Language Generation for Broad-Coverage Explainable Cognitive Systems

Marjorie McShane

Ivan E. Leon

Language Endowed Intelligent Agents Lab

Cognitive Science Department | Rensselaer Polytechnic Institute

Introduction

- This paper describes key design features of the natural language generation (NLG) capabilities we've recently begun developing within the OntoAgent cognitive architecture

- Natural language understanding (NLU) in OntoAgent is described in **Linguistics for the Age of AI** (McShane and Nirenburg; MIT Press, 2021; open access)

- This paper answers the natural question: **and what about NLG?**

The OntoAgent content-centric cognitive architecture

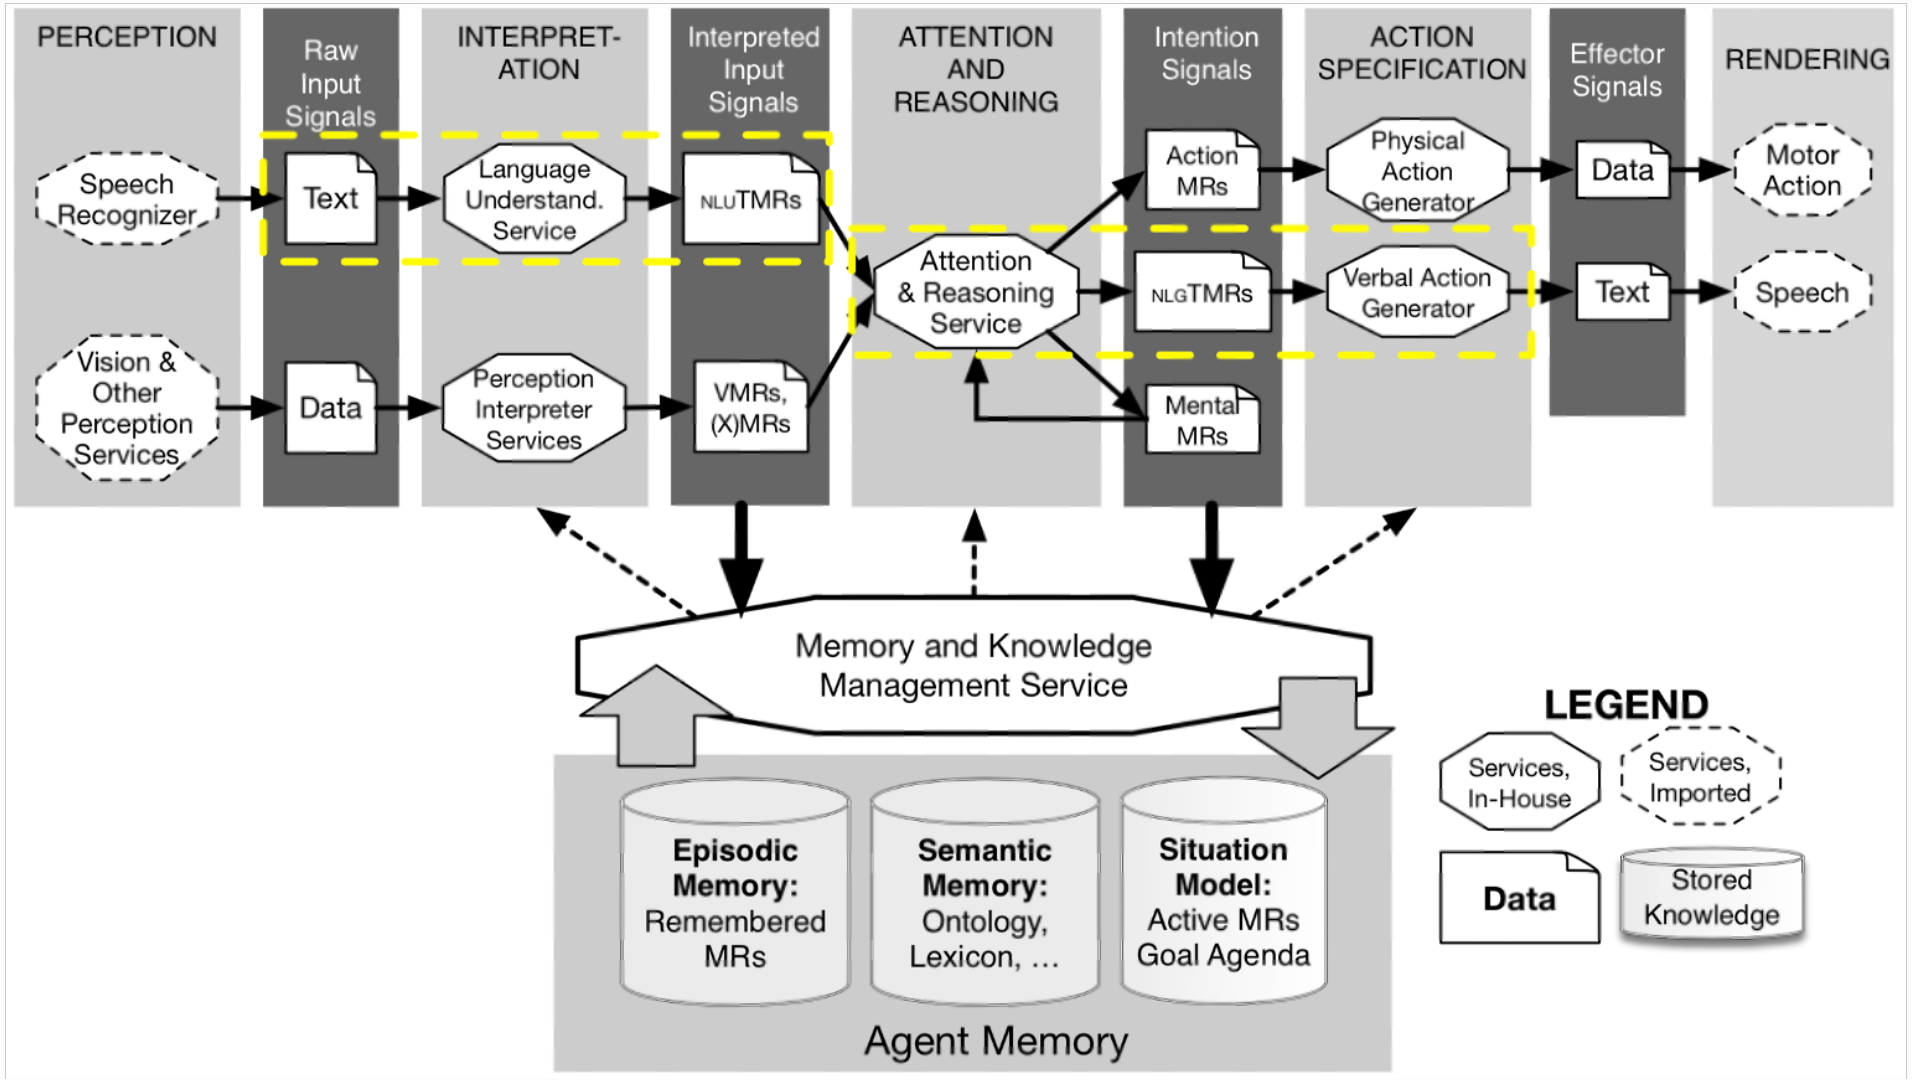

Content Specification

- The prerequisite for NLG is content specification: deciding what meaning is to be conveyed
- In applications, this is provided by agent reasoning
- However, in the current state of the art, nearly all cognitive systems have very narrow coverage, so they can provide few content specifications

Content Specification (cont.)

- Our goal, by contrast, is to solve the problems of broad-coverage NLG (polysemy, paraphrase, coreference resolution, etc.)
- Where will the content specifications to foster such R&D come from?
- **Our NLU system**

