

Reconstructing Mirativity: Evidentiality, Information, and Deixis*

Abstract

The goal of this paper is to investigate the linguistic, psychological and cognitive properties of utterances that express the surprise of the speaker, with a focus on how grammatical evidentials are used for this purpose. This is often labeled in the linguistics literature as *mirativity*. While there has been a flurry of recent interest in mirativity, we still lack an understanding of how and why evidentials are used this way, and an explanation of this effect. In this paper I take steps to filling this gap by claiming that the mirativity associated with grammatical evidentials is one of the many linguistic reflexes of the more general cognitive process of surprise. I approach this by analyzing mirativity, and the language of surprise more generally, in a schema-theoretic framework enriched with the notion of *new environmental information*. I elaborate on the field methodological issues involved with testing the mirative use of an evidential and why they are used this way by connecting mirative evidentials to the broader phenomenon of *deixis*.

Keywords: Grammatical evidentiality, Mirativity, Deixis, Information Theory, Schema Theory, Bayesian Surprise

1 Introduction

Broadly speaking, new information that is not easily assimilated into a speaker's current situational awareness is often coded differently linguistically than information that can be adapted into this awareness. This coding – often referred to as *mirativity* – is the linguistic reflex of what we commonly interpret as surprise. Languages employ a variety of grammatical strategies for registering the surprise of the speaker. In languages such as English, there are lexical items that entail the surprise of the speaker: 'I'm *surprised* you arrived on time'. However, even a cursory survey into this area reveals that the expression of surprise is subject to considerable cross-linguistic variation. As an example of this, imagine that Fern and Gwen are preparing a surprise birthday party for Sheila. Gwen doesn't expect Sheila home for another

*Special thanks to my Gitksan consultants Barbara Sennott, Leiwa Weget, Louise Wilson, and my Turkish consultants Hande Ergun and Hakan Ergun. For the useful advice and support, thanks to Tom Bever, Massimo Piatelli-Palmerini, Dane Bell, and the members of the Cognition, Language And Neuroscience Lab at the University of Arizona. Research on the Gitksan language was made possible by a grant from *The Endangered Languages Documentation Program* (SOAS) awarded to the author. Data that is not cited is from fieldwork, and all errors and any possible misinterpretations of secondary data are solely my responsibility.

hour. After an hour Gwen looks outside and sees Sheila's car in the driveway. In this context a speaker of Gitksan (an endangered indigenous language of the Tsimshianic family spoken in northwestern Canada) will likely use the sentence in (1):

- (1) *n'akw*=hl wixw=s Sheila
EVID=CND arrive=PND Sheila
'Looks like Sheila is here.'

Gwen does not actually *see* or know that Sheila has, in fact, arrived; rather, Gwen infers that she has arrived because Gwen has indirect sensory evidence for her arrival: Sheila's truck is in the driveway. Gwen linguistically expresses her inference using the grammatical evidential *n'akw* (glossed as EVID) – the morpheme in this language that lexically encodes an inference from sensory evidence (Peterson 2010a). Now, adjust this scenario slightly: Alvin's family is preparing a surprise birthday party for him. Gwen doesn't expect him home for another hour. However, in the middle of their preparations Alvin walks through the door. In this context Gwen exclaims:

- (2) *n'akw*=hl wixw=s Alvin
EVID=CND arrive=PND Alvin
'[Looks like] Alvin is here!'

The sentence in (2) is grammatically identical to the sentence in (1). However, the indirect evidential meaning of *n'akw* seem to somehow be 'short-circuited' in this context: Gwen *actually sees* that it is the case that Alvin is here. This 'misuse' of *n'akw* does not lead to infelicity, instead, in *direct* evidence contexts *n'akw* has a mirative translation in addition to its evidential one: Gwen is mentally unprepared for and surprised at Alvin's early arrival.

Linguists have long since recognized the connection between evidentiality and surprise, and recently there has been a surge of interest in mirativity.¹ As such, we are beginning to get a clearer picture of the semantic features of mirativity as a linguistic category and how it is manifested cross-linguistically. However, mirativity has yet to face a dedicated, theoretical investigation. This leads to a number of basic questions that motivate this paper: what are the contextual and cognitive factors that drive the mirative meaning in (2). There is a robust connection between grammatical evidentiality and surprise – one of the main manifestations of mirativity in the world's languages. Why are grammatical evidentials used for this purpose? Finally, although we are now equipped with many rich descriptions of mirativity and detailed generalizations from a functional/typological perspective, there is still no clear field methodology that is dedicated to testing the mirative use of a grammatical evidential.

The aim of this paper is to fill this gap by explaining how mirativity is a linguistic reflex of a general cognitive process involving *surprise*, one the core human emotions. However, because the empirical scope of this study is how and why grammatical evidentials are used to express surprise (mirativity), it is necessary to examine in detail evidential meaning on its own. This is done using primary language data from two unrelated languages, Gitksan and Turkish, and set within the context of the current views on evidentials and mirativity. I show that only certain kinds of evidentials can be used to express surprise; specifically, in systems with two or more grammatical evidentials, the evidential that is most specialized for sensory

¹For example, the reader is referred to a special issue of *Linguistic Typology* (2012, Vol. 16, No. 3).

evidence is usually the one that expresses mirativity. This predicts that reportative evidentials cannot be used this way, as they typically encode information from linguistic sources, not sensory ones. This empirical generalization motivates the need to distinguish between what I call *environmental information* and *linguistic information*. In this paper I also discuss the methodological considerations involved with testing surprise and mirativity, something that has not been previously done in the literature. Finally, I offer an explanation of why grammatical evidentials are suited to expressing surprise by connecting this to deixis: mirative evidentials are used to ‘point’ or draw the attention of conversational participants to surprising information.

The paper is structured as follows: in Section 2 I review and evaluate two of the predominant theoretical approaches to surprise; the first is *schema theory*, which provides a cognitive framework for how surprising information affects our background knowledge and current situational awareness. The second is an exposition of a Bayesian definition of surprise. This offers a quantifiable view of what surprise is in terms of how information effects a change in the probability distributions representing our hypotheses about the world and current situations. Neither of these theories have been put to use in explaining the language of surprise (or mirativity), but I show how schema theory is ideally suited to the task. I discuss the potential use of a Bayesian approach, especially as it has been used in experimental settings examining the physiological responses to surprising stimuli. In section 3 I introduce grammatical evidentiality, with a focus on the evidential system on Gitksan. I also focus on how evidentials that encode sensory evidence are specialized for expressing mirativity. In section 4 I introduce the concept of *new environmental information*, a concept that frames states, events and actions in terms of information that is new to the speaker – a key contextual driver of surprise and, by extension, mirativity. New environmental information is then embedded within a schema theoretic analysis of surprise. In this section I also outline one of the field experiments used to test the mirative use of *nakw* in (2). In section 5 I present an analysis of mirative evidentials as deixis. Section 6 concludes, and outlines some future avenues of research.

2 Quantifying surprise

Aside from the landmark studies of Slobin and Aksu (1982); Aksu-Koç and Slobin (1986) on Turkish, little attention has been given to the link between the linguistic encoding of surprise and the psychology of surprise. What is known about surprise originates mainly from the literature in psychology (for example Meyer and Niepel 1994; Meyer et al. 1997; Niepel et al. 1994; Reisenzein 2000). More recently surprise has also been the focus of numerous investigations in the cognitive and computer sciences (Schützwohl 1998; Itti and Baldi 2005, 2006; Lorini and Castelfranchi 2007), language acquisition (Kagan 2002), and artificial intelligence (Macedo et al. 2009).

As a starting point, surprise can be characterized in three complementary ways. First, surprise is one of the core human emotions, along with happiness, sadness, anger, disgust, contempt, and fear (Ekman 1980, 1984; Izard and Malatesta 1987).² These core emotions are biologically determined responses whose expression and recognition – typically through a combination of verbal and non-verbal cues seen in association, or what is called a *syndrome* –

²See Schützwohl (1998) Reisenzein (2000) and for an comprehensive overview on the history and evaluation of research on surprise.

is fundamentally the same for all individuals, regardless of ethnic or cultural differences. Secondly, surprise can be viewed from an evolutionary and cognitive perspective: the concept of surprise is central to sensory processing, adaptation and learning, attention, and decision making. Thirdly, surprise can be viewed as a probabilistic or information-theoretic phenomenon: any sudden variance in the expected values within a dynamic system, such as the flow of information, can be quantified as surprise within that system.

We can gain a novel insight into the linguistic realizations of surprise (and more specifically, mirative evidentials) through engaging in this research. I do this in two parts: first, I examine the predominant cognitive analysis of surprise, *schema theory*, and then show how the language of surprise fits into this model. Secondly, I provide an exposition of the predominant probabilistic approach to surprise, and how it has been used in an experimental setting to measure the physiological correlates of surprise. I evaluate these two approaches and discuss how this might be applied to testing the language of surprise in an experimental setting.

2.1 Surprise as a cognitive process

The properties of the emotion of surprise have been effectively charted through the use of *schema theory* (i.e. Mandler 1984; Rumelhart 1984). A central feature of this theory are *schemata*, typically defined as non-linguistic formal objects that are organized knowledge structures representing concepts such as situations, objects, events, and actions at various levels of abstractness. Schemas can be thought of as categorical rules or scripts that we use to interpret the world and manage the relationships among them; they are also implicated in memory retrieval biases and the planning of responses to new information (DiMaggio 1997; Kleider et al. 2008). Information that does not fit into these schema may not be comprehended, or may not be comprehended correctly. Sensory information (i.e. visual, auditory, tactile, and olfactory) is processed according to how it fits into these schema, which are used not only to interpret our environment through the inferences performed on this information, but also to predict outcomes of events or situations occurring in our immediate environment. However, the world is often a highly dynamic place, and in order to maintain a proper representation of reality there is a continuous feedback loop, where one continuously checks (mostly unconsciously) whether their schema matches the inputs coming from the surrounding environment (Scherer 1984). Schema theory is widely applied in the areas of social and cognitive psychology and is ideally suited to the task of explaining why new information (encoded by grammatical evidentials) is surprising. Thus, we make take a step forward in explaining mirativity.

Before we can examine this process in more detail, it is necessary to sharpen the concept of ‘sensory information’, as the use of grammatical evidentials such as *n̄akw* in (1) and (2) represent the encodings of the different kinds of sensory information that are attained through our environment. In section 4 I explain in detail the empirical and linguistic motivations for what I call *new environment information* (NEI). In a nutshell, NEI has two features: first, the information that is gained through sensory inputs must be environmental in nature. In other words, NEI excludes the kinds of information encoded through the linguistic signal, meaning language itself. Secondly, NEI must be ‘new’. What this means is that NEI must contribute to updating or revising the belief state of the agent. Here is a simple example that shows this: John’s believes his neighbour does not own a dog (John’s current belief state). However,

one afternoon John hears a dog barking in the neighbour’s apartment. This is *environmental information* (the sound of a barking dog in a particular setting), and this information is *new*: the sound of the barking dog was not previously perceived by John in this situation. As such, John now believes his neighbour recently acquired a dog (John’s updated belief state).

New events can involve a deviation to some degree from *activated cognitive schema*, which are schemata that are immediately relevant and currently activated in the mind of the speaker (akin to situational awareness). This deviation creates a *schema-discrepancy* (Meyer et al. 1997; Schützwohl 1998; Reisenzein 2000). Using these building blocks, a cognitive model in Fig. 1 outlines a series of six mental activities that occur during surprise.

