

On Errors, and Other Things

What do we want to know?

The goal of the class is to learn about the structure of the language under investigation, which we will express in the form of general statements about the language. At some point, the skeptical question may arise “How do we know that this is what’s really going on?”. The purpose of this essay is to say a little about the “How” question. This essay focuses on the fact that statements may not be completely accurate; my point here is that rather than despair over our doom because errors are possible, clever scientists do their best to make sure that their statements are as accurate as they can be. How do you do that? By not doing things that cause errors. So I think it helps to spend a little time thinking about what causes errors. Hence this essay.

It’s probably obvious what an “error” is, but just to be clear about this, it’s where a statement describes a particular situation, only that situation isn’t “real” or isn’t a fact. A simple example is the statement “English has no voiceless velar fricative [x]”, which describes a situation regarding English pronunciation — only, what really happens in English is different from this, because of what English speakers do when they pronounce “Loch” [lɔx], “Bach” [bax] and “masks” [mæsx]. Data such as [lɔx] contradict the statement, and our statement is supposed describe the data. There are two aspects to finding this error. The first is grasping the predictions of the statement, i.e. [lɔx] shows that the statement is wrong. That seems pretty straightforward in this example and we won’t concentrate on figuring out this correspondence between statements and facts. The second aspect, which we will concentrate on here, is the data. In that case, someone should ask, what is data (singular *datum* probably for a couple more generations)?

A datum, such as a blip on an electron microscope, a formant value, a temperature, or a transcription, is a representation of something in reality. It’s not the same as “the thing itself”, but it has a systematic relationship to the thing. For speech, the bare thing itself would be an event of speaking, which happens and then is permanently done with. With a pronunciation, an audio recording may give you a reasonably accurate representation, but the recording is still not the thing itself. For the non-phonetic aspects of the language thing — meanings for instance — we are extremely far from being able to create comparable representations of brain states. There is no problem in dealing with a representation of the thing rather than the thing itself, as long as the representation of the thing is accurate. If you had an accurate representation of the thing, you could in principle give an objective rule stating the relation between the thing itself and its representation, where every property of the thing is in the representation, and the object can be perfectly reconstructed from the representation. However, there is always something about the thing itself which isn’t in the representation (even in an audio recording) — representations are always inaccurate to some extent. We want that inaccuracy to be minimal.

Now we get to the normative part of the essay. The pure theoretical analyst deals with existing data, which are presumed to be facts, and provides an interpretation of the data in terms of general linguistic laws. The theorist’s arguments can be inspected and evaluated by anyone in terms of how well they match the known data. The field worker is in a special position compared to the theorist, because the field worker creates the data¹ that is the grist for the theoretical mills: the data are typically presumed to be correct. A field worker who presents data to the world implicitly war-

¹Does it seem strange to say that the field worker “creates the data”? Perhaps, but it is an accurate way to phrase it, since data do not just magically materialize on the pages of journals and books. Someone has to take raw language events and make them into data.

rants that he has taken whatever precautions are necessary to make the data as accurate as possible. Not perfectly accurate, which is impossible, but correct to the degree of accuracy that is possible with the technology being used.

Science is incremental, in that researchers assume (unless there is compelling reason to think otherwise) that existing data are indeed accurate, for which reason it is unnecessary to completely reproduce all prior experiments before adding something new to what we know. But suppose that the data which you present to the world are highly inaccurate, well beyond the level of accuracy that prudent researchers can easily reach. Then general conclusions (theoretical principles) drawn on the basis of such data will have significantly reduced validity, and will not reflect reality: this is counter to the purposes of science. It is therefore essential to the progress of science that you make sure your data are reliable.

Suppose you wish to know if subjects precede or follow verbs in a language. You find a speaker and ask “How to you say ‘The child wept?’”. The speaker replies *alilay anaamwo*, and says “*alilay* means ‘wept’ and *anaamwo* means ‘child’”. You rush to publish this fact, and for some odd reason this fact becomes important to linguistic theory, to the point that many theories are rejected and departments are closed, because modern theories can’t explain these data. In reality, the speaker spoke two different sentences of single words, and that the correct translations are “(s)he will sing”; “it was a dwarf”. In this language, the way you really say “The child wept” is *mono yemba*, where *mono* is the proper word for “child” as subject, and *yemba* means “wept”. The data which you presented were very inaccurate, and the consequences for the discipline were far-reaching.

Unfortunately for you, this horrid error was eventually revealed as a result of later research on the language by the careful field worker, Smith, who published a series of mocking articles that rightly ridiculed you for your slovenly research, and you end up being laughed out of the business. Furthermore, your error had horrific consequences for other linguists, who rejected perfectly good theories of language because those theories *seemed* to make incorrect predictions about this language. This disaster would have been averted if it were known that your data were wrong — but how do we (the community of linguists) know this in advance? You could criticise the editors and reviewers and say that they should have caught this, but that’s only reasonable if the editors and reviewers actually know something about the language. That is rarely the case — there are plenty of experts on English, French and German, very few experts on Zulu, Navaho or even Arabic, and for most languages (Kikerewe, Konni, Campa) only one person — the author — has direct experience with the language. We therefore trust that the fieldworker’s data are as accurate as feasible. In this case, that trust was misplaced: you did not take appropriate steps to assure that your statement accurately describes the facts of the language. What should you have done; what was your error?

