

Connectionist F-structure Transfer

Ye-Yi Wang and Alex Waibel

Carnegie Mellon University

Abstract

A traditional transfer system in machine translation maps between language structures and an intermediate representation. Our connectionist transfer system maps from f-structures of one language to f-structures of another language. It encodes the intermediate representation implicitly in neural networks' activation patterns. The system is learnable, therefore it does not need any effort in hand-crafting the representation and mapping rules. Experiments show the system has good scalability and generalizability performance.

1 Introduction

Most of the current machine translation systems adopt an indirect strategy that maps between languages and an intermediate representation. The *interlingua* model (Nirenburg et al 1987) uses a language-independent intermediate representation. Design of the representation requires cross-linguistic expertise. The intermediate representation in a *transfer* model (White 1987) is language-dependent. Its design is relatively easier. However, multiple such representations are required for a multi-lingual translator. Both models rely upon hand-crafted mapping rules, which demand tremendous human effort.

The difficulties appeal for automatic learning mechanisms for intermediate representations and mapping rules. Chrisman (1991) proposed a connectionist confluent influence system that acquired the distributed inter-language representation of sentences during its learning to achieve the tight coupling between the representations of sentences in two different languages.

The approach was hard to scale up for larger tasks or to generalize for unseen inputs, mostly due to its over-simplified representation of sentences.

We present here a connectionist mapper. It can learn the transfer from a source language (English) LFG f-structure (Bresnan 1982) into its corresponding target language (German) f-structure. It does not need explicit intermediate representation or mapping rules. Instead, the connection patterns of the neural networks implicitly encode the rules and representation.

The domain of our task was the Conference Registration Telephony Conversations. It covered a wide range of topics related to conferences, such as registration, cancellation, hotel reservation, conference information inquiry, etc. The lexicon for the task contained about 400 English and 400 German words in root forms. About 300 pairs of f-structures of the English and German sentences were available from symbolic parsers.

A machine translation system for the Conference Registration task consisted of three parts: a parser deriving the f-structure from an input source language sentence, a mapper generating a target language f-structure from its source language counterpart, and a text generator producing a target language sentence from its f-structure. According to our experience, mapping between f-structures was the most difficult part, which required the hand-crafting of an intermediate representation and the rules that map between f-structures and the intermediate representation. An automatic transfer system is thus desirable. The system should have the following properties:

Learnability: The system should be able to learn the structure transfer automatically from paired samples. It should not require hand-crafting of any explicit representations and mapping rules.

Scalability: With limited retraining, the system should be able to deal with larger tasks with an expanded lexicon.

Generalizability: The system should have satisfactory performance on unseen inputs.

2 F-structure Representations

An f-structure is a structured functional representation of a sentence or a phrase. It is composed of a head, terminal features, and sub-structures. For the f-structure in Figure 1a, *SEND is the head. The contents in the inner brackets are the sub-structures, whose *grammatical relations* or *roles*¹ are labeled next to the brackets. The rest parts in Figure 1a are the terminal

¹We will use *Grammatical relation* interchangeably with the term *role*.

features. A sub-structure can be referred to with its grammatical relation or its phrasal category (NP, VP, ...). Thus the sub-structure $[_{subj} *YOU]$ can be called either a *SUBJECT sub-structure* or an *NP sub-structure*. The SUBJECT, RECIP and OBJECT sub-structures are the three *immediate sub-structures* of the top level f-structure in Figure 1a, because there is no intervening structure between these sub-structures and the top level f-structure. The DET sub-structure is an immediate sub-structure of the OBJECT sub-structure. If A is an immediate sub-structure of B, then B is the *parent structure* of A.

A symbolic f-structure cannot be presented to a neural network directly. Figure 1c-f illustrates how an f-structure can be coded as a network's input. Below are the terms used for the representation.

