Musical Relevance: a Computational Approach

Edoardo Acotto (acotto@di.unito.it)
Università di Torino - Dipartimento di Informatica, Corso Svizzera 185
10149, Torino ITALY

Daniele P. Radicioni (radicion@di.unito.it)
Università di Torino - Dipartimento di Informatica, Corso Svizzera 185
10149, Torino ITALY

Abstract

This study is a first attempt at formalizing the concept of *Musical Relevance* from a cognitive and computational perspective. We elaborate on Sperber and Wilson's Relevance Theory, and extend it to account for musical cognition, involving both listening and understanding. Our claim is that the application of the concept of *Cognitive Relevance* to music would permit us to partially explain hearers' behavior and composers' choices. A computational model of Musical Relevance could also contribute to the formulation of a general computational theory of musical cognition. In turn, formulating an algorithm to compute Musical Relevance can shed light on the computational nature of the broader cognitive principle of relevance. We propose to unify Relevance Theory with the Generative Theory of Tonal Music, in order to compute Musical Relevance. We started implementing a system to test the proposed approach over simple examples and report about the results in a preliminary experimentation.

Keywords: Relevance Theory; Computational Approach; A Generative Theory of Tonal Music; Tonal Pitch Space.

Introduction

The *Relevance Theory* was initially formulated as cognitive-pragmatic theory of communication (Sperber & Wilson, 1986); lately it has been viewed and developed as a general theory of human cognition (Wilson & Sperber, 2004; Carruthers, 2006). The relevance of an input to an individual (or a cognitive system) is defined as the ratio between the cognitive effect and the processing effort. In the authors words:

"*a.*) Other things being equal, the greater the positive cognitive effects achieved by processing an input, the greater the relevance of the input to the individual at that time. *b.*) Other things being equal, the greater the processing effort expended, the lower the relevance of the input to the individual at that time."

An input is relevant for a cognitive agent in a given context when it can be related with the information registered in memory and accessible, and when this relation yields a "positive cognitive effect".[1] Relevance of an input is a continuous (non-categorical) variable. The concept is comparative and non quantitative: e.g., "*x* is more relevant than *y*, for *P* in the context *C*".[2] The greater the cognitive effects are, the greater

the relevance of a given input is (*ceteris paribus*); on the other side, the smaller is the processing effort, the greater is the relevance of a given input (*ceteris paribus*).

It is a matter of fact that Relevance Theory has a lot of opponents. For example, in a footnote Jerry Fodor let us know that according to his opinion a Relevance Theory doesn't even exist: "As for a theory of relevance, saying that if we had one it would solve the frame problem is as pointless as saying that if we solved the frame problem, that would give us a theory of relevance: Both are true, of course, because 'assessing relevance' and 'framing' are two terms for the same thing. [. . .] If cognition is to attain true beliefs with any efficiency, it's got to be the case both that what's importantly relevant is generally in the frame, and that what's not importantly relevant generally isn't. Maybe meeting these conditions is tractable within the assumptions of Classical theories, but I don't know of any current proposal for a cognitive architecture, Classical or otherwise, that seems likely to tract it" (Fodor, 2000, p. 114). Fodor takes correctly Sperber and Wilson's theory as a semantic-pragmatic theory of linguistic comprehension, and he poses the question of how a cognitive system can attain true beliefs in an efficient way. However, if we conceive that Relevance Theory can be a general theory of (human) cognition, we have to remark that not all mental representations have a truth-based semantics, and mental representations of music seem to be a good candidate for representation without truth value (Acotto, 2011 (in press)). So, the efficiency of the cognitive system faced with non-semantic representations has to be analyzed with different criteria than those that Fodor has in mind.

One chief problem with Relevance Theory is the difficulty to formalize it: however, in a restricted and formal domain like music, this seems to be possible and psychologically plausible. Modeling musical relevance we have to shift from a "subjective" concept of relevance to an "objective" one: instead of modeling the musical relevance for a given individual in a given context, we'll model the relevance for an idealized listener familiar with the Western tonal music idiom. That is, we are presently concerned with a restricted subset of all possible music.

In their original formulation of Relevance Theory Sperber and Wilson (1986) do not propose any method for calculating relevance, so we had to provide relevance with a quantitative counterpart to design a computational model. This is a key contribution of the present work.