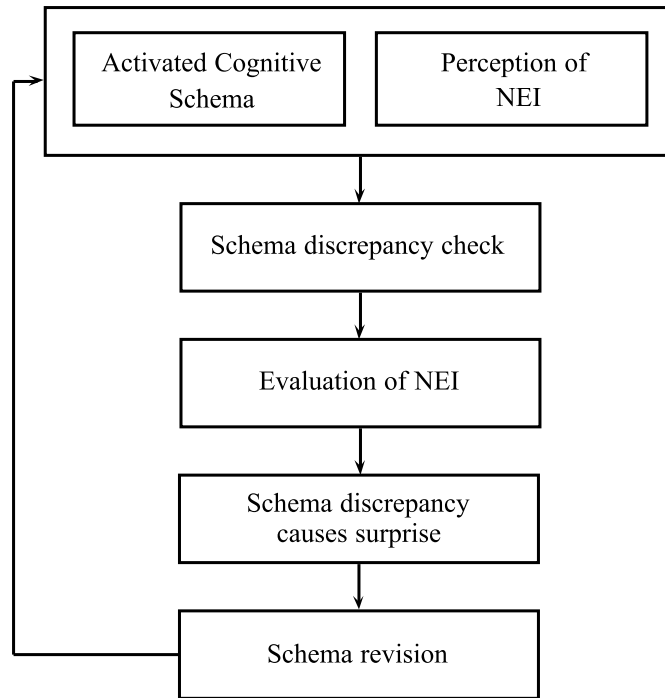


Figure 1: A Cognitive Model of the Mental Processes Elicited by Surprising Events (adapted from Reisenzein 2000: 265)

The initial state can be thought of as the state of the speaker at any given moment; it represents the current situational awareness (active cognitive schema), along with the capabilities of perception (sight, sound, tactile etc.) that constantly monitor NEI. Surprising NEI is evaluated within this process as exceeding some threshold value of schema-discrepancy. In other words, upon perceiving some NEI, one or more of the person’s current expectations is violated to a significant degree. The consequence that follows is that the person experiences surprise. Motivated by the feeling of surprise, the person analyzes the surprising event, consciously verifying its unexpectedness, assessing the cause of the event, and deciding whether some response or change in behaviour or future course of action is necessary. If necessary, a verbal response is initiated.

The schema discrepancy check is responsible for monitoring and evaluating the active cognitive schema: when NEI is perceived by the individual it is evaluated as either being assimilable or unexpected. If it passes under a certain threshold, then it feeds back and up-

dates the active cognitive schema. However, if this NEI counters an individual's previous expectations, or is an event that they are not mentally prepared for, then it triggers a schema discrepancy. It is at this step that surprise is an elicited response to a discrepancy. The NEI is then eventually assimilated and fed back into the active cognitive schema. Consider an analysis in (3) along these lines using the scenario from above:

- (3) 1. **Schema:** *our neighbour does not own a dog; animals are not permitted as pets in our apartment complex; our neighbour's child is allergic to animal hair.*
2. **Activated cognitive schema:** *John is sitting in his apartment one evening, relaxing and enjoying the peaceful quietness.*
3. Unexpected events may involve a deviation to some degree from an activated cognitive schema: **schema discrepancy check** and **evaluation of NEI** (the auditory perception of the barking of a dog)
4. **Schema discrepancy causes surprise:** *The sound of a dog barking in the neighbour's apartment*
5. **Revision:** *"Oh, our neighbour must now have a dog."*

Note that it's not necessary that it John's neighbour actually has a dog; there could be other possibilities, such as the sound of barking coming from the TV set, or perhaps someone with a dog is visiting John's neighbour. All that is required is that John *believes* (until further revision) that their neighbour is now in possession of a dog. Given the facts that comprise John's background beliefs (the *schema*), this may be surprising to John.

This also where we get our first glimpse at the connection between evidentiality in John's revised belief. Imagine the following conversation between John and his wife Mary:

- (4) M: What was that commotion next door?
- J: Our neighbour **must** now have a dog.

John could choose to assert *our neighbor has a dog*, since the NEI consists of evidence that contributes to his revised belief that their neighbour now owns a dog. However, it may be the case that John believes only that their neighbor *must* have a dog, given the same NEI. In John's response in (4) John does not know for sure that their neighbour owns a dog, but infers this from the NEI. What is relevant here is that this is the evidential use of the epistemic modal auxiliary *must*, which grammatically registers this inference (cf. von Stechow and Gillies 2007; Matthewson et al. 2007; Peterson 2010a). This hints at the connection between grammatical evidentiality and NEI that is elaborated on below.

Schema theory provides us with a way of calculating the cognitive steps that elicit the emotion of surprise. It is also where we can explicitly account for the asymmetric relation observed above between new information and surprise: old information is a part of a person's active cognitive schema and is used to appraise NEI, which may or may not lead to schema discrepancy. When it does, surprise is elicited; when it doesn't, it simply updates an individual's active cognitive schema. However, this is only part of the picture: the framework set out in Figure 1 treats surprise as a cognitive process; it isn't explicit as to the *presentation* of surprise. Thus, it is important to probe further these properties in order to shed light on why the elicitation of surprise can be a prominent enough feature to be linguistically encoded.

2.2 The Linguistic Correlates of Schema Discrepancy

The emotion of surprise is often viewed as a *syndrome*, which is commonly defined as the association of recognizable physiological and behavioural features that often occur together to form a pattern. The patterns that comprise the surprised reaction manifest themselves on three distinct levels simultaneously, given in (5) (Reisenzein et al. 1996; Meyer et al. 1997; Schützwohl 1998; Reisenzein 2000):

(5) The syndrome of surprise

1. THE SUBJECTIVE LEVEL: the subjective feeling of surprise; verbal exclamation of surprise.
2. THE PHYSIOLOGICAL LEVEL: different cortical response wave patterns; changes in heart and respiration rates; increased neural activation; etc.
3. THE BEHAVIOURAL LEVEL: subsequent curiosity/exploratory behaviour; specific facial expression; interruption of ongoing activities; etc.

It is at the subjective level that we find a channel for the linguistic expression of surprise. Unlike the physiological, and perhaps to a lesser extent, the behavioural levels, the subjective level mostly involves volitional responses to a surprising event. This may come in the form of simple spontaneous vocalizations (i.e. ‘wow!’, ‘oh!’, etc), or more complex constructions at the utterance level. This volitionality at the subjective level is reflected in the variety of linguistic strategies a speaker can use to express surprise.

(6) Context: *You and Gwen are preparing a surprise birthday party for Alvin. You don't expect home for another hour. However, in the middle of your preparations, Alvin walks through the door*

1. **Schema:** Surprise birthday parties are only surprising if the birthday boy or girl doesn't know you're planning it, and you can surprise them when they show up.
2. **Activated cognitive schema:** The party is for Alvin, the birthday boy; the party preparations are underway while Alvin is *not* there
3. **Schema discrepancy check:** Alvin is there

The schema discrepancy check in (6) results in a schema discrepancy and thus elicits surprise in this context. As such, the sentences in (7) are a sample of the variety of linguistic choices a speaker has to express surprise when Alvin unexpectedly shows up at a party:

- | | |
|---|---|
| (7) <i>Alvin made it!</i> | (Surprised intonation marked by '!') |
| <i>I don't believe Alvin made it!</i> | (Referencing the schema-discrepancy) |
| <i>Looks like Alvin made it!</i> | (Evidential verb, with/without '!') intonation) |
| <i>Wow, Alvin's here!</i> | (Surprise vocalization, plus '!') intonation) |
| <i>I'm amazed Alvin made it!</i> | (Verb of surprise) |
| <i>That can't be who I think it is!</i> | (Root modality, negation, plus '!') intonation) |
| <i>What a surprise (Alvin's here)!</i> | (Wh-exclamative, verb of surprise) |

In addition to giving us a glimpse of the intra-linguistic diversity in marking surprise, the utterances in (7) also show us how a single schema discrepancy (the perception of Alvin’s arrival) can have a multiplicity of linguistic realizations, for example, in the way information structure marking intonation (‘!’) can co-occur with an evidential verb. All of these expressions of surprise are a result of the same schema discrepancy: the speaker was unable to assimilate the NEI they were faced with in a specific context. This is a crucial point that is often equivocated on in the mirativity literature (Hengeveld and Olbertz 2012 make a similar point): in the English sentence ‘Alvin’s here!’, the speaker is *not* surprised at the proposition denoted by the sentence (Alvin’s here), rather, the speaker is surprised by the actual presence of Alvin – the NEI. The speaker then makes a subjective, linguistic choice in the special marking a sentence that registers this surprise – among the many other choices in (7). In the case of ‘Alvin’s here!’, this marking is a specialized intonational pattern.

2.3 A Bayesian definition of surprise

There is a relevant line of research on quantifying surprise in an experimental setting involving visual gaze analysis and explained using Bayesian probability. It is worth briefly exploring this approach, as a probabilistic view of knowledge and the evidence that changes that knowledge is relevant to the study of the language of surprise.

In determining what they call a unit of surprise (a ‘wow’ in their terminology) Itti and Baldi (2006) develop a Bayesian definition of surprise based on the quantifiable difference between an observer’s posterior and prior distributions of beliefs over a class of models or hypotheses about the world, which are updated as new sensory input is acquired. To begin with, there is a direct point of contact between NEI and what they describe as ‘new sensory input’. Their Bayesian model is schematized in Figure 2, which has been adapted to include NEI and the variable (ι) over instances of NEI:




Figure 2: Bayesian schema of prior and posterior belief distributions given NEI.

Using a standard Bayesian probability theorem, Itti and Baldi (2006) assume that an observer’s background knowledge is a probability distribution: an observer’s *prior distribution* of beliefs $\{P(M)\}_{M \in \mathbb{M}}$ relative to a model M in a set of possible models \mathbb{M} . This indicates one’s preconceived hypotheses about the world or current situation. For example, we can create a sample model space \mathbb{M} in (8) based on the scenario involving John’s neighbour from the section above:

- (8) $\mathbb{M} = \{\text{our neighbour does not own a dog; animals are not permitted as pets in our apartment complex; our neighbour’s child is allergic to animal hair}\}$

The probability of $P(M)$ is affected by the observation of NEI in a model space \mathbb{M} , defined as the variable (ι). When an observer encounters and processes NEI her prior belief distribution is transformed from $P(M)$ to the *posterior distribution* of beliefs $P(M|\iota)$: the probability of M given (ι) after (ι) is observed and assimilated into an observer’s situational awareness. This is accomplished through the application of Bayes theorem (adapted from Itti and Baldi 2006: 1296):

$$\forall M \in \mathbb{M}, P(M|\iota) = \frac{P(\iota|M)}{P(\iota)}P(M) \quad (1)$$

The probability distribution $P(\iota|M)$ is the probability of observing (ι) given what is known in the model space M . This is proportional to $P(\iota)$, or the *marginal likelihood* or evidence represented as (ι). Together $\frac{P(\iota|M)}{P(\iota)}$ then represents the impact of the new environmental information (ι) on $P(M)$. As such, the posterior probability of a model is determined by a combination of the inherent likeliness of a model (the prior) and the compatibility of the observed evidence with the model (the likelihood).

