreading, and is therefore not targeted by the constraint.

(15) SPREAD-R[αF]:
For all non-initial vowels, for each feature value [αF] on the phonological level, assign one violation if there is not a rightward-pointing projection representation originating at that feature value belonging to a rightward adjacent vowel.

(16) SPREAD-L[αF]:
For all non-final vowels, for each feature value [αF] on the phonological level, assign one violation if there is not a leftward-pointing projection representation originating at that feature value belonging to a leftward adjacent vowel.

It is impossible to satisfy both SPREAD-R and SPREAD-L simultaneously. For example, in order for the input /+ –/ to satisfy SPREAD-R, the initial vowel must project its underlying form, violating SPREAD-L. In order to satisfy SPREAD-L, the final vowel must project its underlying form, violating SPREAD-R.

I assume that the SPREAD constraints are evaluated directionally (Eisner, 2000)[3]. Directional evaluation breaks up the evaluation of the constraint in terms of each segment in the input. For example, SPREAD-R is evaluated from the leftmost vowel first, followed by each additional vowel, rightwards. This provides a formal mechanism for denoting the exact locus of each violation of spreading, similar to the effect of constraint violation found in the ALIGN family of constraints, in which (for example) failure to spread rightwards is worse at the beginning of a word because the fewer vowels that intervene between the disharmonic segments and the end of the word, the better the harmonic feature is aligned to the end of the word. Directional evaluation computes violations differently for each vowel in the input. The violation profile for an epenthetic vowel is the same as the closest vowel with an underlying form. Thus, epenthetic vowels do not affect the severity of disharmony for each individual failure to spread.

The tableau in (17) illustrates the use of directional evaluation for the input /+ – – –/. Each candidate incurs a single violation of SPREAD-R. Candidate (a.) fails to spread to the first vowel, candidate (b.) fails to spread to the second vowel, and candidate (c.) fails to spread to the final vowel. With directional evaluation, each locus of spreading

[3] Precursors to directional nature of evaluation can be found in the original formulation of Optimality Theory (Prince and Smolensky, 1993/2004) in which there may be weights on the location of a violation. For example, a violation at the beginning of the word could have a more severe violation than the middle of the word, which may be more severe than a violation at the end of the word.

is divided into a separate evaluation, with the ranking of the first locus for spreading ranked highest. Failure to spread from left-to-right is worse at the leftmost locus of violation. This produces iterative behavior of vowel harmony because left-to-right spreading must start at the leftmost edge, even if spreading elsewhere in the word would be advantageous. For ease of reference, a subscript is placed next to each violation of SPREAD to indicate its location. This location also serves as a tag for severity of the violation. In future tableaux, only subscripts will be used.

Representations in the tableaux are arranged such that the top line in GEN indicates the underlying form, the second line indicates the projection level, and the third line represents the pronunciation level.

(17) Directional Evaluation of SPREAD-R

/+ – – –/	SPREAD-R		
	V1-V2	V2-V3	V3-V4
(a.) / + – – – / ↓ ↓ + – → – → – [+ – – –]	$*_1$!		
(b.) / + – – – / ↓ ↓ + → + – → – [+ + – –]		$*_2$!	
(c.) ☞ / + – – – / ↓ ↓ + → + → + – [+ + + –]			$*_3$

Note that while the violations for SPREAD appear as if they were different constraints, it is not possible to interleave another constraint into the violations of loci of evaluations for SPREAD.