You forgot a fundamental fact of science, that errors are possible; and as the creator of the data, you have a responsibility to be sure they are accurate before you unleash these data on the world. As it happened, when you asked the speaker for this example, he was having a bad hair day and just was not paying any attention. If you had rechecked the data (knowing that errors are possible), your informant might have then provided you with the correct form *mono yemba* which you didn’t know about; you would have discovered the error in translating the word meaning “it was a dwarf” as the general noun meaning “child” if you had looked at more examples involving “child”. You could have also caught the tense-aspect anomaly (translating a future tense as a past tense), if you had investigated the tense-aspect system of the language more carefully, and noticed that your speaker is somewhat confused about the meanings of tenses in English.

So how do you make sure that your data are accurate? The answer lies in another fundamental principle of science, that the universe obeys consistent laws, and therefore experimental re-

sults should be reproducible. If the sentence “The child wept” really were *alilay anaamwo*, you should be able to get that result again and again — any number of times — when you ask a speaker this sentence. If you can’t reproduce the results, you must suspect that the datum is an error (which would prevent you from publishing it as accurate, thus saving your career and the whole enterprise of linguistics). Field workers spend a lot of time rechecking, to make sure that their data accurately represents the language. If you get unexplained variation in your data, you should suspect that there is an error of some sort. The error might be extremely subtle and yet revealing — for example apparently random switching between forms might mean that the language makes a distinction which is not accurately being represented in the English translations that you are using (e.g. in Kikerewe there are different past tense for things that happened yesterday vs. before yesterday, and my initial glossing of *akatééka* as meaning the same as *ateekilé* was an error).

Who should you blame?

Fieldwork involves at least two people, the analyst (the linguist) and the language consultant, either of whom can contribute to error. The linguist can contribute in many ways (by mumbling, being vague, not understanding the potential complexity and nuances of a question, incorrectly interpreting a response, being obnoxious...). Linguist-contributed unreliability is to be avoided, and in the ideal world can be controlled for perfectly through sheer willpower (see the section on perfect control and sheer willpower, t.b.w.). The language consultant can also be a source of errors; the linguist strives to recognise these problems, work around them, understand how they arise and how they influence the probability of having correct knowledge.

Informants are not evil people if they generate inconsistent data, nor are linguists sloppy scientists if their actions accidentally contribute to errors. My guesstimation is that most inconsistencies (after the initial stage where transcription errors abound and the linguist almost always has to take the rap for it) are due to linguist/speaker interactions — you don’t know exactly what to ask or how to ask for it effectively, and the speaker is doing her best to figure out what you’re “really” looking for, but guesses wrong some of the time. Assigning blame doesn’t help the situation (that answers *that* question). Even if the informant is the one who makes the mistake, it’s the responsibility of the linguist to ascertain that the data presented to the world are correct: so we can automatically blame the linguist for not having caught the mistake. While we’re assigning blame, it will also make things easier (for our discussion) if we attribute the source of the error to the informant. Pure linguist’s mistakes are easier to detect and stamp out, though doing so may require professional help. The hands-down most common unambiguous linguist mistake is in transcription; so I’ll just dismiss it for the moment by saying that it take lots of practice with languages to get to the point that your phonetic transcriptions are reliable. Some people can do good transcriptions fairly easily, other people are, well, nearly hopeless. There are other mistakes that linguists make, such as writing the example down wrong, or writing the translation down wrong.

A fundamental distinction made in linguistics is the opposition *langue* ~ *parole*, competence ~ performance, or perhaps language ~ speech. An assumption made by very many linguists² is that speakers of languages have ‘grammars’ in their heads. Grammars are the cognitively real things in heads which provide an explanation for a substantial portion of human linguistic behavior. The key here is ‘substantial’: there are other factors which interact with grammars in explaining the entirety of human linguistic behavior. To use the standard terms, grammars are (definitionally) about ‘com-

²I have no idea if the percentage is 99%, 51% or less than that, but I’m sure it’s more than 25% which justifies the term ‘very many linguists’.


petence’, and ‘performance’ is about how the grammar actually is applied (both what is produced, and a theory of the relation between what is produced and linguistic competence).

In describing a language, we are trying to figure out what the grammar looks like, and our evidence is the behavior of speakers. If nothing else got in the way, speakers would be perfectly consistent and their behavior would perfectly reflect the nature of their internalized mental grammar. Things do get in the way. The sole cause of error is ‘performance’: grammar are never in error. If only we could get rid of that pesky performance (and get rid of gravity while we’re at it). A grammar is like the ideal frictionless plane, or a quark: you never see one directly, but it exists. In terms of competence vs. performance, what then is an error? It’s a mismatch between what the grammar would give, and what the speaker actually does. For linguistic fieldwork where the speaker tells the linguist “Here is how you say X”, you have to add the fact that the speaker must comprehend the linguist’s question in the manner that the linguist intended it, and usually also has to find the conceptual correspondence between the intent in the language of elicitation, and the language being investigated. So there are plenty of opportunities to introduce error.