---

[1] "A positive cognitive effect is a worthwhile difference to the individual's representation of the world: a true conclusion, for example. False conclusions are not worth having. They are cognitive effects, but not positive ones" (Wilson & Sperber, 2004).

[2] On the comparative/quantitative notion of relevance, and on Carnap's distinction of comparative and quantitative concepts, see (Sperber & Wilson, 1986).

The formulation of an algorithm to compute Musical Relevance would lead to improve the general computational nature of the cognitive principle of relevance: "Human cognition tends to be geared to the maximization of relevance" (Wilson & Sperber, 2004). If Relevance Theory is empirically plausible, and if the musical mind yields a kind of thought comparable with other forms of mental life, Relevance Theory can apply to the musical thinking as well. In order to explore such hypothesis, we propose to put together Relevance Theory and Lerdahl and Jackendoff's Generative Theory of Tonal Music, GTTM (Lerdahl & Jackendoff, 1983).[3]

The GTTM describes the musical comprehension of a hearer familiar with the Western tonal idiom. It postulates the existence of mental representations of music, structured on five levels: first the musical surface, than two horizontal structures (meter and grouping), and finally two hierarchical structures, the time-span reduction and the prolongational reduction, which can be formalized as binary branching trees (Hamanaka, Hirata, & Tojo, 2006). Generative Theory of Tonal Music finds in Lerdahl's Tonal Pitch Space theory a partial readjustment (Lerdahl, 2001), especially concerning the formalization and the quantification of the musical dimensions.

Although other musicological theories exist that are related to musical salience (e.g., by Deliège (1996)), we chose the notion of relevance by Sperber and Wilson because it seemed to be more naturally suited to a computational implementation. A major assumption is that it allows formalizing and quantifying musical relevance via the computation of the musical cognitive effect and of the processing effort. We presently do not explore the connections of our work with related investigations of a notion similar to Musical Relevance and grounded on information theory (Conklin & Witten, 1995; Pearce, Conklin, & Wiggins, 2004): we defer to future work the exploration on such links.

The paper is structured as follows: we first qualify Musical Relevance as the ratio between the cognitive effect and the processing effort, and explore both cognitive (musical) effect and (musical) processing efforts. We then provide an example to show how such concepts fit to the musical context. Then a preliminary experimentation is illustrated, and the results are reported and discussed. Finally we conclude by pointing out present limitations and future works.

Computing Musical Effect

According to Relevance Theory, in order to be more relevant than another, a music excerpt has to offer a greater cognitive/emotional effect than another one requiring the same pro-

cessing effort; alternatively, a musical excerpt has to require a minor processing effort than another one that yields the same effect. The Musical Relevance (MR) is defined as the ratio between the Musical Effect (ME) and Processing Effort (PE): that is, $MR = ME/PE$.

GTTM individuates three types of tonal tension: *surface*, *sequential* and *hierarchical* tension. Some experimental tests have been carried out, confirming that sequential tension is not sufficient to represent the effective musical understanding, and that hearers perceive hierarchical tension as well (Lerdahl & Krumhansl, 2007). The Musical Effect yielded by the tonal tension is complemented by the tonal attraction:[4] in other words, the "forces" that constitute musical effect are both *tensional* and *attractive*. In order to calculate the musical effect some rules can be applied, that were devised as Tonal Pitch Space Rules (Lerdahl, 2001). The following rules can be used to compute the musical effect.

Surface tension rule

$$T_{diss}(y) = scale\_degree + inversion + non\_harmonic\_tones \tag{1}$$

where the tension score of the target chord y is computed as the sum of three elements *scale_degree*, *inversion* and *non_harmonic_tones*. *scale_degree* is 1 if $3^{\wedge}$ or $5^{\wedge}$ is present in the melodic voice, 0 otherwise; *inversion* is 2 if $3^{\wedge}$ or $5^{\wedge}$ in the bass, 0 otherwise; *non-harmonic tone* is 3 if a pitch class is a diatonic non-chord tone, 4 if it is a chromatic non-chord tone, 0 otherwise.

Sequential tension rule

$$T_{seq}(y) = \delta(x_{prec} \rightarrow y) + T_{diss}(y) \tag{2}$$

where y is the target chord, x_{prec} is the chord that immediately precedes y in the sequence, $T_{seq}(y)$ is the tension associated with y, and $\delta(x_{prec} \rightarrow y)$ is the distance from x_{prec} to y.