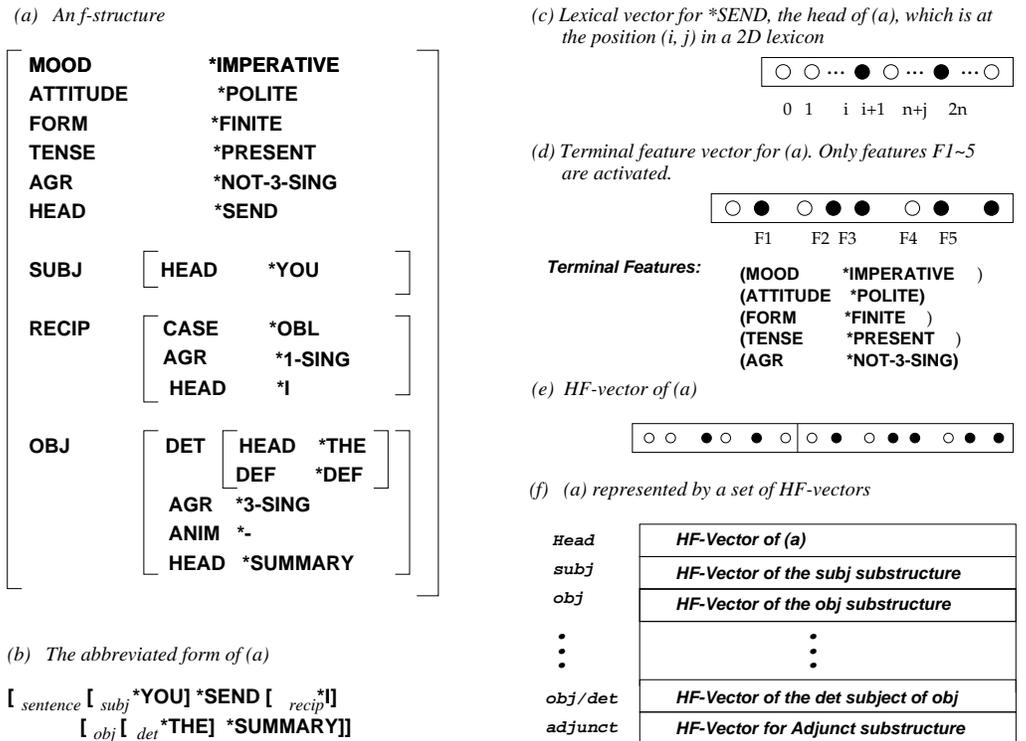


Figure 1: F-structure Representation. (a) an f-structure. (b) abbreviation. (c). lexical vector. (d). terminal feature vector. (e) HF-vector. (f) f-structure represented by HF-vectors.

A **lexical vector** is used to code a lexical item. Assuming that every lexical item is an entry in a two-dimensional space instead of a one-dimensional word list, we need two indices to specify the position of a lexical item in the space. Lexical vector is a 0-1 vector with exactly two elements being 1 (being activated). The positions of the two activated elements in the vector specify the two indices for an item in the 2D lexicon (Figure 1c)².

The **terminal feature vector of an f-structure** codes the terminal features of the f-structure. Each element of the vector corresponds to a feature-value pair like (TENSE *PRESENT). The vector, again, is a 0-1 vector with the activated elements indicating that their corresponding feature-value pairs are terminal features of the f-structure (Figure 1d). Since there are altogether around 60 different values for all the features used in the f-structure, the length of the terminal feature vector is around 60.

The **HF-Vector of an F-structure** is the concatenation of the lexical vector of the head and the terminal feature vector of the f-structure (Figure 1e).

Thus an f-structure can be represented by its HF-vector and its sub-structures' HF-vectors (Figure 1f).

3 The Mapper

A mapper is a push-down transducer that consists of:

1. a symbolic controller that assigns an f-structure transfer task to a neural network and interpreting the network's output. According to the interpretation, it recursively assigns the sub-structure transfer tasks to the related networks, and assembles these networks' results to the target f-structure;
2. seven neural networks that map phrasal f-structures between two languages. Each network is constructed for a phrasal category in the target language: IP(sentences), VP, NP, AP, PP, DP(determiners), and MP(miscellaneous, for phrases like "hello", "oh", etc.).