What do we not want to account for in a grammatical description?

There are many extragrammatical factors which influence linguistic behavior, and they are not exclusively factors pertaining to language. An understanding of these factors is not part of the study of grammar, but it is crucial to knowing what *is* in the grammar, versus not in the grammar; it may also be important to understanding how a grammar developed historically to become the way it is at present. Certain performance factors are well-known (I hope). It is very likely that the sentence “The horse the man the child I knew saw bought died” will be rejected by nearly all speakers of English (I have heard that some speakers, ones with graduate training in linguistics, have no problems with such sentences and produce them frequently). Although the sentence is rejected by normal speakers, it nevertheless follows the rules for construction of sentences. Consider the relation between the sentences “I knew the child”, “The child saw the man” and “The child I knew saw the man”. The relation can be stated as one for combining two sentences into one complex one: $N_i V_i N_j + N_j V_j N_k \rightarrow N_j N_i V_i V_j N_k$. All we have done in the sentence “The horse the man the child I knew saw bought died” is apply that rule three times.

Why is this sentence rejected? Rejection / acceptance of sentences is different from grammaticality / ungrammaticality. A sentence is grammatical if and only if it is produced by the grammar. A sentence is acceptable if and only if it is accepted. Grammaticality is an idealization pertaining to abstract grammar; acceptability is about speaker behavior which is related to the structure of the grammars, as well as other factors. Grammars can produce sentences that are rejected by speakers. It would be a crazy theory of grammar which did not allow this unacceptable sentence to be generated, since the restrictions on the rule of relative clause formation would end up being very complex, possibly unstatable in any theory. The reason that we can get away with this apparent hand-waving is that there is an explanation for the speaker’s negative reaction that comes from outside of the theory of grammar itself. The problem stems from the fact that it is very difficult to undo the tangle of references which allows you to figure out that “the child” is the subject of “saw” and the object of “knew”, represented schematically as:



The presumption — one that has carried us a fair distance in this discipline — is that an independent theory of sentence processing can actually make explicit the reason why such structures are hard to unravel, and thus are rejected.

Even among theoreticians who work on the theory of grammar, there isn't uniform agreement on what is in the domain of grammatical theory and description. Another example of a factor that most linguists would treat as being outside of the theory of grammar is anything that depends crucially on real-world truth. Thus, the sentence "Earth is the fifth planet from the sun" is false, but it is meaningful and grammatical, and a description of English should not care that in reality Earth is the third planet in the solar system, as far as I'm concerned. Nor should a grammar say that the sentence "My toothbrush is trying to kill me" is ungrammatical because toothbrushes are non-sentient and "try" requires a sentient subject. It is reasonable for a description of English to report that the subject of "try" must be sentient, but this is not a matter of grammar, rather it's a matter of the meaning of "try". But remember: as a fieldworker, you don't have direct access to what's in the grammar, rather you try to get at what's in the grammar based on what speakers say about certain sentences, which can reflect grammar, real-world knowledge, or psycholinguistic processing problems. If an informant rejects the sentence "We ate₃ the mice" — this is a language with a lot of verbs for eating whose selection is determined by all sorts of complex things having to do with syntax and tense-aspect — you need to consider whether the rejection is because of the grammar of eat₃, or the fact that these people don't eat mice.

Another example of 'noise' introduced into data by extragrammatical factors is data caused by the existence of other people in the world. Some people pronounce "route" as [rawt] and others pronounce it as [ruwt]; some people say "ain't" as the negative present tense copula; some people can say "Tom and me went shopping" and others say "Tom and I went shopping". My grammar does not allow such expressions as "stand on line", "this car needs washed", "boiled cod, I wouldn't eat" or the pronunciation [ru:t] for "root" (the correct pronunciation is [rot]). I have become familiar with these structures which are used by other speakers of English, I understand them, and it is possible that I might even utter such things given high enough levels of fatigue and social pressure, but if you are writing a grammar of my language, you don't want to include an erroneous production of the word [ru:t] as part of the grammar of my language, since it's not in my language.

It is unreasonable to expect a speaker of 'English' to know whether there exists any speaker of English anywhere in the world who will accept such-and-such an utterance (though it's a good bet that almost any speaker of English would recognise that the sentence *omukama yaafwiile* is not a grammatical sentence of any dialect of English). The only thing that a speaker can be sure of is their own grammar (and as you know, they don't consciously know the grammar itself). A speaker might make a reasonable guess about other people, but they can be surprisingly wrong. I was surprised to discover, even after being a linguist for over decades, that my own mother actually made a distinction in the pronunciation of the vowels in "cot" and "caught". If you had asked me in 1992 whether my mother pronounced the words the same I would have said 'yes', and I would have been wrong. The flip side of the coin is, alas, that it's also unreasonable to expect a speaker of a language to be completely ignorant of what other speakers of the language do; for this reason, you might find a speaker who says "you could pronounce that as [ru:t] or [rot]", even if it turns out that they say [rot] but they've heard the other pronunciation which is why they pretend there is a choice.