Hierarchical tension rule

$$\begin{aligned} T_{loc}(y) &= \delta(x_{dom} \rightarrow y) + T_{diss}(y); \\ T_{glob}(y) &= T_{loc}(y) + T_{inh}(x_{dom}) \end{aligned} \tag{3}$$

where y is the target chord, x_{dom} is the chord that directly dominates the prolongational tree; $T_{loc}(y)$ is the local tension associated to y; $\delta(x_{dom} \rightarrow y)$ is the distance from x_{dom} to y; $T_{glob}(y)$ is the global tension associated to y; $T_{inh}(x_{dom})$ is the sum of the values of the distances inherited by the chords that dominate x_{dom}.

Melodic attraction rule

$$\alpha(p_1 \rightarrow p_2) = \frac{as_2}{as_1} \cdot \frac{1}{n^2} \tag{4}$$

where p_1 and p_2 are pitches, with $p_1 \neq p_2$; $\alpha(p_1 \rightarrow p_2)$ is the attraction of p_1 to p_2; as_1 is the anchoring strength of p_1 and

---

[3]It is noteworthy that Musical Relevance model is not directly committed to the GTTM for computing the effect and the effort. E.g., we could employ different theories descending from (Meyer, 1956), such as (Narmour, 1990, 1992) and (Huron, 2006). However, Narmour's Implication/Realization theory does not account for the *hierarchical* structure (i.e., binary branching tree) of music perception (Margulis, 2005, p. 688), and the same holds for Huron's Expectation theory.

[4]The model by Lerdahl and Krumhansl is a quantitative theory of tonal tension made out of four components: "1. A representation of hierarchical (prolongational) event structure. 2. A model of tonal pitch space and all distances within it. 3. A treatment of surface (largely psychoacoustic) dissonance. 4. A model of voice-leading (melodic) attractions" (Lerdahl & Krumhansl, 2007).

as_2 is the anchoring strength of p_2 in the current configuration of the basic space; n is the number of semitone intervals between p_1 and p_2.[5]

Harmonic attraction rule

$$\alpha_{rh}(C_1 \to C_2) = c \cdot \frac{\alpha_{rvl}(C_1 \to C_2)}{\delta(C_1 \to C_2)} \qquad (5)$$

where $\alpha_{rh}(C_1 \to C_2)$ is the harmonic attraction of C_1 toward C_2; the constant $c = 10$; $\alpha_{rh}(C_1 \to C_2)$ is the sum of the attraction of the leading voices for all the voices in C_1; $\delta(C_1 \to C_2)$ is the distance of C_1 a C_2, with $C_1 \neq C_2$.

Such rules have found an experimental corroboration in (Lerdahl & Krumhansl, 2007). We assume that these rules represent a good approximation of the musical effect in the overall computation of musical relevance.

For sake of simplicity and because of the greater complexity of music, in this paper we are concerned with melodic music only. Even though this is a clear simplification, and it is the first step of a more complete and complex model, our present work allows us to make experimental tests and to compute a musical relevance score for simple melodies.

Computing Processing Effort

Concerning the *PE*, no methods to calculate it are given nor suggested in (Lerdahl & Jackendoff, 1983) nor in (Lerdahl, 2001). Nevertheless, following GTTM we can surely identify a vertical, hierarchical, dimension of the *PE* represented by the binary branching trees of the musical structure. Against the "concatenationism" hypothesis (Davies, 2011), musical surface is not enough to understand music, and the structural properties of a melody are a key element for its understanding. We can then assume that at least a great portion of the *PE* is involved in detecting the structural properties of the heard melody.


Figure 1: Structure of the toy melody represented in the GTTM notation, by Katz and Pesetsky (2009).

How to compute the binary branching trees that represent the hierarchical musical structure? In their reinterpretation

of the GTTM, Katz and Pesetsky observe that in the GTTM time-span trees the more relevant information is the hierarchical distance from the root of the tree (see Figure 1). This distance is measured through a Root Distance (*RD*) number: "The *RD* number of an event *e* in a structure *K*, $RD(e)$, is the number of nodes that nonreflexively dominate the maximal projection of E (i.e. *eP*) in *K*" (Katz & Pesetsky, 2009).


Figure 2: Structure of the toy melody in the standard linguistic notation.