3.1 Phrasal Networks

A phrasal network has four layers: input, feature, hidden, and output layers (Figure 2a). The *input layer* consists of three parts:

²Viewing the lexicon as 2D reduces the length of the vector used to represent a lexical

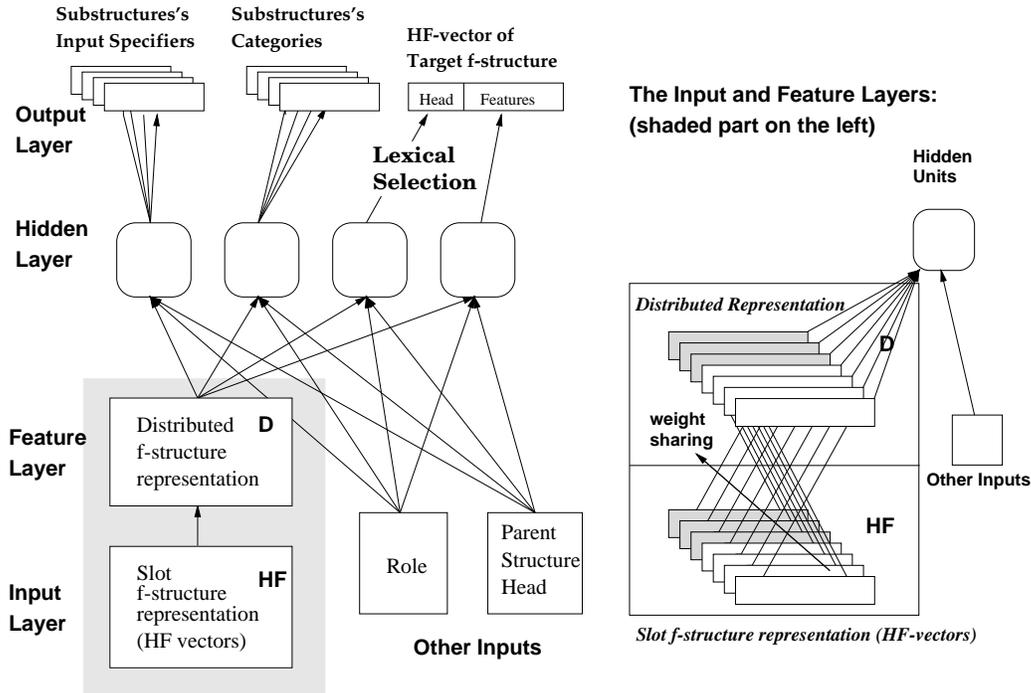


Figure 2: Phrasal Network Structure. (a): the architecture of a phrasal network. (b): details of the lowest two layers. The unshaded slots represent the input f-structure. The shaded ones represent the context f-structure.

Slots of the HF-vectors for an input f-structure and its context (parent) structure. Each slot corresponds to a fixed role.

An input f-structure may have sub-structures of arbitrary depth, but the networks must have fixed number of input slots. Therefore we cannot include all sub-structures' HF-vectors in the networks' input. Instead, we “peel off the shell” of an f-structure — only include the HF-vectors of the immediate sub-structures and their immediate sub-structures in turn for the input f-structure, and the HF-vectors of the immediate sub-structures for the context f-structure. Pre-analysis of the samples reveals the possible roles of the sub-structures that can occur at these levels in f-structures for the seven phrasal categories, and slots are then added to the input and feature layers of the corresponding phrasal networks to take as input the

item from n to $2\lceil\sqrt{n}\rceil$.

HF-vectors of the sub-structures with those possible roles.

Grammatical relation of the input source structure in its context³. This input is a 0-1 vector with exactly one activated element indicating the grammatical relation of the input structure.

Lexical vector of the head of the output f-structure’s parent structure (*p-head*). Sometimes, one input f-structure may be responsible for the generation of multiple target f-structures at different levels. For example, [_{sentence} GOODBYE] corresponds to both [_{sentence} AUF [_{obj} WIEDERHOEREN]] and its sub-structure [_{obj} WIEDERHOEREN] in the training samples. This input serves as a stack pointer, indicating the level at which the output f-structure should be generated.

The HF-vectors at the input layer are the local representations for the words and features in an f-structure. The activation patterns of the slots at the **feature layer** can be viewed as the automatically learned distributed representation of the input HF-vectors (Miikkulainen 1989). The input slots have one-to-one connections to the feature slots (Figure 2b). The slot-slot connections share weights in such a way that the connection from the *i*th unit in slot A at the input layer to the *j*th unit in slot A at the feature layer has the same strength as the connection from the *i*th unit in slot B at the input layer to the *j*th unit in slot B at the feature layer. The weight sharing makes the same HF-vector at different input slots result in the same pattern in their corresponding feature slots.

The output layer of a phrasal network has three parts:

The HF-vector of the f-structure to be generated. From this vector the head and the terminal features of the target f-structure can be recovered.