According to Itti and Baldi this lays the groundwork for a formal measurement of surprise: if the observation of (ι) does not affect the observer’s belief state, then (ι) is not surprising information. In Bayesian terms, the posterior distribution is identical to the prior distribution. However, (ι) is surprising if the posterior differs from prior as a result of the observation of (ι). As such, they claim that the surprise elicited by (ι) can be quantified as a measure of the difference between the posterior and prior distributions. This distance measure is obtained through the application of Kullback-Leibler (KL) divergence (also known as *relative entropy*), a widely used tool in statistics, the applied sciences, and is ideally suited to pattern recognition. In its simplest form, for distributions f and g of a continuous (or discrete) random variable, the KL divergence is defined to be the integral (Kullback 1968; Bishop 2006):³

$$D_{KL}(f||g) = \int f(x) \log \frac{f(x)}{g(x)} dx \quad (2)$$

To illustrate how this works, Figure 3 below is a visualization of KL divergence using simulated values.⁴ The KL divergence in the lower graph is computed based on the sudden change of variance in the higher graph of the mean values along the y axis over time (the x axis). The ‘size’ of the mean variance over time is measured as the spike in the KL divergence graph. This spike represents the ‘surprise’ in the mean variance.

KL divergence has been put to wide application in the applied sciences and a number of psycholinguistic experimental settings, such as speech and image recognition, determining if two acoustic models are similar, and measuring how confusable two words are (Printz and Olsen 2002; Goldberger et al. 2003; Huo and Li 2006; Silva and Narayanan 2006). Itti and Baldi (2005, 2006) designed a number of experiments that track gaze and eye movement in subjects responding to surprising visual stimuli, and applied KL divergence as a distance measure between the posterior and prior distributions conditionalized by (ι); this defines surprise as the function *Surprise* (adapted from Itti and Baldi 2006: 1296):

³An equivalent definition for discrete variables is obtained by replacing the integral with a sum.

⁴Simulated data and graphs adapted from Brooks (2012).




Figure 3: Simulated surprise generated by a sudden change in mean

$$Surprise(\iota, \mathbb{M}) = D_{KL}(P(M|\iota), P(M)) = \int_{\mathbb{M}} P(M|\iota) \log \frac{P(M|\iota)}{P(M)} dM \quad (3)$$

The computation of surprise via the function *Surprise* defined in (3) is inherently recursive through time, as the posterior in one step becomes the prior for the subsequent step. More specifically, at every iteration n , incoming data over time (ι_n) is used to update the prior $\{P(M|\iota_1, \iota_2, \dots, \iota_{n-1})\}_{M \in \mathbb{M}}$ into the posterior $\{P(M|\iota_1, \iota_2, \dots, \iota_n)\}_{M \in \mathbb{M}}$. As such, Itti and Baldi define a unit of ‘wow’ for a single model M as the amount of surprise corresponding to a two-fold variation between $P(M|\iota)$ and $P(M)$, with the total number of ‘wows’ experienced for all models obtained through the integration in (3) (2006: 1297).

The application of the *Surprise* function can be illustrated again using the scenario of the neighbour’s barking dog from the schema-theoretic analysis in section 2.1. Consider again the sample model space \mathbb{M} for this scenario in (8): if the observation of (ι) consisted of the sensory (auditory) evidence of John hearing a dog barking in his neighbour’s apartment, this would create surprise, as John’s posterior beliefs in the hypotheses that make up \mathbb{M} would be dramatically lower than the prior beliefs in these hypotheses. This results in an increased KL

distance measure between posterior and prior over all hypotheses.

Itti and Baldi (2006) adduce experimental evidence that surprise as KL divergence best characterizes what attracts human gaze in natural video stimuli of outdoors daytime and nighttime scenes of crowded environments, video games, television broadcasts, sports, and commercials. In their experiments KL scores were computed by comparing the histograms of values of saccade targets. One advantage of a probabilistic approach over a schema-theoretic analysis presented above is this quantifiability: a precise measure of surprise can be obtained via KL divergence, as conditioned by the perception of evidence (ι) in any given situation. Sketching this out, the perception of the sound of the dog barking in John's neighbour's apartment – the surprising NEI for John – would be represented by the kind of spike in KL divergence, as in Fig 3.

2.4 An evaluation

There are a number of points of contact between a schema-theoretic and Bayesian definition of surprise. Both theories provide way of representing the belief state of the individual (cognitive schema or the probability distribution $P(M)$), and both provide a theory of how this belief state is transformed upon the perception of NEI (schema revision/discrepancy or some measure of KL divergence). So, how do we decide which one is more effective in explaining linguistic surprise? I believe the answer to this rests upon two considerations: the first concerns how we think about what knowledge is.⁵ It is well beyond the scope of this paper to evaluate the conceptual underpinnings and empirical consequences of rule-based/schematic vs. probabilistic view of knowledge. Even if one were to side with a Bayesian definition of surprise there is a more pressing consideration: a Bayesian explanation of surprise is ideally suited to experimental data that can be directly observed and quantified, such as tracking saccade targets in visual scene analysis. However, how can we establish a baseline for knowledge as a probability distribution? In other words, how can we actually quantify this knowledge in order to detect a KL divergence when surprising NEI stimuli is perceived? One possible experiment would be to test for a correlation at the physiological and subjective levels (cf. 5) with surprising NEI stimuli. For example, electrophysiological activity (such as an ERP or galvanic skin response) could be monitored along with any possible linguistic presentations of surprise, and any detectable KL divergence in electrophysiological activity could be used to track for possible correlations between the two when (non)surprising NEI is presented to the subject. This is a promising avenue of research surprise. However, until the methodology and results of this kind of experiment can be developed and evaluated – and in the interests of parsimony – I argue that a schema-theoretic analysis surprise (and mirativity) offers the most explanatory coverage at present.

With a schema-theoretic analysis of surprise in place, I now turn how grammatical evidentials are use to express surprise. This begins with a brief introduction to grammatical evidentiality in Gitksan, followed by a critical evaluation of the state-of-the art views on the link between evidentials and mirativity.

⁵See Lewis (1986) for a probabilistic approach to knowledge, and McCready and Ogata (2007) and Davis et al. (2007) for a probabilistic approach to knowledge and evidentiality utilizing possible worlds instead of model spaces.

3 Evidentiality and Mirativity

It is standardly assumed that a language has an ‘evidential system’ if that language has a paradigm of grammatical elements, where each element of that paradigm uniquely encodes something about the speaker’s type of evidence or source of information for her claim. More specifically, the members of this evidential paradigm would correspond to a systematic arrangement of evidential meanings (cf. Willett 1988; Aikhenvald 2004; Palmer 2006; McCready and Ogata 2007 or Speas 2008). Aikhenvald’s (2004) evidential ontology is arranged around the *source* and *type* of verification that evidential markers provide in making an evidential statement. The source indicates who the statement is attributed to (first hand or second hand), the type indicates how the speaker obtained the information (visually, by inference, through belief, by hearsay, by quoting). Evidential types are organized by the semantic parameters given in (9) (adapted from Aikhenvald 2004: 63):

- (9)
- I. VISUAL: information acquired through seeing
 - II. NON-VISUAL SENSORY: information acquired through the senses other than seeing (i.e. olfactory and tactile)
 - III. INFERENCE: based on visible or tangible evidence or result
 - IV. ASSUMPTION: based on evidence other than visible results (logical reasoning, assumption, general knowledge)
 - V. REPORTATIVE: information in the form of a report without reference to the source of the report
 - VI. QUOTATIVE: information in the form of a report with reference to the source of the report

There is considerable cross-linguistic diversity in how languages encode these semantic parameters: in some languages these are grammatically encoded, while in others evidential meaning is attached to other grammatical and linguistic categories. Finer evidential distinctions may also be grammatically encoded in some languages. Aikhenvald also develops a language typology of systems that parameterize the evidential meanings in (9). Along with this are certain implicational relations between evidential meanings. For example, if a language grammatically encodes evidence in the source of a report, it likely also grammatically encodes information acquired through the senses or inferences based on assumption. Gitksan is this type of language.

3.1 Grammatical evidentiality in Gitksan

Gitksan has three grammatical evidentials, given in Table 1 below:

The examples in (10) constitute a minimal set using a typical stative/intransitive verb, which provides a snapshot of the typical translations of these evidentials given by Gitksan consultants:

- (10) a. mukw=hl maay̓
ripe=CND berries
“The berries are ripe.”

	Gloss	Evidence type
<i>=kat</i>	REPORTATIVE (REP)	Reportative
<i>=ima</i>	MODAL (MOD)	Inference, assumption
<i>nakw</i>	EVIDENTIAL (EVID)	Inference only

Table 1: The evidential system in Gitksan (Peterson 2010)

- b. mukw=*kat*=hl maay̓
 ripe=REP=CND berries
 “[I heard] the berries are ripe.”
- c. mukw=*ima*=hl maay̓
 ripe=MOD=CND berries
 “The berries might/must be ripe.”
 “Maybe the berries are ripe.”
- d. *nakw*=hl mukw=hl maay̓
 EVID=CND ripe=CND berries
 “The berries must be ripe.”
 “Looks like the berries are ripe.”

The meaning of *=kat* supports the standard definition of a reportative: in example (10b), a speaker has come to believe that the berries are ripe through a report from a second hand source. That source could be anonymous or unknown (perhaps a conversation overheard at the gas station), or from someone who was actually out berry picking.

At face value both *=ima* and *nakw* appear to make epistemic claims of varying force about the ripeness of the berries rather than tell us anything specific about the type of evidence a speaker has for asserting that the berries are ripe. However, there are a number of semantic and pragmatic features that distinguish *=ima* and *nakw*. The first involves what can be interpreted as modal force: *=ima* can express either epistemic possibility or necessity (commonly translated by consultants as *must* or *might*), while *nakw* can only express necessity (often translated as *must* and very rarely as *might*). The second difference regards evidence: *=ima* is compatible with any type of indirect evidence, including a simple speculative judgment about the present or future with little or no observable evidence, or a recollection of a past event where the details are no longer clear. The distribution of *nakw* is more restricted: its interpretation requires the speaker to have some type of sensory evidence for her assertion. The context in (11) brings out this difference:

- (11) Context: Gwen is sitting at home talking about going berry-picking. It’s August, and the berries are usually ripe this time of year on the Suskwa (a traditional picking grounds).

This context involves inferential evidence from speculation, general knowledge or experience. Gwen’s use of (10c) with *=ima* is felicitous with this type of evidence, while *nakw* in (10d) is infelicitous – it cannot be used when the speaker is only speculating or inferring

from general knowledge. In a related context, given in (12), where there is some kind of physical evidence available to the speaker – in this case, purple hands – both *n'akw* and *=ima* are felicitous:

- (12) Context: People are arriving home after a day of berrypicking up in the Suskwa. They're carrying buckets of berries, and their hands are all purple.

This physical evidence interpretation is supported by the frequent translation of *n'akw* as “It looks like ...”, and corresponds to Aikhenvald’s inferential evidence from some type of direct, perceptible, physical results. Thus, *n'akw* is felicitous with a specific subset of the inferential evidence types that *=ima* is. Or in other words, *=ima* is felicitous in both inferring from physical and reasoning contexts, while *n'akw* is felicitous only when inference is from sensory evidence.