NLU provides content specifications for NLG:
A practical approach to R&D

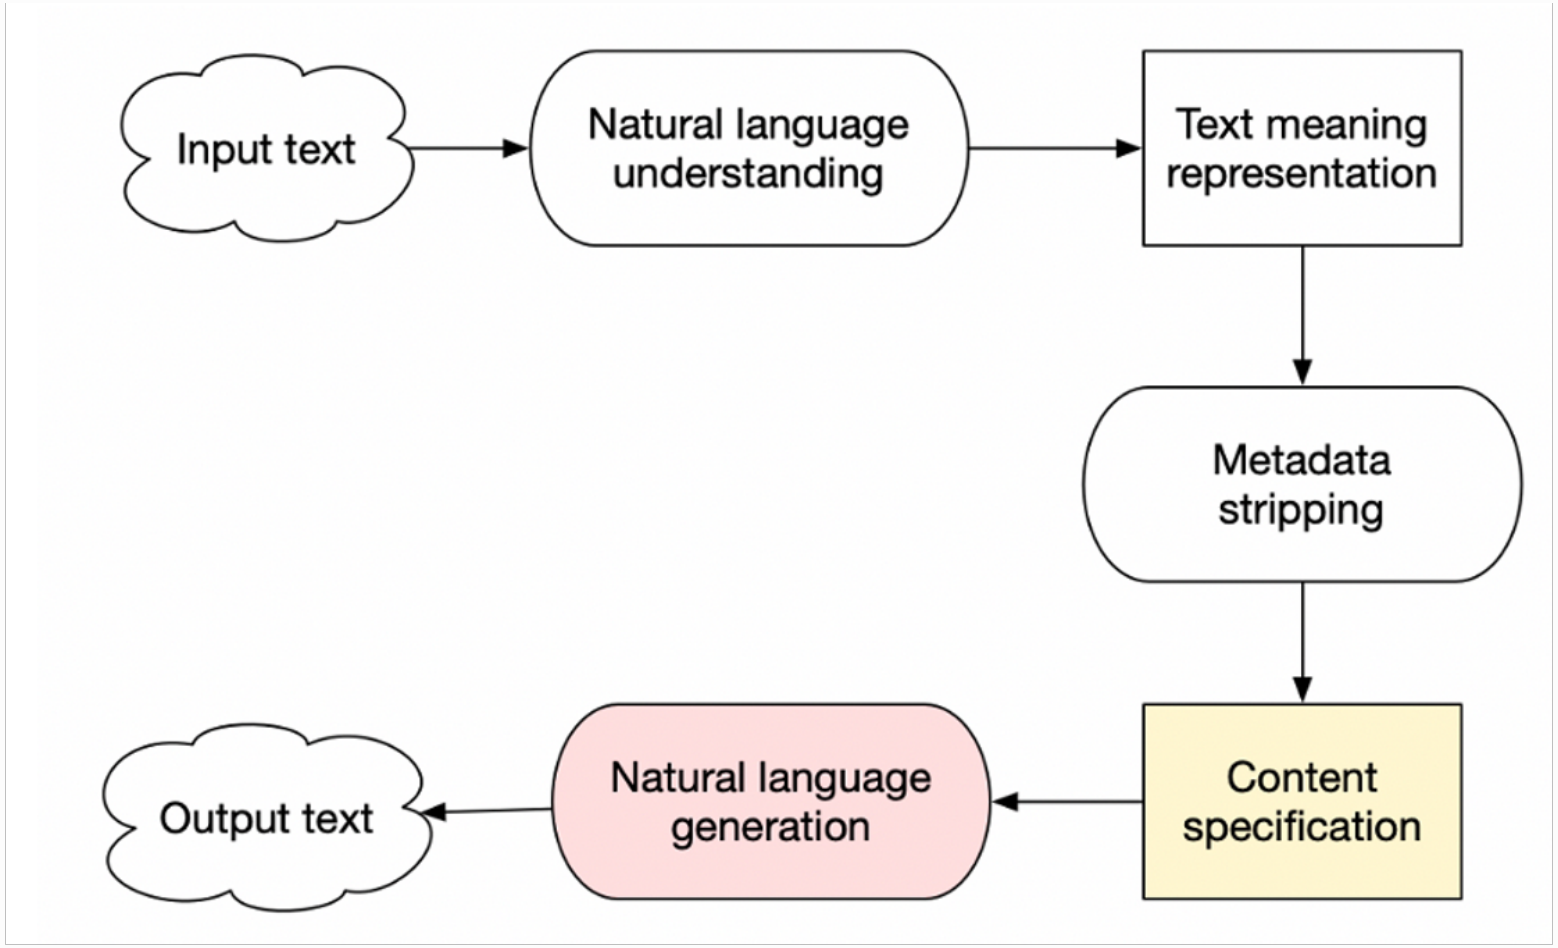

The Architecture of NLG

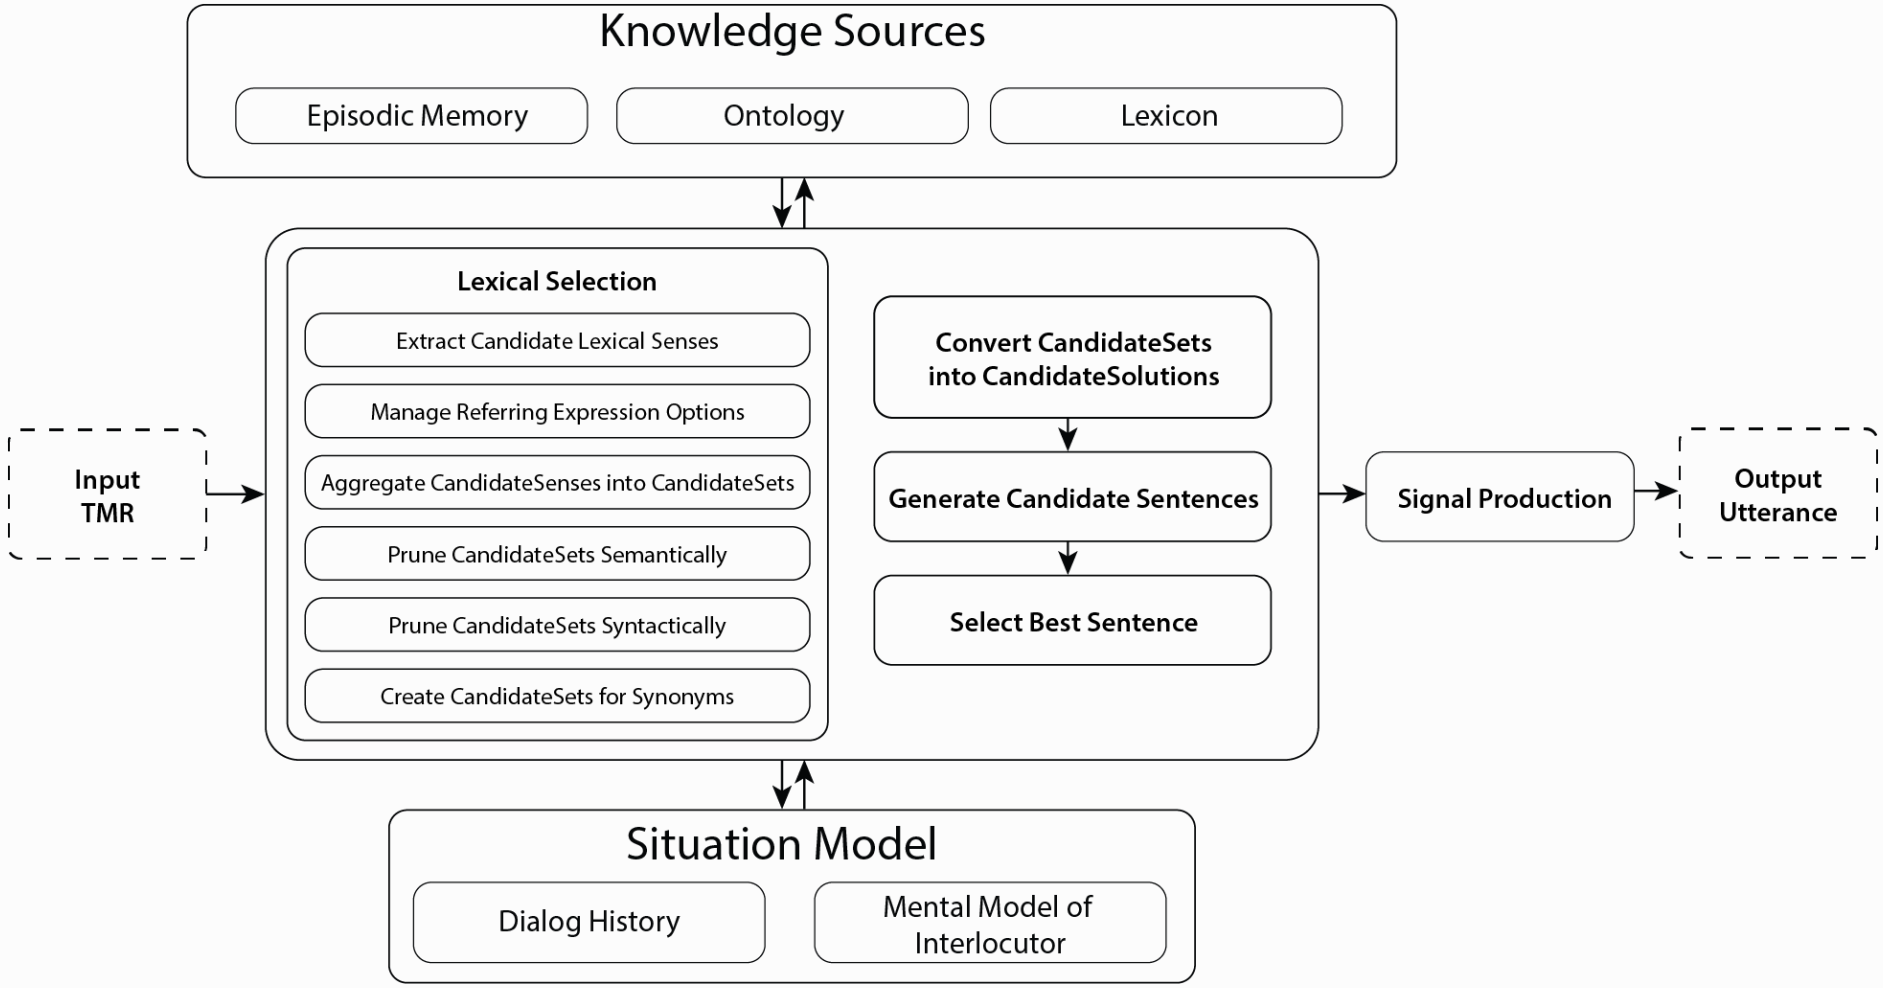

The Architecture of NLG – Lexical Selection

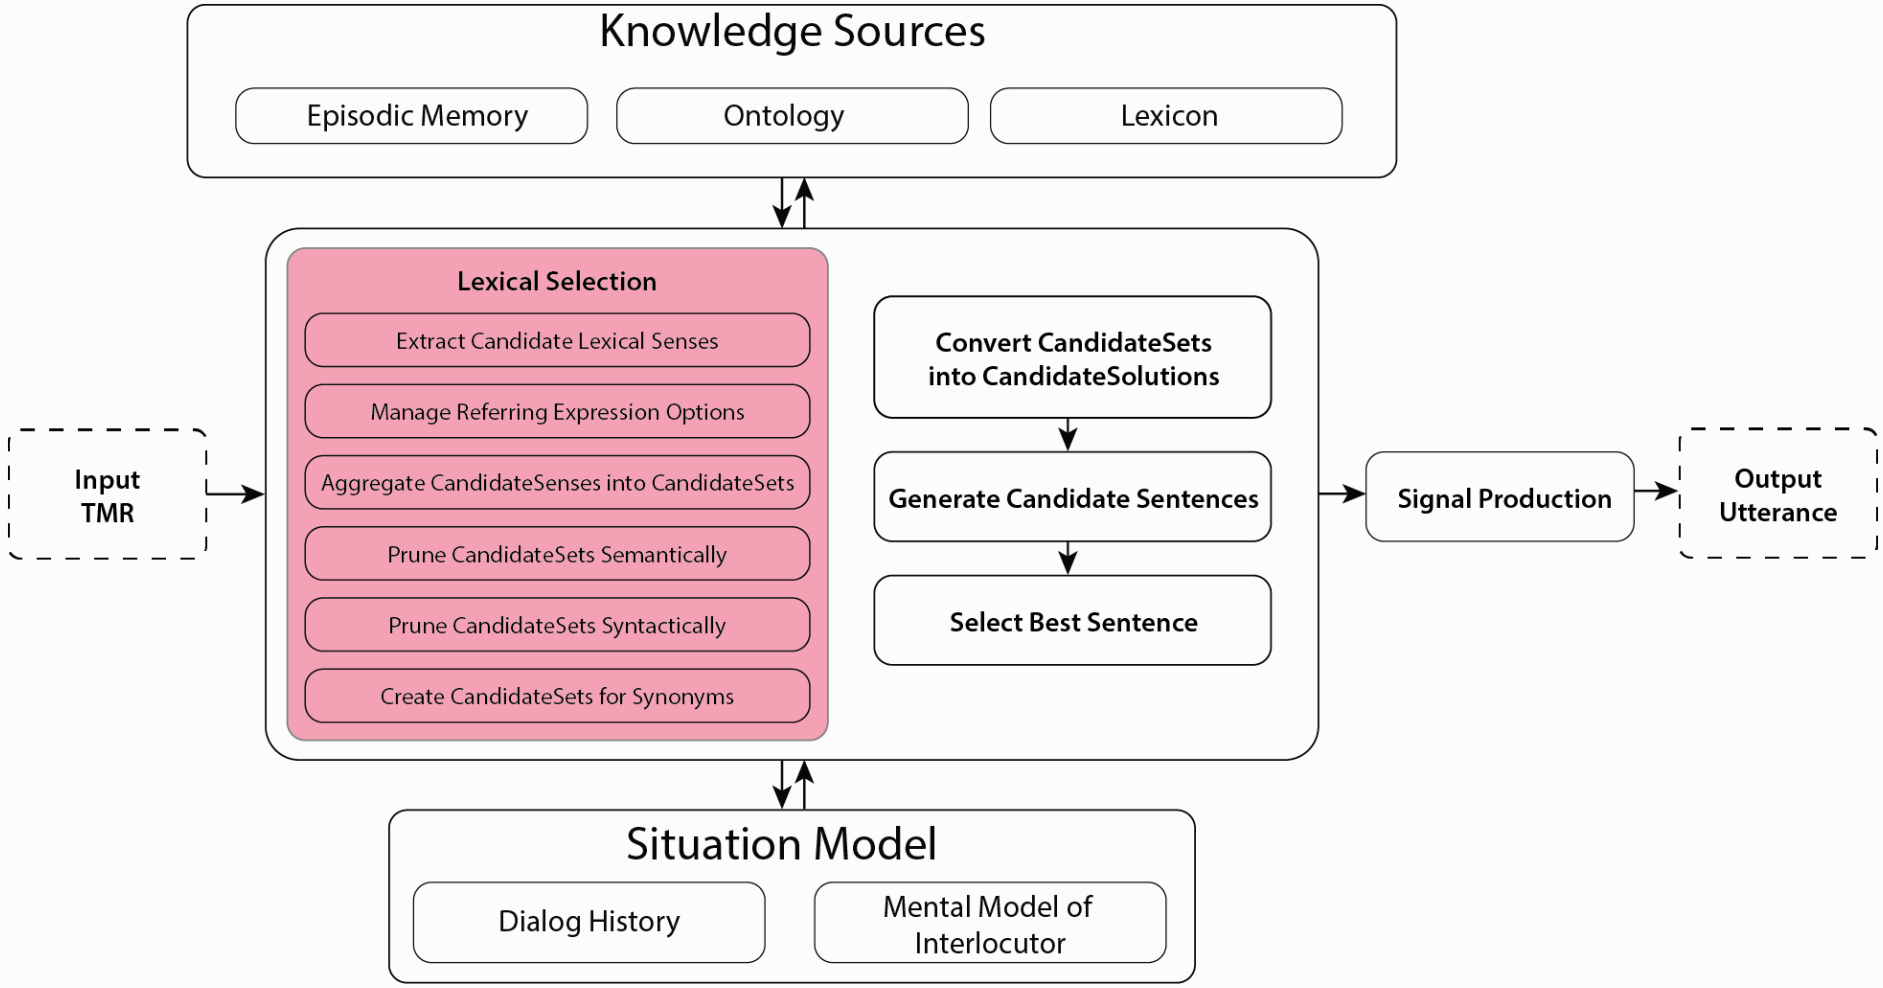

Lexical Selection

- To give just a taste of why this is difficult, words and phrases have syntactic and semantic requirements, listed in the lexicon, that must be fulfilled within their sentence

- Word senses can combine into a sentence only if all their needs are fulfilled by other elements of the sentence

- Below is an informal sketch of the verb ‘fix’ that means FASTEN
- Syntactically it requires a *subject*, *direct object*, and *prepositional phrase* with “to”
- Semantically it requires an **AGENT**, **THEME**, and **DESTINATION**, all of whose fillers must match the semantic constraints specified in the ontological frame for **FASTEN**
- In short, there are a lot of constraints to check and keep track of

Top-level info	Syn-Struc		Sem-Struc	
fix-v2	subj	$var1	FASTEN	
def to attach something	v	$var0	AGENT	^$var1
to something else	directobject	$var2	THEME	^$var2
ex He fixed the bookshelf to the	pp		DESTINATION	^$var4
wall.	prep	$var3 (root to)	^$var3 (null-sem+)	
synonyms attach, fasten, secure	n	$var4		