Since we're after the internal grammar of the consultant, do not be highly concerned, at least right now, about whether your speaker is "representative". There may be some question whether your language consultant is "hopeless" — more about this later — but even if she/he is "not typical", we don't care, because the goal here is not a sociological survey. If you get to the point of

writing a reference grammar then you will be a bit more concerned about sociological issues relating to conformity to norms. But a mental grammar is a mental grammar and is a valid object of study, no matter how much it diverges from the norms imposed by certain sectors of society.³

There is a more controversial question whether there are structures that speakers ‘accept’ which are not generated by the grammar. In one sense, this possibility must be allowed because speakers can quite often deduce the intended meaning from an ungrammatical utterance especially as produced by a non-speaker of the language, a fact that has saved my bacon quite a number of times. An exemplar is the sentence “This man, go go take knife hold goat cut neck blade give meat people”, which you may assume means that somebody snuffed a goat and gave away the meat. That seems pretty obvious; however, in the fieldwork context there will come a point when you ask “Can you say [waka waka lompe fiske]?”, and the informant says “Yes, you could say it”, but what she means is “I guess if you said that, people would figure out what you meant, although clearly it sounds like something said by somebody that doesn’t know the language: the correct expression is [waka lompe waka fiske raspeboller]”. That is: the answer “yes” may not mean “it is acceptable in the language”, but rather “Wow, mzungu speaking our language! How cool is that?!”.

It has also been proposed that the grammar of English does not generate sentences such as “Tom’s difficulty to understand makes him a bad teacher”, and that acceptance of such sentences by some speakers of English is due to ‘performance factors’. In this case, though, it is also plausible that the underlying theory of grammar is wrong, since it is not at all difficult to construct a grammar that generates such a sentence — it’s just a problem for some specific theories of grammar. We will not worry about the latter problem: we will very much worry about the former problem, that of speakers ‘accepting’ sentences which they in fact only believe to be “somewhat interpretable” and are actually not “good” in the language.

We won’t go into great detail about the more philosophical issues about what a grammar should or should not contain. The most important consideration for us is simply this: we only have access to linguistic behavior, therefore to guess at the underlying grammatical system, we have to be aware of — and control for the influence of — various extragrammatical factors. Making an reasoned and informed decision about what is indeed in the grammar consumes a large portion of the fieldworker’s thoughts. The language consultant cannot tell you if a sentence is grammatical, she can only tell you if it is “okay” (is acceptable, if you can say it, and other such expressions). The field worker has to figure out if the rejection (or acceptance) reflects grammar or culture.

How does a speaker’s attitude play a role?

We want to get at an individual’s internalised grammar, but this is a concept understood only by trained linguists (bitter experience reveals that trained linguists don’t all get it). Many people think of language as a “social object” which is somehow “out there” i.e. The Language as possessed by All People speaking The Language. Alternatively, it may be thought of as the language as spoken by the King (normal people speak a cheap knockoff of the Hero’s Tongue). Consultants may thus think in terms of whether *someone* might say X in the language (and this may mean a high-class somebody, or an educated somebody, or — quite typically — an elderly somebody). Related to this is the speaker who thinks in terms of whether others will understand you if you say

³ People can often exploit properties of grammars for various purposes. I happen to know that many people associate the word “ain’t” with ignorance and “calumny” with high-falutinness (huh... would you spell that “falutingness”?), so I may decide to use or avoid these words, depending on whether my goal was to irritate someone or pretend to be educated, or whatever. These kinds of “register” associations are perfectly legitimate objects of scientific inquiry, but they are about something bigger than “what’s in the grammar”, they’re about “how you use it”.

X. A consultant thinking of language in such terms is likely to have more liberal but also less consistent intuitions. The linguist can help out considerably simply by asking questions like “how you *you* say it”, or “which way do you think *you* would do this”, if it seems that the consultant is being unduly influenced by a priori judgements about how they think other people say things. The ability of a speaker to consider themselves apart from all other humans is a variable personality trait, and contributes significantly to consistency of judgements. But consider the empirical consequences of two speakers of the same age, social group and from the same town, with different attitudes. It happens that there are lots of variants of the language in the town, some having to do with age groups, others having to do with where you immigrated from, and then there are caste-related differences. The speaker who answers your questions with attention to the question “how would *I* say this” will probably result in a more uniform corpus; the speaker who thinks in terms of what *some* speaker could say, will most likely be less consistent because at one moment they might think “How would I say this”, at another moment they might think “How would the King say this”, and in general “Is it conceivable that anyone would ever say this”. Not consciously, of course.