Another distinguishing feature between *=ima* and *n'akw* is pragmatic: *n'akw* can be used to express mirativity. Example (13a) shows how *n'akw* can be used in a sentence that represents information which is surprising to the speaker. This effect is restricted to *n'akw*: *=ima* does not have this mirative effect in b:

- (13) Context: Alvin’s family is preparing a surprise birthday party for him. Gwen doesn’t expect him home for another hour. However, in the middle of their preparations Alvin walks through the door. In this context Gwen exclaims.

- a. *n'akw=hl witxw=s Alvin*
 EVID=CND arrive=PND Alvin
 ‘[Looks like] Alvin is here!’
- b. # *witxw=ima=s Alvin*
 arriveMOD=PND Alvin
 ‘Maybe Alvin is here!’

Finally, *n'akw* has a non-literal/metaphorical use (which *=ima* lacks), that often carries an element of surprise for the speaker: in example (14) Gwen does not mean that the batter is actually blind; rather, she is expressing her frustration that the batter keeps missing the ball:

- (14) Context: Gwen is watching a baseball game. The star batter on the speaker’s favourite team keeps missing the ball and striking out, jeopardizing the outcome of the game.

n'akw=hl sins-t
 EVID=CND blind-3

- “He must be blind!”
 “Is he blind or something?”
 “Looks like he’s blind!”

In both its mirative and non-literal uses, *n'akw* maintains its evidential function (a point I elaborate on in section 6): the speaker, Gwen, is making an assertion based on what they infer from the sensory evidence they perceive – in (13a) the fact that she can see Alvin coming through her front door, and in (14) the fact that the batter keeps missing the ball.⁶

⁶The metaphorical uses of *n'akw* are not discussed in this paper; See Peterson (2010b) for details.

3.2 Mirativity

An investigation into the overlap between evidentiality and mirativity is greatly aided by what we already know about evidentials and evidentiality (Chafe and Nichols 1986; Willett 1988; Aikhenvald and Dixon 2003; Aikhenvald 2004; Palmer 2006; Papafragou et al. 2007; Speas 2008; Peterson and Sauerland 2010; a.o.). We can also draw upon the many insights that have been gained on the mirative use of evidentials through the typological studies (cf. Aikhenvald 2004), functional analyses of mirativity in specific languages (Slobin and Aksu 1982; Aksu-Koç and Slobin 1986; DeLancey 1997, 2001, 2012; Friedman 2012; Hill 2012; Lazard 1999; Aikhenvald 2004; Temürçü 2007), and mirativity as the focus of a language description (cf. DeLancey 1990; Dickinson 2000), and also theoretical forays into the syntax and pragmatics of mirativity (Wiklund 2009; Peterson 2010b; Rett and Murray 2013).

However, the history of mirativity finds its roots in the study of evidentiality. For example, Jacobsen (1964) noted that the inferential evidential in Washo can have a semantic extension to include ‘surprise’. Although descriptions of mirativity have since appeared in various language descriptions and in the typological literature, discussions of mirativity usually begin with the observations made on the meanings and uses of the Turkish evidential *-miş* (Slobin and Aksu 1982; Aksu-Koç and Slobin 1986). If a speaker has direct knowledge of Kemal’s arrival (for example, they saw Kemal walk through the door earlier), then a speaker uses the sentence in (15a) marked with the past tense suffix *-di*. If a speaker has only *indirect* knowledge of Kemal’s arrival, perhaps by observing Kemal’s coat hanging in the front hall, or from being told that Kemal has arrived, then the speaker uses a sentence is marked with the inferential evidential *-miş*, as in (15b).

(15) Turkish (Slobin and Aksu 1982: 187)

- a. Kemal gel-*di*
Kemal came-PAST
‘Kemal came.’
- b. Kemal gel-*miş*
Kemal came-EVID
‘Kemal came.’

This can be described as the functionally unmarked use of evidential *-miş*, which is to encode the inference a speaker makes from observable or reported evidence. However, a speaker can also use (15b) if they have direct knowledge of Kemal’s arrival. The interpretative tension caused by using an inferential evidential when the speaker has direct knowledge does not lead to infelicity. Rather, the use of *-miş* in this context is a way of indicating the speaker is ‘not prepared’ for the arrival of Kemal and finds this surprising. Aksu-Koç and Slobin 1986 cast this in psychological terms, and how the alternation in (15) reflects the something of the conscious experience of the speaker: the use of *-miş* represents an experience for which the speaker lacks ‘premonitory awareness’. In section 4 I look more closely at the contextual and interactional factors that draw out mirative meaning if *-miş*.

Mirativity of this kind has been characterized using a variety of related descriptors, such as ‘non-expected’ information (Egerod and Hansson 1974), ‘surprise at unexpected new information based on immediate observation’ (Friedman 2003, p. 197, 200), ‘just discovered’ (LaPolla 2003), and ‘new knowledge’ (DeLancey 2001: 369). Dickinson (2000) distills these

characterizations of mirativity, and claims that it involves two different construals concerning how a speaker's mental state is conditioned by either their immediate experience of an event, or by previous experiences and expectations:

In a mirative system, events and states that cannot be easily assimilated are coded differently than those that easily fall in with the speaker's expectations. One construal is based on the speaker's past experiences of similar situations and his general knowledge. This set of assumptions can range from knowledge about purely physical interactions to assumptions based on cultural and social norms. The second construal is based on the speaker's immediate experience of an event or state. If the immediate situation does not correlate well with the speaker's expectations, the proposition coding the event or state receives special marking (Dickinson 2000: 379).

What is insightful here is how Dickinson's construals ground mirativity in the psychological orientation of the speaker, both in terms of their individual knowledge (their personal experiences and fixed assumptions about the way the world works), and their knowledge of the context (their immediate experience of a state, event, or action). As such, the speaker's psychological orientation in any given context can be transformed upon encountering new information. This can then lead to a sense of surprise, among other similarly related notions such as sudden realization and lack of situational awareness.

There has been a relatively recent resurgence of research into mirativity, with a focus on its categorial status from a descriptive and typological perspective. In a cross-linguistic study Aikhenvald (2012) confirms and many of the previous descriptions of mirativity in the literature, and distills them into five major subtypes of meaning:

- (16) a. Sudden discovery, sudden revelation or realization by the speaker, the audience (or addressee), or the main character
- b. Surprise of the speaker, the audience (or addressee), or the main character
- c. Unprepared mind of the speaker, the audience (or addressee), or the main character
- d. Counterexpectation of the speaker, the audience (or addressee), or the main character
- e. Information new to the speaker, the audience (or addressee), or the main character

Aikhenvald (2012) also gives a survey of how the kinds of mirative meanings in (16) can be expressed as an extension of a variety of grammatical and linguistic categories. Note, also, that the meanings in (16) would adequately describe the various expressions of surprise in the English examples in (7). These examples also show how surprise-as-mirativity can be expressed through specific kinds of structure (i.e. *wh*-exclamatives), specialized intonational contours (i.e. '!'), and lexicalized mirativity (i.e. *amazed*, *surprise*, etc.). This provides evidence for an emerging claim in the literature that mirativity is an independent grammatical category in some languages (Hengeveld and Olbertz 2012).

3.2.1 Sensory evidence

In focussing on evidentials, there are two important generalizations that condition the connection between mirativity and grammatical evidentiality: first, there is a consensus in the literature that if a language uses an evidential to express mirativity, this meaning is never an *entailed* meaning of the evidential (a point I discuss further in section 4.3); secondly, if a language possess more than one grammatical evidential, not all of the grammatical evidentials in that language can be used to express mirativity: it tends to be the evidential that encodes evidence based on sensory information; conversely, mirativity is less associated with non-sensory evidentials such as the reportative (*pace* Aikhenvald 2004, 2012). The latter generalization is developed here and supported by data from Gitksan and other languages that have mirative evidentials.

Aikhenvald (2004: 207) presents a cross-linguistic summary of the grammatical evidentials that have mirative extensions, and extends this study in Aikhenvald (2012). While this summary confirms the basic generalization that mirative meaning arises from the use of an indirect evidential in a direct evidence context (as shown in the Turkish and Gitksan examples above), what is not clear from this summary and the examples cited is what specific kinds of evidential meanings are used to express mirativity in systems that have two or more evidentials. Gitksan has such a three-term evidential system and provides an ideal testing ground for this question. Recall from section 3.1 that *nakw* has a more specialized meaning than *=ima*, as *nakw* encodes a speaker's inference from sensory evidence, whereas *=ima* can also include inferences from general knowledge or past experiences; however, example (13) shows how only *nakw* can be used miratively. This makes sense when we consider more carefully what these two kinds of information mean to a speaker in a particular speech context. Our perception and interpretation of sensory information helps us navigate and make sense of the world; it is a kind unmediated knowledge that is often new to us, revealed in an ongoing, dynamic way as a situation unfolds (NEI). On the other hand, knowledge based on conjecture or previous experience is typically not new. Furthermore, conjecture and surprise seem to be fundamentally incompatible notions (also noted in Aikhenvald 2012: 451). Therefore, non- or 'less-'sensory evidentials would seem to be ill-suited to expressing mirativity, while sensory evidentials are ideally specialized for it.

This claim is supported by observing the use of indirect evidentials in a variety of languages. In Tsafiki the indirect evidential suffix *-nu* encodes 'information inferred from physical evidence' (Dickinson 2000: 407). Adjusting the context so the speaker actually witnesses the event of the motorcycle's arrival, (17) expresses the speaker's surprise:

- (17) '[The] speaker heard what he thought was a car approaching. But when he saw it, he realized it was a motorcycle'

moto jo-*nu*-e
motorcycle be-EVID-DECL

'It's a motorcycle!' (Dickinson 2000: 411)

The speaker knows the proposition (it's a motorcycle) embedded under the evidential is true because they can perceive visually that is true. A similar effect can be observed with the evidential suffix *-k* in Qiang (LaPolla 2003), which encodes inference based on evidence

obtained visually or by some other sense. If we adjust the context such that the speaker actually witnesses the door in the state of being open, (18) expresses the speaker's surprise.

(18) 'The speaker sees that the door is open, but doesn't know who opened it'

dzy de-ʒge-ji-k
door OR-open-CSM-INFER

"The door is open!" (LaPolla 2003)

As in Tsafiki, the speaker knows the proposition (the door is open) embedded under the evidential is true because they can see that is true.

The indirect evidential particle *lō* in Hare (Athabaskan) (DeLancey 1997: 38-40, 2001: 375-378; DeLancey 2012) can also be used miratively when sensory evidence (again, visual) is available to the speaker. What is clear from the context in (19) is that the speaker actually witnesses the guy sitting in the tree; thus the use of *lō* expresses the speaker's surprise:

(19) heee, gúhde daweda! ch'ifi dachída lō
hey, up.there SG.sit.3sg.IMPERF. guy sitting EVID

"Hey, he's sitting up there! The guy is sitting up there!" (DeLancey 2001: 376)

Coming full circle with Turkish: if the speaker actually witnesses Hakan's arrival, (20) expresses the speaker's surprise using *-miş*.