The reinterpretation of Lerdahl and Jackendoff's theory is made within the framework of generative linguistics, where the concept of "projection" plays a central role: "A constituent whose head is *H* is called a projection of *H*, and is conventionally labeled H' ('H-bar') if it is dominated by another projection of *H*; and *HP* otherwise. *HP* is called the maximal projection and H' is called an intermediate projection of *H*. *H* itself is sometimes called the zero-level projection of *H*". As it is showed in Figure 2, the linguistics notation allows a graphic translation of the structure of the GTTM.[6]

We will take into consideration the *RD* number of each sound event as a part of the *PE* (we increase by 1 the *RD* numbers). So, we shall measure the *PE* by using the rules of time-span reduction formulated in GTTM.

An example with melody

Searching for a first implementation of our model, we focus initially on the case of the melodies, in particular the leading voice of the toy melody, and illustrated in Figure 3.

Since we are considering a melody, we only make use of the *melodic attraction rule* (please refer to Equation 4), since the other rules are concerned with events where multiple notes are present at a time.

Music effect Since melodic attraction is between each two musical events, an attraction number is not referred to a single

---

[5]In our implementation we adopted a correction to handle the case of repeated notes, with $p_1 = p_2$, that otherwise would produce a division by zero. Since the underlying rationale is that in case of repeated notes the ratio $\frac{1}{n^2}$ should approach 1, we presently set the value of n to 0.9.

[6]"Variations in the notation with which one expresses a theory can influence one's thinking about the actual topics under investigation. Even when different sorts of diagrams represent exactly the same information (as is the case here), the differences among them may reflect and reinforce differing working hypotheses or hunches about the kinds of phenomena one expects to model. Differences of this sort between GTTM and common practice in linguistics arise in two important domains: the relevance of projection level and the amount of information that project from terminal nodes to the constituents that they head." (Katz & Pesetsky, 2009)

Figure 3: The leading voice of the toy melody.

Figure 4: The three levels of the toy melody time-span reductions.

pitch, but represents a transition between two adjacent pitches *x* and *y*. The attraction values (i.e. the musical effect) between each two notes in the toy melody are as follows:

$$
\begin{aligned}
\alpha(C \rightarrow D) &= 0.125 \\
\alpha(D \rightarrow E) &= 0.375 \\
\alpha(E \rightarrow F) &= 0.667 \\
\alpha(F \rightarrow G) &= 0.375 \\
\alpha(G \rightarrow A) &= 0.007 \\
\alpha(A \rightarrow B) &= 0.25 \\
\alpha(B \rightarrow C) &= 2.0 \\
\alpha(C \rightarrow D) &= 0.125 \\
\alpha(D \rightarrow C) &= 0.5
\end{aligned}
$$

Processing Effort In order to calculate the *PE*, we compute the *RD* numbers by following the binary branching tree: for each pair of events the *PE* will be the average of the two *RD* numbers (augmented by a unit). To compute processing effort implies the possiblity to automatically "reduce" a given melody to its more fundamental schema, as it is shown in Figure 4.

This kind of reduction relies on a set of "preference rules". However, these are not easily implemented because if there are cases where multiple rules are triggered at the same time, unfortunately in the GTTM no criteria are proposed to resolve such conflicts. Deepening the computation of the *PE* component of *MR* will be one major focus of future work.

Musical Relevance We can now calculate the musical relevance of the transition from one event to another in a melody (without considering, for the moment, the relevance of similarity). As each note has a value of PE, but the attraction value

Figure 5: Excerpt from the Piano Sonata in C minor (KV. 457) by W.A. Mozart.

Figure 6: Excerpt from the "Petit pièce" from the Opus 68 n.5 by R. Schumann.

is between two events, we'll calculate the average PE of each couple of notes, and then we'll calculate the relevance (the ratio between ME and PE) of the passage from one musical event to another.

Experimental Assessment

In order to provide the proposed approach with an experimental assessment, we devised the following experimental setting.

We implemented a system that computes the *Musical effect* (*ME*) component in the equation $MR = ME/PE$. We studied how the computed musical effect differs by varying two simple melodies.

A core hypothesis of the experimentation is that melodies by historical composers maximize *ME*, and *MR* as well.[7] Let us consider the *cognitive constraints* by Lerdahl (1988), that postulates that 'good' music is composed according to the cognitive nature of human mind/brain. By following this theoretical framework we stipulate that the original melodies from historical composers have higher *ME* than 'experimental' variations composed by ourselves.