Variation does not come labeled as to cause, so unfortunately you can't be sure whether you're getting 'noise' or 'data'. Inconsistent, self-contradictory variation is noise. If a speaker say X one day and Y another, that does not mean they are inconsistent, if both X and Y are possible. The contradiction comes in claiming that X is possible and Y impossible on one day, but deciding that Y is possible and X impossible the next day (or, at any rate, later). Variation that is relatively consistent and reproduceable is 'data', proportionate to its degree of consistency — although, it may not be a grammatical datum, so for example you might get a consistent but lunatic judgement that people who pronounce the word as [zya] are lacking in proper moral upbringing. It may be useful to try to understand the nature of such variation, and if you find a repeated pattern of variation, you may want to gently probe the nature of the variability. The message here is that you don't want to be obsessive about controlling the variation. If you find some point of variability, start by simply asking whether the speaker can think of any difference in how you use the two. At this point we'll just leave it at that: see later for more discussion of what questions to ask and how to ask them. Right now I want to focus on the answer. If you get an answer like “it depends on what you are emphasizing”, you probably have hit on a structural property; if you get an answer like “some people say...” then you're probably getting a social intuition. Speakers can often give clues which clarify the nature of the variation.

Real consultant problems

Now it is time to take up the issue of actual consultant problems. They have been rare in my personal experience, but in doing field work that can arise (though it is highly unlikely that you will have to deal with this right now, speaking generically. so this is more a “lesson for the future”).

Language competence is an issue. When we do field work, we are presumably dealing with a fluent native speaker of the language; but what exactly is one, and how do you know if you have one? A really obsessive definition would be “someone who is born into an X-speaking society, grows up all his/her life speaking X and nothing but X, and still uses X as the major or sole language of communication”. I'm not sure how many “native speakers” there are, under that definition, and certainly in the classroom context which we're dealing with, there are none. A speaker of Bukusu does not cease to be a speaker of Bukusu when they move to Nairobi and use English or Swahili as their main languages of communication. Non-use of the language can have a negative effect, but it does not preclude the possibility of using such a speaker as a consultant. It *is* possible to lose ability in a language via disuse (usually measured in multiple decades especially if the lan-

guage is spoken by a persecuted minority who has reason to wish to distance themselves from their heritage, as illustrated by the multi-millenia-long struggle between the *Sapmelaš* and the *Dáru* especially in the eastern kingdom): see the paper on the “startup effect” (t.b.w.) for discussion of deciding whether a consultant has effectively lost the language.

A second problem regarding competence is the “turns out to not actually be a speaker” problem. In multilingual settings people can have partial knowledge of a language that allows them to communicate in a language, without being actual speakers of the language (that’s the way it is with me and lots of languages). A person may be identified as a speaker of Efik, when in fact they speak Ibibio, and only have a partial knowledge of Efik (Efik being the ‘socially dominant’ language). If this happens, your data on tone and vowel length is probably extremely suspect (which is why we switched to Ibibio). Or, one parent speaks Gogo and the other speaks Hehe, and they grew up in a Hehe speaking area so they learn Hehe, but they visit relatives in Gogo-land and thus learn a certain amount of Gogo, but not enough to be considered fluent speakers of the language.

Nonfluent speakers tend to be extremely inconsistent, struggle over the simplest of things and show no signs of improvement, and have high rates of self-admitted errors. Nonfluency is not a ‘typical’ problem in the controlled classroom setting, but it can be a real issue when you are in the field, especially if you’re in ‘the big city’ where the language is not indigenously spoken and there is money involved.

If the person is not actually fluent in the language, they will usually tell you, either directly or indirectly; for instance they will say things like “I only used the language when visiting my maternal grandfather who lives at the other end of the country”. They may simply say “I don’t speak the language very well”; you then have to decide whether that is just bashfulness, or really means they lack fluency. Saying something like “I don’t speak very well” may mean that they are not elegant orators, or they don’t know the classical word for smoked goat-throat-flesh-lappet, like their grandparents do. Try to frame the question in terms of whether they can talk to other speakers of the language and be understood, and can they understand normal conversations. If they express reservations about everyday chitchat, maybe they aren’t fluent speakers. Barring organic disorders, everybody speaks something fluently, so if the speaker implies non-fluency in their best language, they are being modest. It is difficult to know when these problems rise to the level rise to the level that one should give up working with the particular consultant. Fraudulent pseudo-speakers are a possible but pretty rare problem, and can be detected after a fairly short period of elicitation (often with the help of another speaker). Fake speakers can be detected because they are extraordinarily inconsistent.

Assume then that you are working with a speaker who actually has native speaker competence in the language being investigated. The aforementioned issues about ability to give data in the target language can still have an influence even when the speaker is a good speaker of the language. Influence can show up from another language, perhaps the language of elicitation, or the national language, or some other local language which the speaker happens to be bilingual in; or worse (because it is so hard to identify and root out) they may speak another dialect, especially if the speaker grew up in two different areas of the X-speaking territory. If the speaker vacillates between *olugoye* and *omuguha* for “rope”, it might simply be due to bleed-through from another language/dialect that they happen to speak. You have to just accept that this happens. Because your goal is to describe what you get, it doesn’t matter if what you get is “pure” from a historical or social sense. It only matters whether there is a (reasonably) consistent pattern. Consistent influence from outside the grammar of the target language defines one of the fundamental limits on field work (not to mention the theory of grammar vs. performance), because for one speaker, bilingualism may be so

profound that they don't know that they are mixing outputs from two grammars, but for another speaker it may reflect multiple patterns coexisting within a single grammar. Sometimes, you have to say "I don't know for sure".