(20) Hakan gel-miş
Hakan came-EVID

'Hakan came!' (primary language data)

What these examples show is that mirativity arises from the *sensory information* perceived by the speaker in these contexts: they can see (perceive visually) that the proposition embedded under the evidential is true. This is not to say that the evidentials in these languages cannot include other evidential meanings, such as past experiences with similar situations or reports from other sources etc.; rather, the claim here is that if a language possesses more than two evidentials, then the evidential that is more specialized for encoding sensory information is used to express mirativity. This claim is compatible with and sharpens Aikhenvald's (2004) summary of the mirative extensions of evidentials (p. 207). Furthermore, this analysis has the added advantage of making clear predictions that can be tested in other languages with more than two evidentials.

3.2.2 Reportatives and mirativity

One of the consequences of this claim is that reportative evidentials are not specialized for expressing mirativity. This is predicted from the analysis above: evidence in the form of a report is not acquired from the senses, and thus not specialized for expressing surprise. This is confirmed in Gitksan, as the use of *=kat* cannot express the speaker's surprise. However, this claim also challenges a number of the descriptions claiming the mirative use of a reportative evidential. Here I show that this challenge can be sustained in the following way: the descriptions of mirative reportatives in the literature do not show how a reportative evidential

can express mirativity *independently* of any other grammatical category. This shows the importance of distinguishing how a reportative evidential may be *compatible* with some other grammatical element that encodes surprise, from a reportative evidential expressing surprise on its own (as sensory evidentials do).⁷

The Quechuan languages have a three-term evidential system very similar to Gitksan, consisting of DIRECT, REPORTED and CONJECTURAL evidentials (Faller 2002; Adelaar and Muysken 2004). The reportative suffix *-sh(i)* is often cited as having mirative overtones, where the ‘sudden discovery’ kind of mirativity is most salient (cf. (16a)). However, the picture is more complex than this: reportative *-sh(i)* can occur in sentences marked with the “sudden discovery tense” suffix *-na(q)* (Adelaar 1977 cited in Aikhenvald 2012: 451):

- (22) ima-**sh** ga-**naq**, rachak-**shi** kinra-n kinra-n
 what-REP be-3/A/S.MIR toad-REP side-3.POSS side-3.POSS
 çura-naka-ra-:ri-**na(q)** çaski-yubay-si
 place-RECP-PFV-PL-3.A/S.MIR relay.runner-CMPR-ADD
 “What had actually happened? The toads had posted each other on different spots along the track as in a relay-race.”

What this example actually shows is that the reportative evidential is simply compatible with the mirative meaning encoded by *-na(q)* – the reportative does not express mirativity itself.⁸ Another commonly cited case of reportatives-as-miratives is *dizque* in some South American varieties of Spanish, and its Portuguese counterpart *diz que*. Aikhenvald (2004: 206) recounts the story of the girl who was not used to cooking. One day she decides to make *bolinho*; when asked about this the girl responds *tou fazendo bolinho diz que* (“I am making *bolinho*, it is said”), where the *diz que* is used by the girl to express her own surprise at what she finds herself doing. This shows the compatibility issue from a different perspective: whereas *diz que* is a conventional reportative in standard Portuguese, the evidential meaning of *diz que* in these dialects covers all non-first hand types of evidence (Aikhenvald 2012: 298), thus it is simply compatible with information in the form of a report. As such, it cannot be claimed that reportative *diz que* expresses mirativity on its own. Additionally, the primary extended meaning expressed *dizqueldiz que* in these cases is one of irony or sarcasm, with surprise as more of a side effect.

⁷One of the clearest examples of a reportative evidential used miratively is found in Mapudungun (Aikhenvald 2004: 200):

- (21) a. aku-**rke**-y
 arrive-REP-IND
 “S/He arrived, they say.”
 b. fey ti chi domo kalko-**rke**
 that ART woman witch-MIR
 “This woman turned out to be a witch (surprisingly)”

However, Aikhenvald (2012: 200) remarks that it has not been determined which meaning of *-rke* is the most prominent one.

⁸Another complicating factor involving the mirative and reportative data in the Quechuan languages is that most of the examples given are taken from narratives. In fact, Aikhenvald (2004: 202) acknowledges that in Quechua folk tales and traditional narratives always contain the reportative evidential, which functions as a genre marker. This observation itself neutralizes any claim of the reportative expressing mirativity on its own.

grammatical evidentiality and mirativity that requires an explanation; second, given the fact that mirativity can be expressed linguistically and grammatically in myriad of different ways, a prudent methodological move is to single out one area (grammatical evidentiality) in order to provide a clear testing ground for a theoretical explanation of mirativity. This will also provide a stable foundation for expanding the empirical scope. As such, the concern here is not whether mirativity is in fact a linguistic category or not (the consensus leans towards the affirmative), but why and how grammatical evidentials can implicate (but not entail) the meanings in (16).

4 Reconstructing Mirativity

From a broad perspective, the intake of information is a function of our every day life. The perception of states, events, and actions from the world around us feed this intake, as does the presentation of new information from every day linguistic communication. Our encounters with new information often, but not necessarily, trigger a sensation of surprise. However, a necessary ingredient for surprise is new information. This reflects the intuition that we can't be surprised by information that is 'old'. But what exactly does new information consist of? In order to clarify the connection between surprise and new information, it is important to have in place a definition of both information, and information that is 'new'. As a first step, I distinguish two kinds of information: *environmental information* and *linguistic information*.⁹

4.1 Evidentiality and environmental vs. linguistic information

The notion of environmental information is an adaptation of a notion from information theory, as defined in (24):

(24) *Environmental information:*

For an agent with prior knowledge k , two systems α and β are coupled in such a way that α 's being of type or in a state, event, or action $F(\alpha)$ is correlated to β being of type or in a state, event, or action $G(\beta)$, thus carrying for the information agent the information that β is G if the agent can infer that β is G from α being F together with k (but could not from k alone).

(Adapted from Barwise and Seligman 1997; McKinlay 2008)

Consider the following scenario as an illustration of (24): you drive by John's house to ask a favour of him, but you don't know if he is at home or not. Standing in his driveway, you can see John's light (α) which is in the state of being on ($F(\alpha)$). The state of his light provides information from the environment about John (β) in the state of being home ($G(\beta)$). This simple scenario shows how environmental information is defined relative to an observer (an information agent), who is supposed to have no direct knowledge that β is G (i.e. John is home). Rather, a kernel of information emerges from the inference that is performed that α and β are coupled in a kind of way: the fact the John's light is on provides information that John is home. This inference is performed by an agent that must have some prior knowledge

⁹This distinction (and variations on it) is of course a classic one, with roots in Bar-Hillel and Carnap (1953).

(*k*) of John's habits with conserving electricity in order for the observation that 'John's light is on' to be meaningful. This information is less meaningful to someone who lacks this prior knowledge, such as a man out walking his dog.

When an agent becomes a speaker within a speech context, this state of affairs can be linguistically realized in the one of many different ways, a sample of which is given in (25). These utterances are conditioned by the content of $F(\alpha)$ coupled with *k*, and include, for example, a simple assertion, modal auxiliaries such as *must* and *might*, evidential verbs such as *looks like* and *I see*, and the raising verb *seem*:

(25) *John is home; John must be home; John might be home; Looks like John's home; I see that John's home; It seems John's home; etc.*

As we learned about the meaning of the grammatical evidential *n'akw* from the previous section, the speaker of (26) does not know first hand whether it is true that Alvin is in fact home or not. Rather, the speaker is performing an inference based on the information available in her environment.

(26) *n'akw=hl ta'a=s Alvin*
 EVID=CND at.home=PND Alvin
 'Looks like Alvin is home.'

In languages which have grammatical evidentials such as Gitksan, this inference is encoded by an evidential, which encodes what specific type of evidence this is. In (26), the use of *n'akw* grammatically encodes the fact the speaker is making an inference from perceptual evidence (visually perceiving that Alvin's light is on).

Standing in contrast to environmental information is what I call *linguistic information*. This is intended include any kind of information that engages the speaker's language faculty. The most straightforward example of linguistic information, as defined here, are the assertions speakers make in conversational exchanges. Linguistic information can include other kinds of illocutionary forces and sentence types, or even more indirect forms of linguistic communication such as conversations overheard at a bus stop, or an article printed in a newspaper. Grammatical evidentials contrast between environmental and linguistic information. In (27), a speaker makes a statement that John is home. He believes this, not because he has any direct knowledge of John being home, but because he overheard someone at the gas station saying so. This is grammatically encoded in Gitksan using the reportative evidential =*kat* (cf. (10b)):

(27) *ta'a=kat=s John*
 at.home=REP=PND John
 '[I hear(d)] John is home.'

The distinction for mirativity of both environmental and linguistic information and the distinction between them is crucial: whereas environmental information can be used miratively (i.e. the surprise expressed through evidentiality), linguistic information typically cannot be. This connects directly to the claim made in the previous section: if a language has more than two grammatical evidentials, mirative meaning is expressed by the evidential that encodes sensory evidence (environmental information). This entails that sentences containing reportative evidentials such as (27) cannot be used to express mirativity, as information in the form of a report is linguistic information, not environmental information.

4.2 What it means for environmental information to be new

With these two definitions of information in place, the next step is to determine what it means for information to be ‘new’. In a nutshell, what we are concerned with is new *environmental* information, and not new *linguistic* information. The following series of scenarios help illustrate how the distinction between environmental vs. linguistic information, and why this is important.

Sheila and a friend have been planning on going berry-picking for some time, but she has been waiting for the berries to get ripe enough. Sheila’s mother Leiwa has just returned from the berry picking grounds and calls her on the phone; Leiwa’s says ‘The berries are ripe.’ This is new information to Sheila, but according to the definitions of information introduced above, this statement is new *linguistic* information: ‘the berries are ripe’ is an assertive speech act. The addressee interprets this sentence as a commitment on the part of the speaker that the berries are, in fact, ripe. It is an entirely linguistic transaction. In Gitksan, this assertion is expressed in (28):

- (28) mukw=hl máay
ripe=CND berries
‘The berries are ripe.’

Now consider a variation on this scenario: Sheila observes people returning from the forest with buckets full of ripe berries; their boots are muddy, and their hands are purple. Crucially, the assertion in (28) is infelicitous in this context: Sheila does not actually *know* if the berries are ripe or not. Rather, Sheila infers from the evidence – the environmental information – she observes (i.e. full buckets, purple hands). Sheila’s use of the evidential *nákw* in the sentence in (29) is felicitous in this scenario:

- (29) *nákw*=hl mukw=hl máay
EVID=CND ripe=CND berries
‘It looks like the berries are ripe.’ ‘The berries must be ripe.’

But what precisely makes environmental information new? As a basic characterization, environmental information is new when it has just come into awareness and not yet been integrated to the current knowledge state of a speaker (cf. Dickinson 2000). However, this characterization can be sharpened by identifying two conditions on environmental information that make it new, as defined in (30):

- (30) *Environmental Information* (cf. (24)) is *New* relative to an information agent *A* in a context *c* iff
- (i.) It has not been previously observed or perceived by *A* in *c*, and
 - (ii.) It is spatio-temporally bound to the context *c* as acquired by *A*.