In this setting, we expect the system to compute lower *ME* for our variations; also, we expect it to be able to distinguish between grammatical and 'ungrammatical' variations, by assigning lower scores to ungrammatical ones.

Experimental setting We selected the music excerpts illustrated in Figures 5 and 6. Such pieces were chosen in order to capture (and test the system in) two widely different experimental conditions. In particular, they can be thought of as two paradigmatic examples of themes opposite in spirit. The first one is rather percussive and jumps over the main degrees of the C minor key. On the other side, the second one is a typical *cantabile* theme: it is more regular under a rhythmic viewpoint, and the melody mostly moves stepwise.

---

[7] In accord with the given definition, music relevance (*MR*) grows –*ceteris paribus*– as the music effect (*ME*) grows, and vice versa it decreases –*ceteris paribus*– as the processing effort (*PE*) grows. Since we added new events by interleaving existing events with new ones, this makes the input more complex. Then we know in advance that new nodes will produce further levels to the reduction tree, thus increasing the *PE* component. That is, we know *a priori* that by increasing *PE* and by decreasing *ME*, the final *MR* will result reduced.

Table 1: The Musical Effect scores for the six considered pieces.

Excerpt	Score
Mozart excerpt	0.4259
Mozart Var 1	0.3888
Mozart Var 2	0.3267
Schumann excerpt	0.3182
Schumann Var 1	0.2460
Schumann Var 2	0.1877

Therefore, different musical aspects are accounted for by the considered excerpts.

We then elaborated two variations for each excerpt (Figure 7). In both cases the first variation (indicated as *Var 1* in Figure 7-*a)* and 7-*b)*) is only slightly different from the original excerpt. As regards as the second variation, the Mozart excerpt has been modified in a ungrammatical fashion (see *Var 2* in Figure 7-*a)*), whilst Schumann excerpt has been modified through the insertion of musically plausible notes that transform it into a rhythmically regular arpeggio (see *Var 2* in Figure 7-*b)*).

The implemented system takes in input the excerpts encoded as MIDI files, and computes the associated musical effect –through the formula in Equation (4)– as the sum of the melodic attraction between each two music events:

$$ME_{excerpt} = \left[\frac{\sum_{i=1}^{excerpt.length-1} \left(\frac{as_{i+1}}{as_i} \cdot \frac{1}{n^2} \right)}{excerpt.length} \right]$$

By adding new notes we expected a reduction in the musical effect. Furthermore, since the second variation of each input excerpt was more different from each 'original' source, we expected to observe a decrease in the musical effect and, relatedly, in musical relevance.

Results In accord with our intuition, the implemented system computed the maximum ME scores for the original excerpts, reduced scores for each first variation, and the lowest scores for both second variations. The final figures are reported in Table 1.[8] Provided that this experimentation represents only the very first step towards a psychological validation (that would require considering in how far the results approach human responses), the results seem to corroborate Lerdahl's hypothesis. Tonal music is governed by an attraction-based syntax. This sort of attraction, which is maximally exploited by composers and which is maximal in the original music excerpts, is at least partly grasped by the proposed model. Further work is needed to investigate whether and how classical western tradition as a whole 'incorporates' a criterion to maximize the Musical Effect (independently of the associated processing effort).

---

[8]The material employed in the experimentation (MIDI files, printable scores and Lilypond sources) along with the results file is available at the URL
http://www.di.unito.it/~radicion/datasets/cogsci12/ .

Also, if we compare the original excerpts (Figure 5 and 6), the system accounts for the greater 'dramatical' salience of Mozart's excerpt (which is a first theme of a C minor sonate). Schumann's excerpt is a simple piece: its value results perhaps from the balance between its simple effect and its structural simplicity.

Conclusions

The paper illustrates a modeling attempt, and an initial implementation of a complex phenomenon such as relevance-guided music understanding. The presented implementation only accounts for the musical effect; coping with the computation of the processing effort is left for future work.

Due to such limitation we considered for experimentation only melodies with differing surface, but with similar underlying structures. Notwithstanding this limitation, the preliminary experimentation provided some evidence that the musical effect captures meaningful aspects of Western tonal music.