The Sub-Structures’ Input Specifiers. It consists of slots of 0-1 vectors. Each slot has at most one element being activated. And each slot corresponds to a sub-structure of a specific role of the target f-structure. The role of the sub-structure is implied by the position of the slot in the output layer. Each vector of *sub-structures’ input specifier* is of the size ($\|input\ layer\ slots\| + 1$). For an output slot in *sub-structures’ input specifier*, if it has one activated element, then the sub-structure with the corresponding role should be included as a part of the desired output f-structure. The position of the activated element in the slot indicates the input sub-structure (as specified by the slot number in the input layer) that is the counterpart of (and therefore is responsible for the generation of) the tar-

³Slot position only indicates the role of sub-structures, not the role of the input structure, since the HF-vector of the input f-structure with different roles always occupies the first slot.

get sub-structure, or *nil* when no input sub-structure is a counterpart of the output sub-structure. If a network does not activate any element in an output slot, then the slot's corresponding sub-structure should not be expected as a part of the desired target f-structure.

The Sub-Structures' Categories. It consists of slots of 0-1 vectors. There can be at most one element being activated in each slot, specifying one of the seven phrasal categories for the corresponding target sub-structure.

According to a network's output, the controller can build sub-structures recursively by assigning subsequent sub-structure mapping tasks to the networks of the categories specified in the output of *sub-structures' categories* at the output layer. The input f-structures of those mapping tasks are specified in the output of *sub-structures' input specifiers*. By combining the recursively built sub-structures and the head and the terminal features from the output HF-vector, the desired target f-structure can be produced.

4 An Example

The following example illustrates how the system works.

Source Sentence: *I would like to register for the conference*
 Source F-structure: [_{sentence} [_{subj} I] WOULD [_{xcomp} [_{subj} I] LIKE [_{xcomp} [_{subj} I] REGISTER [_{pp-adj} FOR [_{obj} [_{det} THE] CONFERENCE]]]]]]
 Target Sentence: *Ich wuerd mich gerne zur Konferenz anmelde*

(0) IP network

Input:

source: [_{sentence} [_{subj} I] WOULD [_{xcomp} [_{subj} I] LIKE [_{xcomp} REGISTER]]]

Output:

head: NIL

subs: sentence <WOULD VP>⁴ (1)⁵

features: (MOOD *DECLARATIVE)

F-structure assembled by the controller:

[_{sentence} [_{subj} PRONOUN] WERDE
 [_{xcomp} [_{subj} PRONOUN] [_{adj} GERNE] ANMELDEN [_{obj} PRONOUN]
 [_{pp-adj} FUER [_{obj} [_{det} DER] KONFERENZ]]]]]

(1) VP network

Input:

source: [_{sentence} [_{subj} I] WOULD [_{xcomp} [_{subj} I] LIKE [_{xcomp} REGISTER]]]
 context: NIL
 role: sentence
 p-head:⁶ NIL

Output:

head: WERDE
 subs: subj <I NP> (2)
 xcomp <LIKE VP> (3)
 features: ((CAT V) (PERSON 1) (MODAL +) (FORM FIN) ...)

F-structure assembled by the controller:

[[_{subj} PRONOUN] WERDE
 [[_{subj} PRONOUN] [_{adj} GERNE] ANMELDEN [_{obj} PRONOUN]
 [_{pp-adj} FUER [_{obj} [_{det} DER] KONFERENZ]]]]

In step (0), the controller first activates the IP network with the *source* input f-structure. There is no *context input* for the IP network, since the sentential f-structures are the top level f-structures in our task. From the network’s output, the controller knows that the head of the IP is NIL⁷. It also generates the sentential feature (*MOOD *DECLARATIVE*). And it interprets the output as that the only sub-structure of the sentence is a German VP, whose English counterpart is the (non-proper) sub-structure with the head *WOULD*⁸. Therefore it builds the target f-structure framework [NIL (MOOD *DECLARATIVE) [_{sentence} *]], and activates the VP network in step (1). Upon receiving the VP sub-structure returned from step (1), it combines that sub-structure with the f-structure framework, and collapses the NIL-headed f-structure to form the assembled f-structure shown as the output in step (0).

⁴The sub-structure’s input specifier and category are combined into a tuple here.