What is a reasonable expectation?

The biggest struggle for the budding fieldworker is developing an appreciation for what is a reasonable expectation to have of a language consultant. It may be reasonable to expect that a speaker of Malayalam with a PhD in linguistics can explain how passives work in their language, or can list the phonemes and their major allophones, but that is only true thanks to specialised training.

Do not expect your speaker to perform a linguistic analysis of their language. That is what you are supposed to be doing. There are no shortcuts. Absolutely do *not* try to coerce your speaker into doing a theoretical analysis (e.g. do not teach the speaker about consonantal extrametricality and ask "is that consonant extrametrical?"). Maybe after many years working with a speaker you can train her to do linguistic analysis (maybe even convince her to go to grad school in linguistics), but we won't mess with that possibility here because that is in the remote future tense. Don't do it in this class.

The pleasant surprise is that quite often speakers can, nevertheless, tell you important things about their language which are quite correct and quite enlightening, so they may be able to say "this vowel is longer than that one" or "you could say that if you leave out the *e*-", or "no, you'd have to use the other past tense". Unfortunately, speakers can also say things that are totally wrong with the same kind of conviction. Therefore you should take any kind of analytic information coming from a speaker as a suggestion well worth considering, but it is just a suggestion. Ultimately, you have to assemble the evidence which forms the justification for the claim (even if the consultant gives you the breakthrough hint). To clarify, this is about volunteered analysis, not data. If a speaker says 'X and Y sound different', that is a fact which needs to be dealt with (if it is consistent). If a speaker says '[mai] is one syllable and [moe] is two syllables', that is an analysis which goes beyond unbiased and objective statements about the language which any speaker of a language can be expected to have: it depends on a theoretical construct which is not part of "naive knowledge", and is acquired only by direct education. It may be correct, but the correctness would be established on the basis of raw data, such as minimal pairs where the speaker can say "X and Y sound different", or some theoretically-grounded argument that you construct based on how verb conjugation works. Or, it might be supported by something as simple as the "say it slow" test which reveals a big difference in how you stretch out [mai] versus [moe].

Saying words

The most basic task that should be doable by any speaker of a language is saying words. Don't confuse this with performance on the task "say the word for 'rat'". In order to "say the word for 'rat'", the speaker has to understand what a 'rat' is, access the word in their mental lexicon, and pronounce it. In a case-marking language, an additional decision may have to be made as to what case-form to use — the nominative, the vocative, or what? Just because you yourself may have a clear mental image of what a 'rat' is does not mean that everyone can translate from English to their language. People can get confused about the technical distinction between 'rat' versus 'mouse', not to mention 'mole' and 'rabbit'; and you yourself might be confused. If you explain "you know, rat, the one with the big fluffy ears", you are exacerbating the problem. Field workers often have bags of preserved rats, mice and voles for such demonstration purposes. Pictures are a poor substitute: real rats are not 1 inch long and 2-dimensional.

Imperfect fluency in the language used for elicitation can cause problems. You have to consider the possibility that *you* don't speak the elicitation medium fluently (causing problems), or that the consultant doesn't. And worse, your dialect may be different from the informant's, especially in the case of English. It is unlikely that any informant can give you the word for "machete" pronounced [məʃeri], although depending on where they come from they might pronounce it [mæçət] or else call the object a "cutlass" ("machete" is, as far as I know, a strictly American English word, though an informant who has lived in America for a while could have learned the word). Whatever semantic distinction (if any) you make between "I have sung" and "I sang" may have a totally different set of connotations in the dialect of your informant. Don't waste time insisting that "I have sung" means X and "I sang" means Y, if that's not what it means to the informant; although, if the informant is curious about this word "machete", feel free to explain it. The point is, words (what the word is and its exact meaning) can differ considerably between dialects, so don't assume that your understanding of a word is the only one possible (and certainly don't assume that your understanding of the word is morally superior).

The problem may not be a question of comprehension, but one of intrinsically vague definitions or lack of well-known common words. It is in the area of "technical vocabulary" that one is likely to have the greatest chance of running into a translation problem. When an informant translates the word [gabagul] as "warbler", they may be correct, but trust that belief only if the informant is a graduate student in ornithology, writing a dissertation on warblers. Very often, technical terms (such as plant and animal identifications) are thanks to some teacher who may have been wrong (this explains why one consultant was taught that the long-nosed pig-like animal is an "anteater", when it is actually an "aardvark" — a term not used in many dialects of English except those where "A" is for "aardvark"). Similarly, there is a term "calabash" which in a strict historical sense refers to a dried gourd typically used to hold fluids, but is sometimes extended to mean "any container of the appropriate shape and size, used for liquids" which allows for the existence of "bamboo calabash", "clay calabash". A few people have the technical knowledge that allows them to distinguish cedar, pine, hemlock, fir, larch and spruce, but most people just refer to those things as 'pine trees'. They may know these words, and just not know exactly what kind of tree they are (it was late in life that I learned what kind of tree a larch is, even though it was a word I knew for years before).