3.2 SPREAD[+F]

The SPREAD constraints in (15) and (16) can be modified to account for cases in which a dominant feature value spreads. In these cases, the presence of a particular feature value (e.g., [+ATR]) induces harmony that typically will involve spreading from suffixes to stems, as well as from stems to suffixes. Following previous analyses of dominant-recessive harmony systems (Archangeli and Pulleyblank, 2002, Bakovic, 1999, Bakovic, 2000, Mahanta, 2007, Noske, 2000, Orie, 2003), I assume that dominant-recessive harmony is induced by two constraints: a constraint requiring spreading of a particular feature value e.g., SPREAD-R[+ATR], and an identity constraint preventing underlyingly dominant segments from changing their feature value (e.g., in previous analyses

Max[ATR] (Archangeli and Pulleyblank, 2002, Orie, 2001, Orie, 2003), ID[+ATR], ID[ATR] & *[–ATR] (Bakovic, 1999, Bakovic, 2000)). In dominant-recessive harmony, the dominant feature value of the harmonic feature is included in the spreading constraint. This is essentially the same as the SPREAD constraints that do not specify a specific feature value. The difference here is that the spread constraint only applies if there is a [+F] feature value in the projection level, and all vowels in the domain of the constraint are required to be [+F]. These are defined in (18) and (19) below.

(18) SPREAD-R[+F]:
For all non-initial vowels, for each feature value [+F] on the phonological level, assign one violation if there is not a rightward-pointing projection representation originating at that feature value belonging to a rightward adjacent vowel.

(19) SPREAD-L[+F]:
For all non-final vowels, for each feature value [+F] on the phonological level, assign one violation if there is not a leftward-pointing projection representation originating at that feature value belonging to a leftward adjacent vowel.

If both SPREAD constraints are ranked above faithfulness, harmony will be bi-directional. If featural identity ranks in-between the two spread constraints, harmony will be dominant in a particular direction. While the majority of dominant-recessive harmony languages are bi-directional, there are cases of harmony where a specific feature value spreads in a single direction. For example, languages with a privative feature (e.g., [NASAL] or [ROUND]) spread the specific feature value in one particular direction.

Unlike the generic version of SPREAD, both SPREAD-L[+F] and SPREAD-R[+F] can be satisfied simultaneously. This is because violations for SPREAD[+F] are assigned from the source of spreading, the [+F] vowel (if present). For example, /– + –/ ⇒ [+++] satisfies both SPREAD-L[+F] and SPREAD-R[+F] because [+F] spreads to the left from the source, as well as right from the source.

3.3 Constraints on Non-Participating Vowels

As noted above, transparent vowels satisfy SPREAD, but at the expense of creating a mismatch between the projection level and the pronunciation level. This mismatch is penalized by the RECIPROCITY family of constraints, which enforces uniformity of feature values between the pronunciation level and the projection level. Segments whose projection features do not match their corresponding pronunciation features incur a violation of RECIPROCITY. This includes mismatches that occur when the pronunciation feature has deleted, but the projection still remains. Deleted vowels always incur a violation of RECIPROCITY because deleted segments always leave a trace behind at the projection level.

(20) RECIPROCITY (REC):
Assign one violation for every feature value [αF] that does not have a corresponding value [αF] at the phonetic level.

RECIPROCITY is evaluated as a comparison between the projection level and the pronunciation level, and is completely independent of the input. RECIPROCITY violations can occur when the underlying feature value changes (21) but also when there is no change to the underlying form (22).

(21) Violation of RECIPROCITY

[–ATR]	Underlying Form
↓	
[–ATR]	Projection/Phonological Level
[+ATR]	Pronunciation/Surface Level

(22) Violation of RECIPROCITY

[–ATR]	Underlying Form
[+ATR] ← [+ATR]	Projection/Phonological Level
[–ATR]	Pronunciation/Surface Level

However, if the underlying form changes its feature value, RECIPROCITY may be satisfied if the projection and pronunciation values match even if there is a change in the underlying form of the vowel. If a vowel is projected by its surface form, there is no way for that vowel to violate RECIPROCITY, as the feature value of the surface form dictates the feature value of the projection level.

(23) Projection from the Surface Form Necessarily Satisfies RECIPROCITY

[–ATR]	Underlying Form
[+ATR]	Projection/Phonological Level
↑	
[+ATR]	Pronunciation/Surface Level

While RECIPROCITY governs identity between the projection and the pronunciation levels, featural identity (ID[F]) governs the relationship between the underlying form and the projection level. This version of featural identity is stricter than standard ID[F]. ID[F] may be violated even if the feature value of the underlying form is identical to the feature value of the projection, but the vowel does not project the underlying form.