⁵The number in the parenthesis indicates the subsequent step of network activation for this sub-structure.

⁶*P-head* is the head of the target f-structure’s parent structure

⁷NIL-headed f-structure happens only when there is only one sub-structure or when there is an *xcomp* sub-structure. The NIL-headed f-structure must collapse into the only sub-structure in the first case, or into the *xcomp* sub-structure in the second case. All terminal features and other sub-structures are moved into the *collapsed-into* sub-structure during collapsing.

⁸The network actually specifies the slot at the input layer instead of the lexical item *WOULD*.

In step (1), the input *source* was determined in step (0), since the *sentence* sub-structure’s head was “*WOULD*” according to the IP network’s *sub-structure’s input specifier* in step (0). The *context* input is NIL because the *source* f-structure does not have a parent f-structure. The input *role* has the value *sentence* because the slot position of the output sub-structure in step (0) implies the grammatical relation of the sub-structure is *sentence*. The input *p-head* is NIL because the head of target f-structure in step (0) was NIL as specified by the output HF-vector there.

The VP network maps the input f-structure to its German counterpart by specifying (a) the head of the German VP structure *WERDE* and the terminal features of the German VP structure in the output HF-vector, and (b) the input specifiers and the categories of the sub-structures of the target German VP f-structure. To build detailed sub-structures for this VP f-structure, the controller will activate the NP network with the input of the English sub-structure with the head “*I*” and the VP network with the input English sub-structure with the head “*LIKE*” in the subsequent steps, and combine the sub-structures returned from these subsequent steps into the f-structure framework $[[_{subj} *] \text{WERDE} [_{xcomp} *]]$. The combined structure is then returned to step (0) to be integrated into the top level f-structure framework there.

5 Training, Testing and Performance

From the 300 sentential f-structure pairs, we extracted all the German NP sub-structures, their grammatical relations and their parent structures’ heads. We labeled their English counterparts⁹. These were all the information required for the training of the NP network. About 700 samples for the NP networks were created this way. The training samples for the other networks were prepared in the same way. The NP network had the most samples, while the MP network had the least of 89 samples. Standard back-propagation was used to train the networks. We also tried the information-theoretical networks (Gorin et al 1991) to generate the head of a target structure in the HF-vector, which required less training time and achieved comparable performance as the network trained with pure back-propagation algorithm (Wang 1994). The training took 500 to 2000 epochs for different networks, and the training time ranged from one hour to three days on DECStation 5000. The mapper achieved 92.4% accuracy on the

⁹An NP’s counterpart is not necessary to be an NP.

training data¹⁰.

Learnability: The connectionist f-structure transfer described above did not require any hand-crafted rules or representations. The structure transfer was learned automatically. By clustering the distributed representations of words learned by the networks, i.e., the activation patterns of a feature slot when a lexical item was presented to its connected input slot, we had some interesting findings about what was learned by the networks. One of them was that the feature patterns for English nouns in the DP network were clustered into three classes, which reflected the three genders of German nouns: the German translations of the words in each class were roughly of the same gender. Another finding was about the classification of verbs. When we clustered the feature patterns for verbs in the VP networks, we found some intransitive verbs like *register* in the same class as most of the transitive verbs. This seemingly strange classification is not odd at all if we consider the fact that the German translation for *register*, “*anmelden*”, is a transitive verb. These two independent findings reveal the networks’ ability to discover some linguistic features of the target language and use it in the representation of an entity of the source language which does not possess those features. This is exactly what a symbolic transfer are supposed to do: using an intermediate representation which reflects the linguistic features of the two languages in question (even if one of the languages may have degenerated form for a specific feature,) and thus being able to make a “transfer” at both the lexical and structural level into corresponding structure in the target language. Our system learned the intermediate representation automatically, although it was not expressed explicitly in symbolic forms but encoded in the networks’ activation patterns. Because the development of this representation was integrated into the process of automatic learning of f-structure mapping, it tended to include in the intermediate representation the important language specific linguistic features which were directly relevant for the ultimate purpose of structure transfer. In the other words, the learning of the intermediate representation was focused on the purpose of improving the transfer performance. This is one of the biggest advantage of this approach over the hand-crafted intermediate representation.