Furthermore, there may not be a distinction between 'rat' and 'mouse' in the language. There may also be no single word for 'rat' and instead there may be words for 'grey cane rat', 'house rat', 'spotted forest rat' and so on. It is the linguist's job to think of such factors (which might explain why the speaker pronounces the word *agapako* one time and *omipi* the second, because she was thinking of different kinds of rats). It's pretty unlikely that a speaker will say "You can pronounce the word as *agapako* or as *omipi*". More likely they will tell you that there are two words, one pronounced *agapako* and the other *omipi*; they might tell you when you use one versus the other. At any rate, you ought to be able to get a speaker to utter a word, even if you can't be positive what it means and what the context is for appropriate useage.

Judging words

Pronouncing words is really the only thing you can take as a given (assuming you are asking for a word that the speaker knows: if they don't actually know the word, you can't expect a lot of consistency in terms of how it's pronounced). A close runner-up would be distinguishing possible and impossible words. For example, you might ask a speaker if "possible+im" could be a word of English, and it's pretty likely that they'd say "no", but there might be agnostics out there who can't tap into their intuitions and put you out of your misery, so they might say "I don't know". Why

would anyone ever say that? Well, it's not very different from asking if unicorns, yetis, 5-legged insects (I don't mean the kind that have undergone a 6-to-5 conversion job at the hands of a small child) or blue-skinned mongooses exist. If you have experience with a thing, you know it exists. If you don't have experience with it, you don't know that it *doesn't* exist. In that case, people appeal to all sorts of concepts to answer the question; they may decide that "I don't know" is the most that they can responsibly say. Fortunately, most people are able to avoid this trap whereby they can't ever say "no". I think everyone probably *can* do it, unless they get too philosophical and start worrying about negative-existential proofs.

The main problem area regarding judgements about words is unfamiliarity with the word. It's very unlikely that a real speaker of a language would not know a word for "nose" (not knowing is one thing, not remembering instantly is another), but they may not know some particular version of the word like [šnaz] as opposed to the more commonly used [bizər], and therefore may not be able to tell you the accusative for it. This generally affects specific roots (including anything that might be construed as "technical vocabulary", so that ironworkers often don't know the terms for fishing implements), and the uncertainty may be general or restricted to some one area where the answer is not trivially figureable out. This is probably the most common form of inability to decide about words.

Speakers can have "fuzzy" intuitions about words which don't have to do with specific roots; thus they may not be able to give you a good answer to any word of the form "plural subjects didn't cause each other to Verb for me", no matter what the root is, or "2nd person dual possessor form of illative plural" for nouns in general. This kind of fuzziness correlates with frequency multiplied by transparency, which is to say, if a given form is rarely used, and if the possibility of the combination, or the shape of a morpheme, is not totally trivial, then you have an increased chance of getting unclear judgements. Inflectional properties tend to be more resistant to fuzzy judgements because most words have the relevant marks and you use them all the time. Chaos breeds chaos: if there is a high degree of lexicality in a given area of a language, a speaker is likely to have a sub-conscious (or conscious) awareness of this low predictability. Then in some closely related area of the language which has low frequency of use, a feeling of uncertainty is likely to also exist.

Grammaticality

As long as you avoid the term "grammatical", speakers can be expected to judge whether a given sentence is "okay" or not, that is, "grammatical". The best way to ask about sentences is "Can you say X?"; avoid the term "grammatical" because most people think that means "in conformity with the official rules in the grammar books" which may be seen as a test of the speaker's knowledge of formal grammar. The problem is sorting the reactions and turning them into something like a yes/no distinction. This is a pretty huge topic (to be covered separately), but the basic fact is that people can say things like "yes", "no", "maybe", "not as good", "even better" and so on. The speed of the reaction and the forcefulness of the reaction are good indicators: if a speaker starts snorting in derisive laughter before you finish saying the sentence, you probably have an ungrammatical sentence. Expressions like "I suppose you might..." or "Some people could..." are indicators of marginality.

Strangely, I'm just going to leave this topic right now, because it needs separate and vastly more extensive treatment. The main point is, you should expect (apparently) uncontrolled variation in reactions, depending on a lot of mystical-seeming things, like what sentence you asked 5 sentences ago. Converting "reactions" into grammaticality values is the proverbial smoking gun.