(24) ID[F]:
Assign one violation for every segment whose underlying form does not project the feature [F] onto that segment, whose projection comes from its pronunciation form or a neighboring segment.

ID[F] is violated by any segment that is projected by its surface representation or the projection of one of its neighbors. ID[F] may be satisfied even if the surface form is not identical to the underlying form (but this does violate RECIPROCITY, described above). Because the underlying form projects this vowel, there is no violation of ID[ATR].

I also assume an asymmetric identity constraint, which applies to values of a specific feature value (typically the dominant feature value in a dominant-recessive harmony language such as [+ATR]). This asymmetric faithfulness (defined in (25) below) is violated with the [+F] feature is present underlyingly but not projected in the intermediate representation.

(25) ID[+F]:
Assign one violation for every segment with [+F] as the underlying form does not project the feature [+F] onto that segment.

The final type of constraint that interacts with vowel harmony is featural markedness. For simplicity, as a stand-in for all feature co-occurrence constraints, I will assume a single feature co-occurrence constraint in the following analysis of vowel harmony: *[+ATR, –HIGH][4].

(26) *[+ATR, –HIGH]:
Assign one violation to each vowel that is [+ATR] and [–HIGH] on the pronunciation level.

This markedness constraint determines whether a harmony language has non-participating vowels. A high-ranked featural markedness constraint yields non-participating vowels.

4. Vowel Harmony in Turbid Spreading

Together, the above constraints form the grammar necessary to induce vowel harmony. I assume that vowel harmony is induced by a markedness-faithfulness constraint interaction such that harmony applies when SPREAD[F] outranks ID[F]. When *[+ATR, –HIGH] outranks both SPREAD[ATR] and ID[ATR], non-high vowels will not participate in harmony. The ranking of RECIPROCITY with respect to SPREAD decides whether the non-participating vowel is transparent or opaque. If RECIPROCITY is ranked above SPREAD, the non-participating vowel is opaque. If RECIPROCITY is ranked below SPREAD, the non-participating vowel is transparent. For this paper, I ignore the effects of MAX and DEP. I also assume a high-ranked faithfulness constraint to the feature [HIGH] (e.g.,

[4] While I have been concerned with the representation of a single feature value ([ATR]), I assume that all vowels have three-leveled representations for each feature, ATR, high, etc. For ease of exposition, these other feature representations will not be included unless necessary.

ID[HIGH]) such that no vowels will change their height feature to satisfy the SPREAD constraint.

In (27) (a.), the initial underlyingly [+ATR] vowel spreads to the final vowel, which violates ID[F]. Candidate (b.) also undergoes spreading, but pronounces [–ATR], incurring a fatal violation of RECIPROCITY. Candidate (c.) does not undergo spreading, and therefore violates the SPREAD constraint. Notice that projection from the pronunciation level (candidate (d.) is harmonically bounded by candidate (c.) because it violates faithfulness in addition to the SPREAD constraint.

(27) Vowel Harmony and Turbid Spreading (Opacity Ranking)

/i ɪ/	*[+ATR, –HIGH]	REC	SPREAD[ATR]-R	ID[ATR]
(a.) ☞ / + – / ↓ + → + [+ +]				*
(b.) / + – / ↓ + → + [+ –]		*!		*
(c.) / + – / ↓ ↓ + – [+ –]			*!	
(d.) / + – / ↓ + – ↑ [+ –]			*!	*

The relative ranking of RECIPROCITY and SPREAD determines whether a non-participating vowel is transparent or opaque. In (28) below the non-participating vowel /ɛ/ is flanked by [+ATR] to the left and a [–ATR] vowel to its right. The transparent candidate (b.) spreads [+ATR] all the way through the non-participating vowel.