(30) defines what I call *new environmental information* (NEI). In order to illustrate the effects of the conditions in (30), consider another scenario: Alvin is driving into the town to run some errands. He passes by Bob’s house and observes smoke rising from Bob’s *wilpsehon* (smokehouse). This observation provides environmental information (evidence) that Alvin uses to infer that Bob is smoking fish. It is also new information, as Alvin did not previously observe the smoke. In this context Alvin says to the friend accompanying him:

- (31) *n'akw*=hl se-hon-(t)=s Bob
 EVID=CND CAUS-fish-3=CND Bob
 'Looks like Bob is smoking/preparing/doing up fish.'
 'Bob must be smoking fish.'

The evidential statement in (32) satisfies the first condition of (30). The second condition on the spatial-temporal boundedness is required to account for an effect that is predominant with evidentiality and the linguistic encoding of inference that is provided by the new environmental information. Upon returning from town, Alvin again passes by Bob's house, and again observes the smoke. (32) is still felicitous, but the environment information is no longer new. Using this same scenario, the evidential sentence in (30) is markedly less felicitous if Alvin utters it to his friend just as they are returning to Alvin's house, well out of sight of Bob's smokehouse. This is because the felicity of evidential *n'akw* in (30) requires that the new information provided by the sight of smoke coming out of the smokehouse to be spatially and temporally bound to the context of the utterance of (30).

Defining NEI achieves two things that are relevant for the examination of mirativity: first, it provides a way of contrasting information that a speaker acquires through their own experience of the world (environmental information), and the information acquired by others (linguistic information). This distinction can be linguistically coded in languages with evidentials such as Gitksan, as (32) shows:

- (32) a. *mukw*=*kat*=hl *maay*
 ripe=REP=CND berries
 '[I heard] The berries are ripe.'
 b. *n'akw*=hl *mukw*=hl *maay*
 EVID=CND ripe=CND berries
 'It looks like the berries are ripe.'
 'The berries must be ripe.'

The reportative sentence in (32a) is felicitous if a speaker overhears a conversation at the coffee shop that the berries are ripe and ready for picking – the speaker is inferring the berries are ripe based on linguistic information in the form of a report – as encoded by the reportative evidential =*kat*. The sensory evidential *n'akw* in (32b) is infelicitous in a reportative context such as this. Rather, in (32b) the speaker is inferring the berries are ripe based on environmental information in the form visual evidence (i.e. people returning with muddy boots and purple hands). Because =*kat* is an evidential, we might expect it to be used miratively. However, this is not the case: (32a) *cannot* be used to express surprise. As described above, it is this interpretative tension – the use of a non-first hand evidential with the first hand knowledge acquired as sensory evidence – that drives the mirative use of an evidential.

Secondly, there is an asymmetrical relation between NEI and surprise: surprising information always entails NEI, but NEI isn't always surprising. Intuitively speaking, it is difficult to imagine a situation where some piece of information we're already aware of can also be surprising. For example, the utterance of (32b) may be no more surprising than its English translation: both involve some kind of NEI, which may not be surprising at all if the current month is August, and the berries are usually ripe at this time.¹⁰ So, why is NEI sometimes

¹⁰It is possible to recreate a sense of surprise at a piece of information that has already been mentally assimilated.

surprising?

A cognitive analysis of surprise presented in section 2 connects directly to Dickinson's (2000) construals of mirativity in the following way: we now have the theoretical means to explain the cognitive features that underpin mirativity, specifically, a speaker's current mental state (active cognitive schema), and how this state is transformed when faced with NEI that does not correlate well with the speaker's current mental state (schema discrepancy). Schemata not only provide structures for a speaker's existing network of assumptions and beliefs, but they also structure knowledge about oneself, grounded in both in the present and past experiences of similar situations, and general knowledge (i.e. *k*-knowledge). Within any context these schemata can be activated at any time and used to evaluate NEI.

It is now clearer that the term 'surprise' and the other characterizations of mirativity in (16), are just a handle for what is essentially a single mental event (in a series of events) that occurs when a schema discrepancy is triggered. This single event corresponds to a variety of related emotions and psychological notions, including (but not limited to) surprise, unpreparedness, and unexpectedness. Thus, Aikhenvald's ontology of mirative meanings in (16) now form a natural class.

An analysis of mirativity as schema discrepancy predicts that, because all of these notions are derived from a single event (the discrepancy), these notions should not be mutually exclusive from each other. For example, a mirative marker in one language may have surprise as its dominant meaning, while implicating a secondary sense of unexpectedness, or vice versa. This variability can be viewed as a reflection of the individual nature of some schemata, the universality of others, and the always unique nature of a discourse context. Additionally, because NEI is what triggers the schema discrepancy, NEI is necessarily one of the meanings of mirativity, and, in fact, in some languages it may even be the dominant interpretation. This explains the observation that surprising or unexpected information always entails NEI.

As another illustration of this prediction, consider a situation in Turkish, where Hande and her friend walk through the door and see Hakan sitting on the couch. Hande has one of two linguistic options for responding to this situation in (33):

- (33) a. Hakan bura-da, şimdi gid-ebil-ir-iz
Hakan here-LOC now go-MODAL-PRES-1pl
'Hakan's here, so we can go now.'
- b. Hakan bura-da(y)-miş, şimdi gid-ebil-ir-iz
Hakan here-LOC-EVID now go-MODAL-PRES-1p
'Hakan's here, so we can go now.'

In (33a), using the simple past, Hande expected Hakan to be there, since they made previous plans for the three of them to go out together; the fact that Hakan is there is not new information to her. In (33b), using evidential *-miş*, Hande was *not* expecting Hakan to be there. Hande and her friend had an extra ticket to the theatre, and were looking for a third

lated by a speaker. For example, one might say 'You're smoking!' as a surprised reaction (encoded by 'surprised' intonation in English) upon encountering a friend who recently stopped smoking. Having assimilated this surprising new information, one can later report this event to another friend by exclaiming 'he was smoking!' (also with the surprised intonational contour '!'). However, I believe this is akin to *reporting* surprise, and not as a reaction of surprise to the previously assimilated information.

person to go to the theatre with; since Hakan happened to be there, they can go together. In this context Turkish consultants report that the sense of surprise with this use of *-miş* in (33b) is secondary, nor is its evidential meaning prominent. Rather, in this context (33b) is simply marking that this information is new to Hande, which may further optionally implicate that this was unexpected for her. Thus, the sentence representing the proposition that *Hakan came* receives special marking using *-miş* instead of the regular past tense. Framed within the cognitive analysis of surprise, a schema discrepancy is triggered by the new information Hande encounters, as, based on her past experience, Hakan is usually at school at this time.

The evidence for the distinction between being surprised at NEI (the schema discrepancy), and the marking of a sentence as being new or surprising (mirative marking) is the one-to-many correspondence shown in the English examples in (7). In the situation of (7) there is only one schema discrepancy – the NEI of the Alvin’s arrival at the party – and many different kinds of propositions that can reflect this discrepancy. The sentences that represent these propositions are where the special linguistic marking of this discrepancy is made, which we label mirativity.

This analysis of the one-to-many correspondence between surprising NEI and the linguistics expressions of surprise also re-enforces the point made about the ‘directness’ of NEI. Returning to the scenario in (25) above, Bill can see John standing at the window of his house, which provides the NEI that John’s at home, which is then linguistically realized through the miratively marked sentence ‘John’s home!’ (perhaps because Bill expected him to be at work). However, the exact same NEI can lead to the mirative expression ‘He’s ready already!’. This NEI is instead surprising because Bill had made an appointment to go out with John, and he’s always late getting ready. This shows the one-to-many correspondence between one kind of NEI and the many kinds of discrepancies that can arise, each of these with their own mirative outcomes.

One question that arises from an analysis of evidence as NEI is: what happens to Aikhenvald’s typology of grammatical evidentials in (9)? More specifically, does an information-theoretic notion of evidence replace this typology? The short answer is no. Whereas Aikhenvald’s typology is concerned with how languages grammatically encode and organize parameterized evidential meanings, the concept of NEI is intended to operationalize these evidential meaning within in a conversational context via a schema-theoretic analysis of surprise. A source/type analysis of grammatical evidentials cannot achieve this on its own. From another perspective, a source/type analysis represents the static meaning of a grammatical evidential, and NEI represents the dynamic meaning of an evidential in conversation. It is worth noting that we see this in other areas of formal grammar. For example, propositions are the semantic core of what we call sentences. However, we can do different things with propositions in conversation, the most common of which is assertion. A declarative clause operationalizes a proposition by adding it, via assertion, to the common ground – the set of propositions that are mutually agreed upon by all of the participants in a conversation. From this perspective, we can take the categories of evidential meanings in Aikhenvald (or any other categorizations of evidential meaning) as the *static* meaning of the evidential, and NEI it’s *dynamic* meaning (Peterson 2009, 2010a).

4.3 Mirativity and surprise in the field

One shortcoming in many of the descriptions and in the literature on mirativity is an adequate description of the methodology used to test mirative meaning. This is especially crucial in languages where mirativity is an extension of another grammatical category, such as evidentiality. As such, because mirativity is an extension of a grammatical category, the null hypothesis is that it is not entailed meaning, and thus should be resistant to direct elicitation. So, what exactly are the contextual factors that drive the mirative use of an evidential (or any other grammatical category)? How can we isolate these in order to develop a field methodology for eliciting surprise?

In section 2.3 I showed how Itti and Baldi (2005, 2006) take an information theoretic approach to explaining the eye tracking patterns of subjects when they respond to surprising events in a dynamic scene (i.e. video stimuli), and I also presented a possible experiment for testing linguistic surprise using KL divergence, as, to my knowledge, there is no study devoted to measuring the linguistic correlates of surprise in an experimental setting. However, there are a number of challenges to designing such a study, the primary one being that the language of surprise is voluntary and presents itself in a myriad of ways linguistically in a language such as English (cf. 7). There are also a number of issues and challenges inherent in eliciting the language associated with surprise when it is only implicated, as in Gitksan, and not entailed as in English. For example, in a field setting the linguist can directly elicit the meaning of the surprised intonational contour (!) and the adjective 'surprised' in predicate position, as in 'I'm surprised Alvin is here!' The following tests show how speakers have clear intuitions regarding the meaning of 'surprised': in (34) its negation leads to a contradiction, and in (35) the kinds of meanings commonly associated with mirativity, such as the senses of unexpectedness and surprise in verbs such as *amazed*, can be challenged by other participants within a speech context:

(34) 'I'm **surprised** Alvin's here, #but I'm **not surprised** he's here.'

(35) A: 'I'm **amazed** Alvin's here!'

B: 'Why are you **surprised**? Didn't you know he was coming?'

The reason why the unexpectedness and surprise meanings can be challenged is because they are entailed and part of the asserted content of the sentences containing them. This is not the case with grammatical evidentials, as mirativity is not entailed by the semantics of evidentials such as *nakw*. As such, we expect that its mirative use cannot be elicited using a direct elicitation methodology such as in (34) and (35).

Consider again the party birthday party scenario from example (2) in more detail. A party at Leiwa's house is underway, and guests are arriving one by one. The party is for Alvin, who is not expected to arrive until later. There are two rooms in the house; Gwen, Sheila and Fern are in the front room, and Holly and Leiwa are in the kitchen preparing the food, separate from the front room. When Alvin shows up unexpectedly early at the door, Gwen registers her surprise with the evidential sentence in (36), which expresses the proposition that 'Alvin is here' (*p*), and that Gwen has evidence for this proposition:

(36) *nakw*=hl *witxw*=s Alvin
EVID=CND arrive=PND Alvin

Evidential interpretation: ‘Looks like Alvin is here.’