Meanings

Another thing that you can almost expect from a speaker is a meaning or use of the word. You can't automatically expect speakers to produce clear and unambiguous translations of words or sentences into English (even if they are fluent speakers of English). The speaker may believe that the animal known scientifically as a "gennet" is a "weasel", because gennets look like weasels and everybody knows the word "weasel". The probability of a good translation is related to the extent to which concepts correlate in the two languages using a very few words (one or two). If a particular verb tense really translates as "completed action" as opposed to "past tense", then you may have a lot of trouble getting a translation into English, because the former is not a concept that has a clear and systematic role in English inflectional morphology. You can eventually express any concept in any language, and therefore you can eventually translate any sentence between the two languages, even if it involves a lot of paraphrasing and explaining.

So if a speaker can't easily tell you what a sentence means, it is most likely that this is a problem regarding translation. It is very unlikely that a speaker would say "Yes, you can say [aba gabo plini morp] but I have no idea what it means". If they say this, they almost certainly mean "I don't know how to explain it in English, but I certainly comprehend the sentence in my own language". (I'm assuming that the sentence is unambiguous: if it turns out that the sentence has two meanings, that could be what the speaker's reaction means). You can probably arrive at the meaning by starting with the main semantic properties, then refining the concepts. There should be no real problem figuring out what the action is, who did it, and who it was done to; though, if the speaker is distracted by trying to figure out this puzzle "what *does* that mean?" they may not be paying attention and may not answer your question "did the *child* do the cooking, or did the *badger*?". The concepts that go into making up the meaning of a sentence can be sorted into a hierarchy of easier and harder to access information. Easiest to get is the basic nature of the action (planting, cutting, falling) and the core thematic roles (the person who does it; the thing that gets chopped; the thing you chop with). Assuming that the language makes systematic time reference distinctions — not all do — you should be able to tell if the sentence says that the action has happened or will happen (the concept of "is happening" is surprisingly slippery in some languages).

It may take more work to clarify subtler concepts; if there are past tenses for events of today versus events of yesterday, the speaker may not be able to verbalise that distinction, so you may be left with two sentences with different verb forms and no clue why they are different. Once you've gotten a dozen or so examples, and assuming that the true nature of the system hasn't made itself clear spontaneously, begin by comparing two sentences that differ minimally in the form of the verb. Ask "do you use these two the same way, or differently?". Maybe they can explain it, which is maximally cool. Maybe they can only say "they're different, but I can't explain it"; that's okay. Maybe they can say "yes, the first means X and the second means Y", only after you gather about a dozen examples it becomes apparently that the speaker is just making up stuff at random; that's a nuisance, because you have to decide whether to filter out that information entirely, or pay some attention. And maybe they'll say "they mean the same thing" when they don't mean the same thing, but because they can't explain the difference they assume that they mean the same thing (since if they really meant different things, surely they could *explain* the difference, and therefore there can't be a difference... unless it turns out that the ability to explain facts of meaning is not the same as implicitly knowing those facts, a subtlety known by most linguists but not by most regular people).

Asking whether two sentences mean the same thing is sort of like asking about underlying forms or syllable-parsings: it requires a theoretical analysis. The concept of "meaning" is a sufficiently general one in the world that it is okay to ask about meaning. In order to decide about iden-

tity of meaning, you have to know what a meaning is. If you were to ask the question “Do sentences A and B have exactly the same truth-conditional semantic properties, referents, and pragmatic properties” you would be asking a more precise question, but one which a normal speaker would not understand. Speakers generally have a concept of “meaning” that attaches higher importance to thematic roles and “essential” properties. Thus the sentence “The boy bit the dog” is unlikely to be seen as meaning the same as the sentence “The dog bit the boy”; “The boy bought a banana” is unlikely to be seen as meaning the same as “The boy bought a knife”. It is likely that if yesterday past and today past is signalled by verb tense, given the sentences “The boy saw_{today} the dog” and “The boy saw_{yesterday} the dog”, the yesterday vs. today distinction will not be initially seen as important enough to constitute a “difference in meaning” (I’ve messed with this in dozens of languages, and only a couple of exceptional speakers were attuned to the time difference from the beginning): people tend to think of such tense properties as not being about meaning, especially when the medium of elicitation does not systematically encode the distinction in the morphology.

Pragmatic properties are the worst, especially if there is some marker or verb tense distinction whose function is to mark a constituent as having contrastive focus. Informants generally have little idea how to convey such information, and (alas) the linguist may have little idea that that is what the speaker is getting at. So-called validation markers, which indicate degree of speaker certainty and personal experience, are also very difficult to translate. Thus if a speaker says of some example either “it depends on what you want to say”, “this is more emphatic”, “you use this to reassure someone” or “you use that only if the other hoe is broken”, this is “unreliability”, in the sense that you have uncontrolled variation (in terms of the meaning). What the linguist needs to do is figure out the unifying property behind the examples, then come up with a reliable way to associate form and meaning (for focus, it could be something as simple as asking “How do you say ‘He dropped the *knife*, not the *fork*?’”). In fact, linguists themselves are not univocal as to whether pragmatics is part of semantics or is separate; so you can’t expect average speakers to have a clear distinction between truth conditional semantics which is “meaning” and properties regarding use and attitude, which is pragmatics (so I’ve been told).