(28) Transparency and Turbid Spreading

/i ɛ ɪ/	*[+ATR, –HIGH]	SPREAD [ATR]-R	REC	ID [ATR]
(a.) / + – – / ↓ + → + → + [+ + +]	*!			**
(b.) ☞ / + – – / ↓ + → + → + [+ – +]			*	**
(c.) / + – – / ↓ ↓ ↓ + – – [+ – –]		$*_1!*_2$		**
(d.) / + – – / ↓ + → + → + [+ – –]			**!	

The forms in candidates (a.), (b.) and (d.) successfully spread the [+ATR] feature to all vowels. Candidate (a.) fails because it produces the ungrammatical non-high [+ATR] vowel. Candidate (d.) fails because it produces two RECIPROCITY violations. Candidate (c.) fails because RECIPROCITY outranks SPREAD. When this ranking is reversed, the opaque candidate will surface, as in (29) below.

(29) Opacity and Turbid Spreading

/i ɛ ɪ/	*[+ATR, –HIGH]	REC	SPREAD [ATR]-R	ID [ATR]
(a.) / + – – / ↓ + → + → + [+ + +]	*!			**
(b.) / + – – / ↓ + → + → + [+ – +]		*!		**
(c.) / + – – / ↓ ↓ ↓ + – – [+ – –]			$*_1*_2$!	
(d.) / + – – / ↓ + → + → + [+ – –]		**!		**
(e.) ☞ / + – – / ↓ ↓ + – → – [+ – –]			$*_1$	

When both SPREAD-R and SPREAD-L are allowed to interact, more options are available for possible winners. I assume that SPREAD-R and SPREAD-L are always ranked with respect to each other such that there is always a default direction for spreading in any language. When both SPREAD constraints outrank ID[F], spreading will apply in the opposite direction, when spreading in the default direction is not possible. This is illustrated in (30) below.

(30) Interaction of SPREAD-R and SPREAD-L

/eɪ/	*[+ATR, –HIGH]	REC	SPREAD [ATR]-R	SPREAD [ATR]-L	ID [ATR]
(a.) / + – / ↓ ↓ + – [+ –]	*!		$*_1$	$*_2$	
(b.) / + – / ↓ ↓ + + [– –]		*!	$*_1$	$*_2$	
(c.) / + – / ↓ + → + [– +]		*!		$*_2$	*
(d.) ☞ / + – / ↓ – ← – [– –]			$*_1$		*

In (30) the only candidate that can satisfy SPREAD also violates high-ranked RECIPROCITY. Because SPREAD-L outranks ID[F], spreading can apply in the opposite direction. In this type of language, rightward spreading is the default, but when this is not possible, leftward spreading may apply.

In the tableaux presented thus far, all candidates that include a vowel that is projected by the surface form are harmonically bounded. This is because a feature projected by its surface form violates faithfulness and cannot participate in spreading. However, it is possible for a candidate projected by its surface form to win an evaluation. This occurs if (i) the underlying feature value of the vowel is marked and cannot surface faithfully, (ii) the vowel cannot get its unmarked feature value from spreading, and (iii) RECIPROCITY is ranked above ID[F]. For example, the input /e/ in tableau (31) below cannot surface faithfully as [e] (candidate (a.)) without violating the high-ranked featural markedness constraint or projected by the underlying form but pronounced as [–ATR] [ɛ] (candidate (b.)) without violating RECIPROCITY. When RECIPROCITY outranks ID[F], projection from the surface form is optimal (candidate (c.)).

(31) Pronunciation-Level Projection

/e/	*[+ATR, –HIGH]	REC	ID[ATR]
(a.) / + / ↓ + [+]	*!		
(b.) / + / ↓ + [–]		*!	
(c.) ☞ / + / – ↑ [–]			*

Because so many conditions are required for a feature to be projected by its surface form, we expect that this representation should only occur in select forms. This expectation is borne out in the factorial typology: few languages allow projection by the pronunciation form.

Dominant-recessive harmony is induced when the spread constraints select a specific feature value (e.g., [+ATR]). If SPREAD-R[+ATR], SPREAD-L[+ATR] and ID[+ATR] are all ranked above ID[ATR], then harmony will be bi-directional.