Mirative interpretation: ‘Alvin’s here!’

p = Alvin is here

$EVID(p)$ = Gwen has indirect evidence for p , and does **not know** that p is true

$EVID_{mir}(p)$ = Gwen has indirect evidence for p , **and knows** that p is true (typically by seeing that it is true)

Now, both Sheila and Fern are present in the hallway. As such, they perceive the same inferential evidence as Gwen (Alvin walking through the door), and also see that the proposition expressed by (36) is true: Alvin is here. Because of this, Sheila and Fern interpret Gwen’s use of the evidential *nakw* in (36) as registering her surprise. However, Holly and Leiwa in the adjacent kitchen *do not* see that p is true, even though they hear Gwen’s utterance of (36) (assuming that they are within earshot). Crucially, because Holly and Leiwa do not know p is true they can only interpret Gwen’s utterance of (36) as a regular evidential statement entirely lacking surprise. Thus, the mirative/non-mirative interpretation of Gwen’s utterance of (36) is simultaneously ‘split’ between two groups of discourse participants. This split is illustrated in Figure 4, where the dotted line means Holly and Gwen can hear the speaker (S), but cannot see Alvin’s arrival into the hallway (the direct evidence denoted by the arrowed line).




Figure 4: S ’s utterance of (36) and implicated (‘split’) mirativity

In other words, the speaker S in Fig 4 has the intent of communicating their surprise through using an indirect evidential in a direct evidence context. Holly and Leiwa assume that S ’s utterance of $EVID(p)$ is made in an indirect evidence, and thus *logically infer* that the speaker has indirect evidence for Alvin’s arrival. On the other hand, Sheila and Fern are aware of Alvin’s arrival (direct evidence), yet must *non-logically infer* another meaning to S ’s utterance of (36). This meaning is the surprise of S . In sum, this field experiment supports the claim that mirativity is not something just expressed, it is also interpreted.

While there is much trial and error in a field experiment such as this, it does provide a basic methodology and a replicable controlled setting for isolating the mirative use of *nakw*, one that should be adaptable to other scenarios and languages that have indirect grammatical evidentials. At the very least, it reveals how in some languages surprise is a non-entailed (non-logically inferred) meaning, and how it can be used to test the different expressions of surprise in English (cf. 7) for entailed/non-entailed meaning.

5 Mirative evidentials as deixis

A longstanding question in the evidential literature concerns *why* grammatical evidentials are used to express surprise. Aikhenvald (2004: 208, 2012: 470) proposes a number of ‘semantic paths’ that yield the mirative readings of grammatical evidentials. Two such paths are given in (37) and (38):

- (37) lack of firsthand information → speaker’s non-participation and lack of control → unprepared mind and new knowledge → mirative reading
- (38) Speaker’s deliberate non-participation → distancing effect → presenting the information as new, unexpected, and thus ‘surprising’

These semantic paths are intended to connect the speaker’s belief state and stance (or attitude) within a conversational context to mirativity. This is a reasonable move for Aikhenvald, as many of these features contribute to the meanings of grammatical evidentials in many languages. However, features such as ‘non-participation’, ‘distancing’ and ‘presenting’ also hint at the interactional quality of mirativity. I claim we can put these qualities to use by connecting mirative evidentials to *deixis*. The basic idea pursued here is that the mirative use of an evidential capitalizes on the meaning of the evidential, and uses it to direct the attention of other conversational participants. This claim has its roots in De Haan (2005), who claims that in many languages there is a connection between grammatical evidentiality and deixis by showing that there is often an overlap in the semantics of evidentials and deictic elements such as demonstratives. An example of this is Nyêlâyû (New Caledonia), which has a paradigm of suffixes that encode both a proximal/distal and visible/invisible distinction from the perspective of the speaker. As such, these suffixes perform a dual function. Aikhenvald (2004: 275) also discusses the diachronic link between grammatical evidentials, and deictic and locative marking in some languages. As such, a deictic treatment of mirative evidentials is a natural triangulation.

It is well known that morphemes can be defined as deictic if their semantic meaning is fixed but their denotational meaning varies depending on time, place and/or contextual information to convey their meaning (Fillmore 1982; Anderson and Keenan 1985; Nunberg 1993; Forbes 2003; a.o.). Pronouns such as *you*, *he*, *they*, *etc* are regarded as typical deictics, as the extension of a pronoun is dependent on the speaker using the pronoun. More specifically, there are two general usages of deixis: gestural and symbolic. Gestural deixis refers to deictic expressions whose understanding requires some sort of sensory (i.e. audio or visual) information provided by the context, such as when pointing at an object referring to it as “this” or “that”. For example, the statement “I want to eat *that* cookie” requires being able to see which cookie is being pointed at behind the counter. However, gestural deixis can include other types of information than pointing, such as direction of gaze, tone of voice, and so on. Symbolic deixis, on the other hand, generally requires only basic spatio-temporal knowledge of the utterance: “*this* is a great country to live in” only requires that the addressee knows from some other means which country the speaker is currently living in.

Mirative evidentials are also amenable to a deictic analysis in the following way: first, like other indexical expressions such as pronouns or demonstratives, evidentials have a semantic meaning that is fixed (at least of the lexical type, such as indirect, sensory, reported, etc.), yet

their meaning varies depending on time, place, the particular evidence in that context, and the person who uses it. It has been argued that certain epistemic modal expressions, such as ‘might’ and ‘must’, are indexical for two reasons: first, because they vary in the type of modality they express from context to context. Secondly, because the use of an epistemic modal in a context is restricted to what the current speaker knows as of the time of utterance (Kratzer 1977; Lyons 1977; Lewis 1979; Papafragou 2006).

Evidential statements become deictic of the gestural kind under certain conditions. Consider the direct evidential verbs in English ‘*Look!* A shooting star!’, and ‘*I see* a shooting star!’, which have the function of pointing out or drawing the addressee’s attention to a shooting star. Indirect evidentials such as *nakw* also take on a gestural deictic meaning under similar circumstances. In the mirative context of (2) above, both the speaker and addressee can see that the state, event or action embedded under the evidential is true: Alvin has arrived. The advantage of this approach is that the original meaning of an evidential is not undermined when it is used miratively, rather, this extended mirative meaning is used for pointing out what the speaker is surprised at, and for drawing the attention to other conversational participants to this.

6 Summary and Future Directions

This paper makes a twofold contribution: first, it contributes to our understanding how emotions such as surprise are linguistically realized – a relatively understudied area. Lexicalist studies have shown that there is considerable intra- and cross-linguistic variation in what kinds of words and structures are used to express core emotions (cf. Harkins and Wierzbicka 2001). For example, it is notable that Gitksan lacks many of the expressive strategies for surprise that are available in English in (7). Whereas English has a plethora of verbs, adjectives and nouns that entail surprise (as revealed by applying the entailment tests in (35)), a Gitksan speaker can only implicate surprise through the use of an indirect grammatical evidential. Conversely, this is grammatical strategy that English lacks, as English does not possess a system of grammatical indirect evidentials. Even words that have an indirect evidential quality like epistemic *must* or the raising verb *seems* cannot naturally be used to express surprise.

The second contribution is made to our understanding of evidentiality and evidentials, through a theoretical explanation of how grammatical evidentials are used to express mirativity. This was done by taking seriously how mirativity is defined through the psychological and cognitive state of the speaker. Not all grammatical evidentials are created equally in this regard: I showed how evidentials that encode sensory information are specialized to express mirativity, which predicts that reportatives cannot. So far, this prediction is borne out in the data. This theoretical explanation, coupled with the detailed analysis and description from the typological/functional literature provides an ideal platform for further research into mirativity. For example, the field test for the mirative use of an evidential confirms that mirative meaning is not entailed by the meaning of the evidential. This strongly suggests that it is *conversationally implicated* meaning. This is a clear theoretical hypothesis which in turn invokes a number of well-defined testable predictions. Is mirativity a cancelable implicature? However, the flip side of this is languages that grammatically encode mirativity: if the mirative meaning is ‘fixed’ to a word or morpheme, we expect it to be subject to tests for entailment (or possibly

conventionally implicated meaning). These are the core (non-theoretical) semantic tests that are required to support any analysis of meaning in any language. We are also in a better position to examine more closely the interactional qualities of mirativity that Aikhenvald outlines in (16): in some languages mirativity is also about the surprise of the interlocutors within a conversational context. This aspect of mirativity was not explored in this paper, but it could bear upon current research on how grammatical evidentials are used by speakers in communicative interactions to negotiate participation and responsibility (Michael 2008). Finally, the stage is set for taking mirativity into the lab and connecting the language of surprise to the physiology of surprise.

Surprise is one of the few universal human emotions. Likewise, mirativity is a linguistic universal: all languages have the means for linguistically encoding mirativity. This paper represents one of the first studies dedicated to the theoretical analysis of mirativity, and also supports this hypothesis. This study is important because of the bridge it creates – which, at the moment, does not exist – between the linguistic encoding of surprise and the purely psychological and cognitive streams of research on the emotion of surprise.

References

- Adelaar, W. F. H. (1977). *Tarma Quechua: grammar, texts, dictionary*. Lisse, De Ridder.
- Adelaar, W. F. H. and Muysken, P. (2004). *The languages of the Andes*. Cambridge, CUP.
- Aikhenvald, A. and Dixon, R. M. W., editors (2003). *Studies in Evidentiality*. John Benjamins, Philadelphia.
- Aikhenvald, A. Y. (2004). *Evidentiality*. Oxford University Press, New York.
- Aikhenvald, A. Y. (2012). The essence of mirativity. *Linguistic Typology*, 16(3):435–486.
- Aksu-Koç, A. and Slobin, D. (1986). A psychological account of the development and use of evidentials in Turkish. In Chafe, W. and Nichols, J., editors, *Evidentiality: The linguistic coding of epistemology*, Advances in Discourse Processes XX, pages 159–167. Norwood, N J: Ablex.
- Anderson, S. R. and Keenan, E. L. (1985). Deixis. In Shopen, T., editor, *Language typology and syntactic description: Grammatical categories and the lexicon*, volume 3, pages 259–308. Cambridge: Cambridge University Press.
- Bar-Hillel, Y. and Carnap, R. (1953). *An Outline of a Theory of Semantic Information*. reprinted in Bar-Hillel [1964].
- Barwise, J. and Seligman, J. (1997). *Information Flow: The Logic of Distributed Systems*. Cambridge University Press, Cambridge.
- Bishop, C. (2006). *Pattern Recognition and Machine Learning*. Springer.
- Brooks, W. (2012). Bayesian surprise as a tool for monitoring sensor networks. Paper presented at the NWQMC 8th National Monitoring Conference.