Scalability: We did a preliminary scalability experiment. We extended the source and target language lexicon by 2%, and made 30 new f-structures with these new lexical items. Trying to scale up from what was already learned, we froze all but the input-feature connections, trained the network

¹⁰A source language f-structure is said to be *accurately* mapped if the generated target language f-structure is exactly the same as desired in the sample.

for about 40 epochs with the new data, then fine-tuned all the connections with old and new data for a few epochs. In doing so, we let the networks first learn the new words to derive their distributed representations, and then learn the structure mapping for the new data later. This approach was based on the observation that a big portion of the new English words were translated to some German words already in the lexicon, which in turn was translated from some English words in the old training data. These old English words were mostly the synonyms of the new English words. By freezing the other connections and training only the input-feature connections, we hoped the networks to be able to develop the distributed representation for a new word similar to the already-learned representations of its synonyms.

This approach greatly reduced the learning time for new words, since the one layer back-propagation was much fast than the full-blown learning. The mapper with the new phrasal networks that were retrained this way achieved 83.3% accuracy on the new data, without affecting the performance on the old data.

Generalizability: A separate set of data was used to test the generalization performance of the system. The testing data was collected from people not associated with our researches. The data was compared with the training corpus, and the sentences that appeared in the training data were removed. An LR parser parsed the sentences to English f-structures. The English were translated into German manually, and the translations were parsed by a German LR-parser. We picked the most probable structure when a parsing result was ambiguous. There were 154 f-structure pairs after we eliminated the wrongly-parsed sentences. The mapper achieved 61.7% accuracy on the testing data. Considering the limited number of training samples, this performance was encouraging. Previous research as in (Chrisman 1991) did not generalize to deal with unseen data.

6 Discussion

The application of the connectionist transfer described in this paper has its restrictions. First, it requires well-formed f-structures for both the input and output sentences. This greatly limits the applicable domain of the approach to well-structured “clean” languages. It is difficult to use this approach for spoken language where performance data like ungrammatical utterances, noises, false starts are pervasive.

Another restriction is that this approach can only achieve satisfactory performance when the input and output languages are similar, in the sense

that the translation equivalents in the two languages mostly have similar recursive f-structures. Although the system can deal with structurally different input/output sentences, like the aforementioned example of [*sentence* GOODBYE] and [*sentence* AUF [*obj* WIEDERHOEREN]], we believe that the performance would drop significantly if drastic structure differences between translation equivalents are very common for the two languages in question. Fortunately, as shown by our data, the structural difference between English and German is not so drastic to ruin our system's performance.

Although we had done some scalability experiment, it is unclear how the system will perform if we increase the lexicon significantly instead of by 2%. Because of the limitation of available data, we found it very difficult to conduct scalability experiments with much more expanded lexicon. We hope that with stable incremental performance, the system can be gradually and easily retrained to deal with more complicated problems.

7 Conclusion

Aiming at the difficulties in symbolic transfer, we have proposed a connectionist transfer system that maps between f-structures of two languages. It can discover meaningful linguistic features by learning. Its performance is promising with respect to learnability, scalability and generalizability.

REFERENCES

- Bresnan, Joan. 1982. *The Mental Representation of Grammatical Relations*. Cambridge: MIT Press.
- Chrisman, Lonnie. 1991. "Learning recursive distributed representations for holistic computation". *Connection Science*, 3(4):345–366.
- Gorin, Allen L., Steve E. Levinson, A. N. Gertner & E. Goldman. 1991. "Adaptive acquisition of language". *Computer Speech and Language*, (5):101–132.
- Miikkulainen, Risto & Michael G. Dyer. 1989. "A modular neural network architecture for sequential paraphrasing of script-based stories". *Proceedings of the International Joint Conference on Neural Networks*. IEEE.
- Nirenburg, Sergei, Victor Raskin & Allen B. Tucker. 1987. "The structure of interlingua in TRANSLATOR". *Machine Translation: Theoretical and Methodological Issues* ed. by Sergei Nirenberg. Cambridge: Cambridge University Press.
- Wang, Ye-Yi. 1994. "Dual-coding theory and connectionist lexical selection". *Proceedings of the 32nd Annual Meeting of the Association for Computational Linguistics, student session*, pages 325–327.
- White, John S. 1987. "The research environment in the METAL project". *Machine Translation: Theoretical and Methodological Issues* ed. by Sergei Nirenberg. Cambridge: Cambridge University Press.