(32) Bi-Directional Harmony in Dominant-Recessive Harmony

/– + –/	SPREAD-L [+ATR]	ID [+ATR]	SPREAD-R [+ATR]	ID [ATR]
(a.) ☞ / – + – / ↓ + ← + → + [+ + +]				**
(b.) / – + – / ↓ – ← – ← – [– – –]		*!		**
(c.) / – + – / ↓ – → – → – [– – –]		*!		*
(d.) / – + – / ↓ ↓ + ← + – [+ + –]			$*_2$!	*
(e.) / – + – / ↓ – + → + [– + +]	*!			*

Note that candidate (a.) does not violate either SPREAD constraint because the [+ATR] feature spreads both to the left and the right. Thus, unlike the generic SPREAD constraint, SPREAD-[+F] can be satisfied for both directions, given that the source of spreading can be in the middle of the word, rather than at a particular edge.

While the initial vowel does not spread, no violations of SPREAD are incurred because this initial vowel is [–ATR] and not subject to the SPREAD-[+ATR] constraint. This is because the source for spreading is assumed to be the vowel that underlyingly possesses the [+ATR] feature.

The relatively small set of constraints: SPREAD, featural markedness, RECIPROCITY and ID can be used to derive a rich set of interactions in vowel harmony, including dominant feature values and directionality effects.

5. Conclusions

This paper has provided a representational approach to vowel harmony using Turbidity Theory, Turbid Spreading. In Turbid Spreading, all features have three levels of representation: an underlying form, a projection (abstract) form and a phonetic (surface) form. These three levels interact such that spreading is initiated by an underlying form and applies through the projection level. Because the pronunciation representation need not share the same feature value as the projection level, vowels may undergo spreading abstractly, but pronounce a different feature, providing an account of transparent vowels. Because this mismatch of pronunciation and projection comes at a cost (violating a RECIPROCITY constraint), some rankings will produce transparent non-participating vowels, while other rankings will produce opaque non-participating vowels.

The Turbid Spreading representational approach to vowel harmony is part of a larger research program for understanding the nature of the typology of vowel harmony processes. Several issues will be addressed in future research including parasitic harmony and interactions with consonants. For example, the representations provided in the present proposal assume only vowel-to-vowel interactions on a single feature value. However, there are many harmony processes that interact with consonants (Turkish is a prime example (Clements and Sezer, 1982), as well as the consonant interactions seen in Sesotho above (Rose and Demuth, 2006)), and many harmony processes that are dependent on multiple features, such as Yawelmani (Cole and Trigo, 1988) and Turkish, where only high vowels participate in round harmony (Charette and Goksel, 1998, Clements and Sezer, 1982, Kirchner, 1993, Polgardi, 1999, Underhill, 1976).

While there are many avenues for future research, as well as many unknowns in the validity of the prediction, the work presented here provides a method for u for directional effects in harmony such that harmony can apply non-directionally or directionally, as well as with a dominant feature value.

References:

Anderson, Stephen. 1980. Problems and presective in the description of vowel harmony. In *Issues in vowel harmony*, ed. Robert M. Vago, 1-48. Amsterdam: John Benjamins.

Archangeli, Diana, and Pulleyblank, Douglas. 2002. Kinande vowel harmony: domains, grounded conditions and one-sided alignment. *Phonology* 19:139-188.

Bakovic, Eric. 1999. Assimilation to the unmarked. *Penn Working Papers in Linguistics* 6.

Bakovic, Eric. 2000. Harmony, dominance and control, Unpublished doctoral dissertation, Rutgers University.

Bakovic, Eric, and Wilson, Colin. 2000. Transparency, strict locality and targeted constraints. *WCCFL* 19:43-56.

Booij, Geert. 1984. Neutral vowels and the autosegmental analysis of Hungarian vowel harmony. *Linguistics* 22:629-641.