- Chafe, W. L. and Nichols, J., editors (1986). *Evidentiality: the Linguistic Coding of Epistemology*. Advances in Discourse Processes XX. Norwood, N J: Ablex.
- Davis, C., Potts, C., and Speas, M. (2007). The pragmatic values of evidential sentences. In Gibson, M. and Friedman, T., editors, *Proceedings of the 17th Conference on Semantics and Linguistic Theory*, pages 71–88. CLC Publications, Ithaca, NY.
- De Haan, F. (2005). Encoding speaker perspective: Evidentials. In Frajzyngier, Z. and Rood, D., editors, *Linguistic diversity and language theories*. Benjamins, Amsterdam.
- de Reuse, W. J. (2003). Evidentiality in western apache. In Aikhenvald and Dixon, editors, *Evidentiality in Typological Perspective*, pages 79–100.
- DeLancey, S. (1990). Notes on evidentiality in Hare. *International Journal of American Linguistics*, 56(1):152–158.
- DeLancey, S. (1997). Mirativity: The grammatical marking of unexpected information. *Linguistic Typology*, 1:33–52.
- DeLancey, S. (2001). The mirative and evidentiality. *Journal of Pragmatics*, 33:369–382.
- DeLancey, S. (2012). Still mirative after all these years. *Linguistic Typology*, 16(3):529–564.
- Dickinson, C. (2000). Mirativity in Tsafiki. *Studies in Language*, 24:379–422.
- DiMaggio, P. (1997). Culture and cognition. *Culture and cognition. Annual Review Of Sociology*, 23:263–287.
- Egerod, S. and Hansson, I. L. (1974). An Akha conversation on death and funeral. *Acta Orientalia*, 36:225–284.
- Ekman, P. (1980). Asymmetry in facial expression. *Science*, 209:833–834.
- Ekman, P. (1984). Expression and the nature of emotion. In *Approaches to Emotion*, pages 319–344. Hillsdale, New Jersey: Lawrence Erlbaum.
- Faller, M. (2002). *Semantics and Pragmatics of Evidentials in Cuzco Quechua*. PhD thesis, Stanford.
- Fillmore, C. J. (1982). Towards a descriptive framework for spatial deixis. In Jarvill, R. J. and Klein, W., editors, *Speech, place and action: Studies in deixis and related topics*, volume 10, pages 31–59. London: Wiley.
- Forbes, G. (2003). Indexicals. In Gabbay, D. and Guenther, F., editors, *Handbook of Philosophical Logic*, volume 10. 2nd edition.
- Friedman, V. A. (2003). Evidentiality in the Balkans with special attention to Macedonian and Albanian. In Aikhenvald, A. Y. and Dixon, R. M. W., editors, *Studies in Evidentiality, Typological Studies in Language 54*. Amsterdam/Philadelphia: John Benjamins.

- Friedman, V. A. (2012). Perhaps mirativity is phlogiston, but admirativity is perfect: On Balkan evidential strategies. *Linguistic Typology*, 16(3):505–528.
- Goldberger, J., Gordon, S., and Greenspan, H. (2003). An efficient image similarity measure based on approximations of KL-divergence between two gaussian mixtures. In *Proceedings of ICCV 2003, Nice, October 2003, vol. 1*, page 487493.
- Harkins, J. and Wierzbicka, A., editors (2001). *Emotions in Crosslinguistic Perspective*. Walter de Gruyter. Copyright.
- Hengeveld, K. and Olbertz, H. (2012). Did you know? Mirativity still exists! *Linguistic Typology*, 16(3):487–503.
- Hill, N. (2012). “Mirativity” does not exist: *hdug* in “Lhasa” Tibetan and other suspects. *Linguistic Typology*, 16(3):389–434.
- Huo, Q. and Li, W. (2006). A DTW-based dissimilarity measure for left-to-right hidden Markov models and its application to word confusability analysis. In *Proceedings of Interspeech 2006 - ICSLP, Pittsburgh, PA*, page 23382341.
- Itti, L. and Baldi, P. F. (2005). A principled approach to detecting surprising events in video. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 631–637, San Diego, CA.
- Itti, L. and Baldi, P. F. (2006). Bayesian surprise attracts human attention. In *Advances in Neural Information Processing Systems, Vol. 19 (NIPS*2005)*, pages 547–554, Cambridge, MA. MIT Press.
- Izard, C. and Malatesta, C. (1987). Perspectives on emotional development I: Differential emotions theory of early emotional development. In *Handbook of infant development*, pages 494–554. Wiley; New York.
- Jacobsen, W. (1964). *A grammar of the Washo language*. PhD thesis, University of California, Berkeley.
- Kagan, J. (2002). *Surprise, Uncertainty, and Mental Structures*. Harvard University Press.
- Kleider, H. M., Pezdek, K., Goldinger, S. D., and Kirk, A. (2008). Schema-driven source misattribution errors: Remembering the expected from a witnessed event. *Applied Cognitive Psychology*, 22:1–20.
- Kratzer, A. (1977). What ‘must’ and ‘can’ must and can mean. *Linguistics and Philosophy*, 1:337–355.
- Kullback, S. (1968). *Information Theory and Statistics*. Dover Publications Inc., Mineola, New York.
- LaPolla, R. (2003). Evidentiality in Qiang. In Aikhenvald, A. Y. and Dixon, R. M. W., editors, *Studies in Evidentiality*, Typological Studies in Language 54. Amsterdam/Philadelphia: John Benjamins.

- Lazard, G. (1999). Pour une terminologie rigoureuse: quelques principes et propositions. *Mémoires de la Société de Linguistic de Paris*, (6):111–133.
- Lewis, D. (1979). Scorekeeping in a language game. *Journal of Philosophical Logic*, 8(1):339–359.
- Lewis, D. (1986). *Philosophical Papers, Vol. 2*. Oxford University Press, New York.
- Lorini, E. and Castelfranchi, C. (2007). The cognitive structure of surprise: looking for basic principles. *Topoi: An International Review of Philosophy*, 26(1):133–149.
- Lyons, J. (1977). *Semantics. 2 vols.* Cambridge University Press.
- Macedo, L., Cardoso, A., Reizenzein, R., Lorini, L., and Castelfranchi, C. (2009). Artificial surprise. In Vallverdú, J. and Casacuberta, D., editors, *Handbook of research on synthetic emotions and sociable robotics: New applications in affective computing and artificial intelligence*, pages 267–291. Hershey, PA: IGI Global.
- Mandler, G. (1984). *Mind and body*. Norton, New York.
- Matthewson, L., Davis, H., and Rullmann, H. (2007). Evidentials as epistemic modals: Evidence from St’át’imcets. *The Linguistic Variation Yearbook*, 7.
- McCready, E. and Ogata, N. (2007). Evidentiality, modality and probability. *Linguistics and Philosophy*, 30(2):147–206.
- McKinlay, S. (2008). Information, knowledge and confirmation holism. In Brey, P., B. A. and Waelbers, K., editors, *Current Issues in Computing and Philosophy (Collected Essays)*. IOS Press.
- Meyer, W.-U. and Niepel, M. (1994). Surprise. In *Encyclopedia of Human Behavior*, volume 3. California: Academic press.
- Meyer, W.-U., Reizenzein, R., and Schützwohl, A. (1997). Toward a process analysis of emotions: The case of surprise. *Motivation and Emotion*, 21(3):251–274.
- Michael, L. (2008). *Nanti evidential practice : language, knowledge, and social action in an Amazonian society*.
- Niepel, M., Rudolph, U., and Schützwohl, A. (1994). Temporal characteristics of the surprise reaction induced by schema discrepant visual and auditory events. *Cognition and Emotion*, 8(5):433–452.
- Nunberg, G. (1993). Indexicality and deixis. *Linguistics and Philosophy*, 16:1–43.
- Palmer, F. (2006). *Mood and Modality*. Cambridge University Press, Cambridge, 2nd edition.
- Papafragou, A. (2006). Epistemic modality and truth conditions. *Lingua*, 116:1688–1702.
- Papafragou, A., Li, P., Choi, Y., and hye Han, C. (2007). Evidentiality in language and cognition. *Cognition*, 103:253–299.

- Peterson, T. (2009). Pragmatic Blocking in Gitksan Modality and Evidentiality. In A. Schardl, M. W. and Abdurrahman, M., editors, *Proceedings of the 38th Meeting of the North East Linguistic Society*, pages 219–232. Amherst, Mass: GLSA Publications.
- Peterson, T. (2010a). *Epistemic Modality and Evidentiality in Gitksan at the Semantics-Pragmatics Interface*. PhD thesis, University of British Columbia.
- Peterson, T. (2010b). Examining the mirative and non-literal uses of evidentials. In Peterson, T. and Sauerland, U., editors, *Evidence from Evidentiality*, volume 28. Vancouver: University of British Columbia Working Papers in Linguistics.
- Peterson, T. and Sauerland, U., editors (2010). *Evidence from Evidentiality*, volume 28. Vancouver: University of British Columbia Working Papers in Linguistics.
- Printz, H. and Olsen, P. (2002). Theory and practice of acoustic confusability. *Computer, Speech and Language*, 16:131164.
- Reisenzein, R. (2000). The subjective experience of surprise. In Bless, H. and Forgas, J., editors, *The message within: The role of subjective experience in social cognition and behavior*. Philadelphia, PA: Psychology Press.
- Reisenzein, R., Meyer, W.-U., and Schützwohl, A. (1996). Reactions to surprising events: A paradigm for emotion research. In Frijda, N. H., editor, *Proceedings of the 9th conference of the international society for research on emotions*, pages 292–296. Toronto: ISRE.
- Rett, J. and Murray, S. E. (2013). A semantic account of mirative evidentials. In Snider, T., editor, *Proceedings from Semantics and Linguistic Theory (SALT) XXIII*, pages 453–472, Ithaca, NY. CLC Publications.
- Rumelhart, D. E. (1984). Schemata and the cognitive system. In Wyer, R. S. and Srull, T. K., editors, *Handbook of social cognition (Vol. 1)*, pages 161–188. Hillsdale, NJ: Erlbaum.
- Scherer, K. R. (1984). On the nature and function of emotion: A component process approach. In *Approaches to Emotion*, pages 293–317. Hillsdale, New Jersey: Lawrence Erlbaum.
- Schützwohl, A. (1998). Surprise and schema strength. *Journal of Experimental Psychology, Learning, Memory and Cognition*, 24(5):1182–1199.
- Silva, J. and Narayanan, S. (2006). Average divergence distance as a statistical discrimination measure for hidden markov models. *IEEE Transactions on Audio, Speech, and Language Processing*, 14(3):890906.
- Slobin, D. and Aksu, A. (1982). Tense, aspect, and modality in the use of Turkish evidential. In Hopper, P. J., editor, *Tense-Aspect: Between Semantics and Pragmatics*, Typological Studies in Language, pages 397–405. Amsterdam: John Benjamins.
- Speas, M. (2008). On the syntax and semantics of evidentials. *Language and Linguistics Compass*, 2(5):940–965.

- Temürcü, C. (2007). *A Semantic Framework for Analyzing Tense, Aspect, and Mood – An Application to the Ranges of Polysemy of -Xr, -DIr, -Iyor and -Ø in Turkish*. PhD thesis, Universiteit Antwerpen.
- von Stechow, P. and Gillies, A. (2007). An opinionated guide to epistemic modality. In Gendler, T. S. and Hawthorne, J., editors, *Oxford Studies in Epistemology Vol. 2*, pages 83–101. Oxford University Press.
- Wiklund, A.-L. (2009). The syntax of surprise: Unexpected event readings in complex predication. In *Working Papers in Scandinavian Syntax 84*, pages 181–224.
- Willett, T. (1988). A cross-linguistic survey of the grammaticization of evidentiality. *Studies in Language*, 12(1):51–97.