Another relevant point for completing the Musical Relevance model involves dealing with musical similarity. Similarities and repetitions in music are a frequent and important structural phenomenon. They affect important musical features like style (Meyer, 1989), but they permit also to affirm that without similarity there would be no music, since similarity is a center of gravity for perception and comprehension (Cambouropoulos, 2009). Similarities and repetitions influence musical effect, and therefore they have impact on the cognitive relevance of a piece of music. The relevance of a musical event E_2 should be a function both of the relevance of the similar event E_1, and of the relevance of E_2 as taken in isolation. The similarity increases the relevance of a musical event; otherwise, similarities should be avoided for the risk to diminish the relevance effect. Starting from existent systems for computing musical similarity (Meredith, Lemström, & Wiggins, 2002; Radicioni & Botta, 2006), in future works we will focus on detecting similar patterns in music pieces.

Acknowledgments

References

Acotto, E. (2011 (in press)). Mental Representations of Music in Cognitive Science. In *Proceedings of the Italian Cognitive Science Society (AISC) Conference.*

Cambouropoulos, E. (2009). How similar is similar? *Musicae Scientiae*, *Discussion Forum 4B*, 7–24.

Carruthers, P. (2006). *The Architecture of the Mind*. Oxford: Oxford University Press.

Conklin, D., & Witten, I. H. (1995). Multiple Viewpoint Systems for Music Prediction. *Journal of New Music Research*, *24*(1), 51–73.

Figure 7: *a)* variations of the Mozart excerpt; *b)* variations of the Schumann excerpt.

Davies, S. (2011). *Musical Understandings: And Other Essays on the Philosophy of Music*. Oxford, UK: Oxford University Press.

Deliège, I. (1996, October). Cue Abstraction as a Component of Categorisation Processes in Music Listening. *Psychology of Music*, *24*(2), 131–156.

Fodor, J. (2000). *The Mind Doesn't Work That Way: The Scope and Limits of Computational Psychology*. Cambridge: MIT Press.

Hamanaka, M., Hirata, K., & Tojo, S. (2006). Implementing "A Generative Theory of Tonal Music''. *Journal of New Music Research*, *35*(4), 249–277.

Huron, D. (2006). *Sweet anticipation. music and the psychology of expectation*. Cambridge: MIT Press.

Katz, J., & Pesetsky, D. (2009). *The Identity Thesis for Language and Music.* draft available at http://ling.auf.net/lingBuzz/000959.

Lerdahl, F. (1988). Generative Processes in Music: The Psychology of Performance, Improvisation, and Composition. In J. Sloboda (Ed.), (pp. 231–59). Oxford University Press.

Lerdahl, F. (2001). *Tonal pitch space*. Oxford: Oxford University Press.

Lerdahl, F., & Jackendoff, R. (1983). *A generative theory of tonal music*. Cambridge: MIT Press.

Lerdahl, F., & Krumhansl, C. L. (2007). Modeling tonal tension. *Music Perception*, *24*, 329–366.

Margulis, E. (2005). A Model of Melodic Expectation. *Music Perception*, *22*(4), 663–714.

Meredith, D., Lemström, K., & Wiggins, G. (2002). Algorithm for discovering repeated patterns in multidimensional representations of polyphonic music. *Journal of New Music Research*, *31*(4), 321–345.

Meyer, L. (1956). *Emotion and meaning in music*. Chicago, USA: Chicago University Press.

Meyer, L. (1989). *Style and Music*. Chicago: Chicago University Press.

Narmour, E. (1990). *The analysis and cognition of basic melodic structures*. Chicago, USA: University of Chicago Press.

Narmour, E. (1992). *The analysis and cognition of melodic complexity*. Chicago, USA: Chicago University Press.

Pearce, M., Conklin, D., & Wiggins, G. A. (2004). Methods for Combining Statistical Models of Music. In *Computer Music Modeling and Retrieval: Second International Symposium, CMMR 2004* (pp. 295–312).

Radicioni, D. P., & Botta, M. (2006). A Methodological Contribution to Music Sequences Analysis. In F. Esposito & Z. W. Ras (Eds.), *Foundations of Intelligent Systems* (pp. 409–418). Berlin: Springer-Verlag.

Sperber, D., & Wilson, D. (1986). *Relevance. Communication and Cognition*. Oxford: Blackwell.

Wilson, D., & Sperber, D. (2004). Handbook of Pragmatics. In G. Ward & L. Horn (Eds.), (chap. Relevance Theory). Oxford: Blackwell.