Charette, Monik, and Goksel, Asli. 1998. Licensing constraints and vowel harmony in Turkic languages. *Structure and interpretation: Studies in phonology* 4.

Clements, George N. 1977. Neutral vowels in Hungarian vowel harmony: an autosegmental interpretation. *NELS* 7:49-64.

Clements, George N., and Sezer, Engin. 1982. Vowel and consonant disharmony in Turkish. In *The structure of Phonological Representations*, ed. van der Hulst and Smith, 213-255. Dordrecht: Foris.

Cole, Jennifer, and Trigo, Loren. 1988. Parasitic harmony. In *Features, segmental structures and harmony processes*, eds. Norval Smith and Harry van der Hulst, 19-38. Dordrecht: Foris.

Eisner, Jason. 2000. Directional constraint evaluation in Optimality Theory. Paper presented at *35th Annual Meeting for the Association for Computational Linguistics and the 8th Conference of the European Association for Computational Linguistics*, Madrid, July.

Finley, Sara. 2008. Formal and Cognitive Constraints on Vowel Harmony, Cognitive Science, Johns Hopkins University.

Finley, Sara. in preparation. Turbid Spreading. Ms. University of Rochester.

Goldrick, Matthew. 1999. Turbid output representations and the unity of opacity. Ms. Talk presented at NELS 30.

Goldrick, Matthew. 2001. Turbid output representations and the unity of opacity. *NELS* 30:231-245.

Hyman, Larry H. 2002. "Abstract" vowel harmony in Kalong: A system driven account. *Theories Linguistiques et Langues Sub-Sahariennes* 8:1-22.

Kirchner, Robert. 1993. Turkish vowel harmony and disharmony: An Optimality Theoretic account. Paper presented at *Rutgers Optimality Workshop I (ROW-I)*.

Mahanta, Shakuntala. 2007. Directionality and locality in vowel harmony: With special reference to vowel harmony in Assamese, Utrecht Institute of Linguistics.

Noske, Manuela. 2000. [ATR] harmony in Turkana: A case of faith suffix >> faith root. *NLLT* 18:771-812.

O'Keefe, Michael. 2007. Transparency in span theory. In *University of Massachusetts Occasional Papers in Linguistics 33: Papers in Optimality 3*, eds. Leah Bateman, Adam Werle, Michael O'Keefe and Ehren Reilly, 239-258. Amherst, MA: GLSA.

Orie, Olanike Ola. 2001. An alignment-based account of vowel harmony in Ife Yoruba. *Journal of African Language and Linguistics* 22:117-143.

Orie, Olanike Ola. 2003. Two harmony theories and vowel patterns in Ebira and Yoruba. *The Linguistic Review* 20:1-35.

Polgardi, Krisztina. 1999. Vowel harmony and disharmony in Turkish. *The Linguistic Review* 16:187-204.

Prince, Alan, and Smolensky, Paul. 1993/2004. *Optimality theory: Constraint interaction in generative grammar*. Cambridge: Blackwell.

Pulleyblank, Douglas. 1996. Neutral vowels in Optimality Theory: A comparison of Yoruba and Wolof. *Canadian Journal of Linguistics* 41:295-347.

Pulleyblank, Douglas. 2002. Harmony drivers: No disagreement allowed. *BLS*:249-267.

Rose, Yvan, and Demuth, Katherine. 2006. Vowel epenthesis in loanword adaptation: Representational and phonetic considerations. *Lingua* 116:1112-1139.

Smolensky, Paul. 2006. Optimality in phonology II: Harmonic completeness, local constraint conjunction, and feature-domain markedness. In *The Harmonic Mind: From Neural Computation to Optimality-Theoretic Grammar, Volume II*, eds. Paul Smolensky, Geraldine Legendre and (eds.), 27-160: MIT Press.

Underhill, Robert. 1976. *Turkish grammar*. Cambridge: MIT Press.

Department of Brain and Cognitive Sciences
University of Rochester
Meliora Hall
Rochester, NY 14627

sfinley@bcs.rochester.edu