The Architecture of NLG – CandidateSolutions

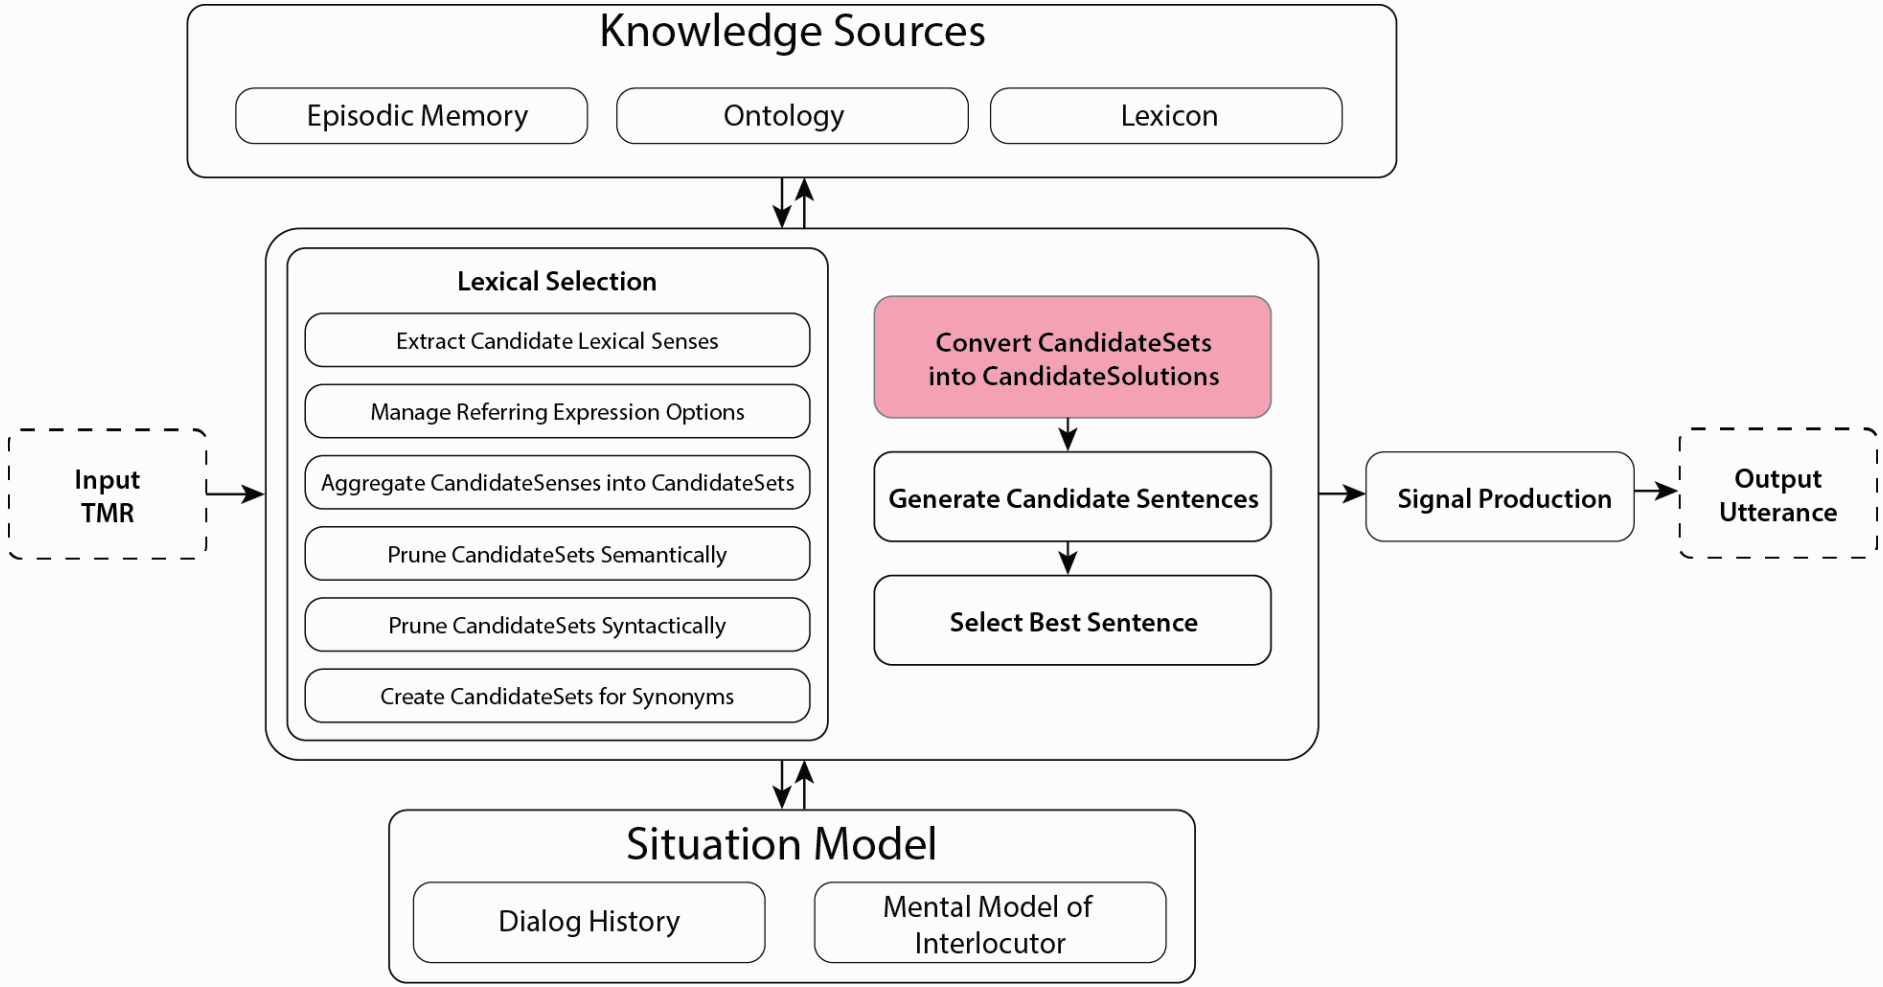

The Architecture of NLG – Generate Candidates

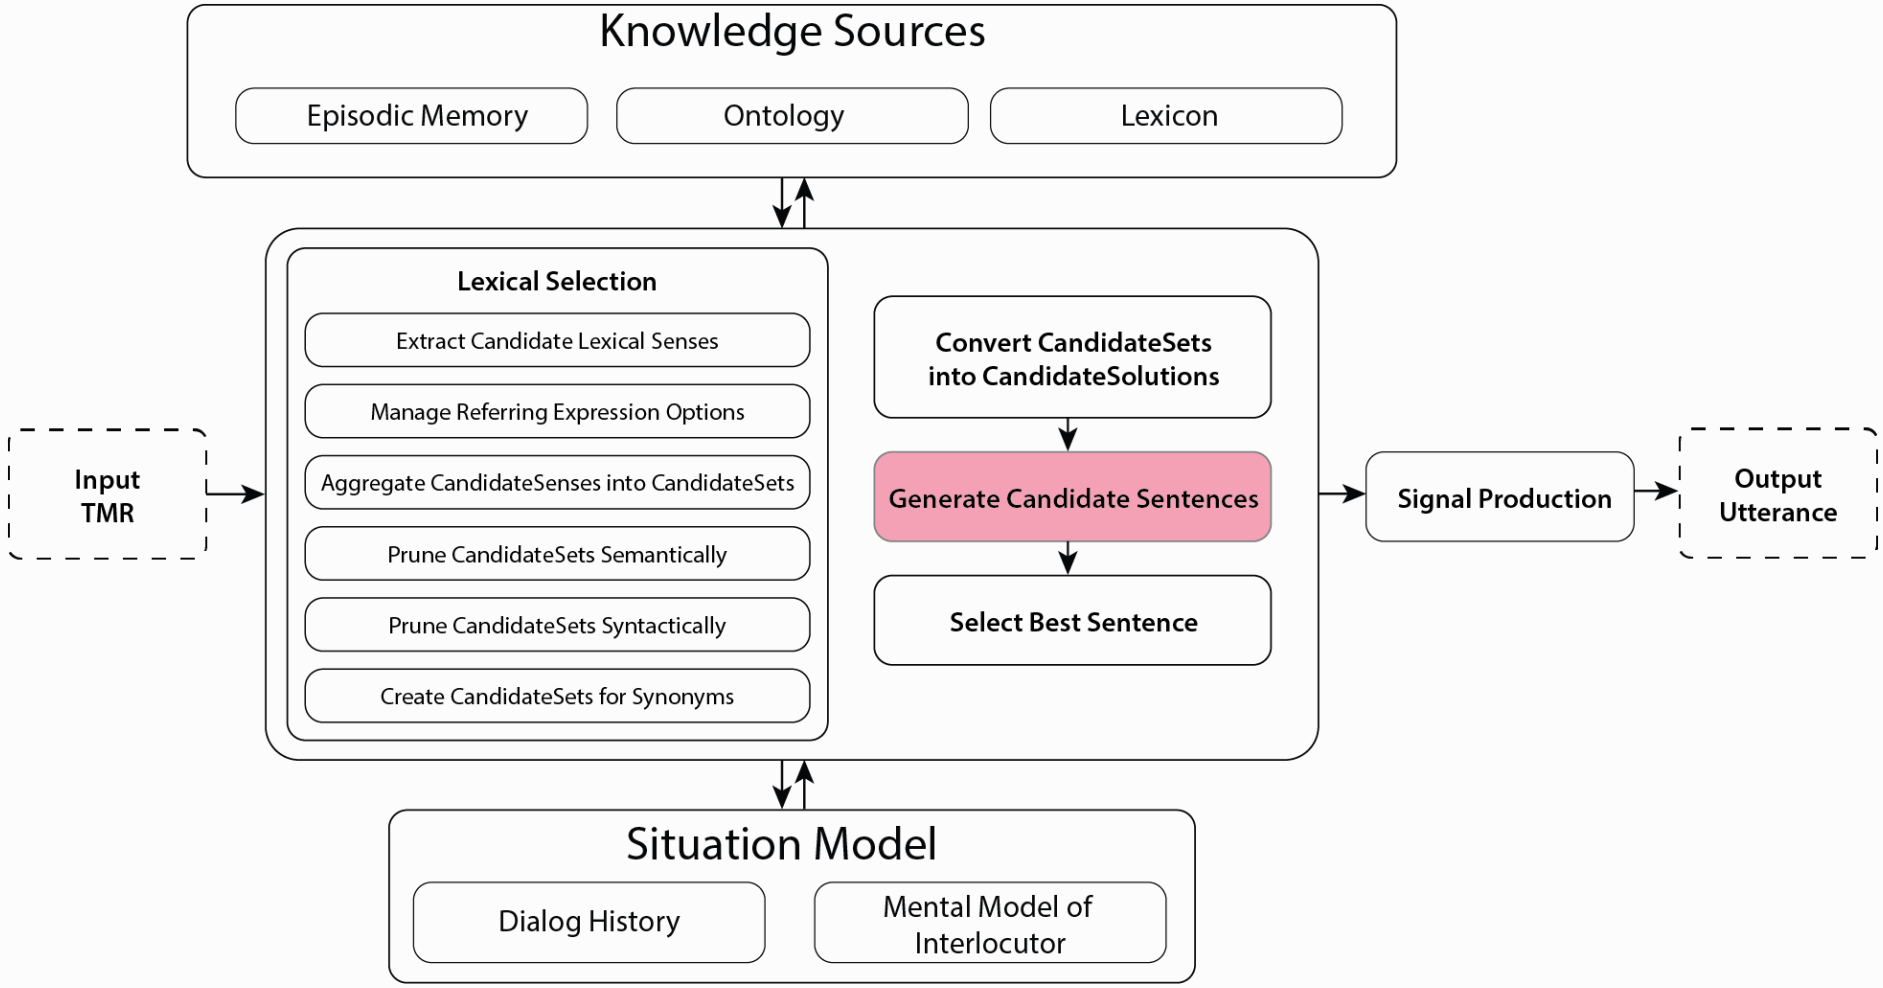

The Architecture of NLG – Candidate Selection

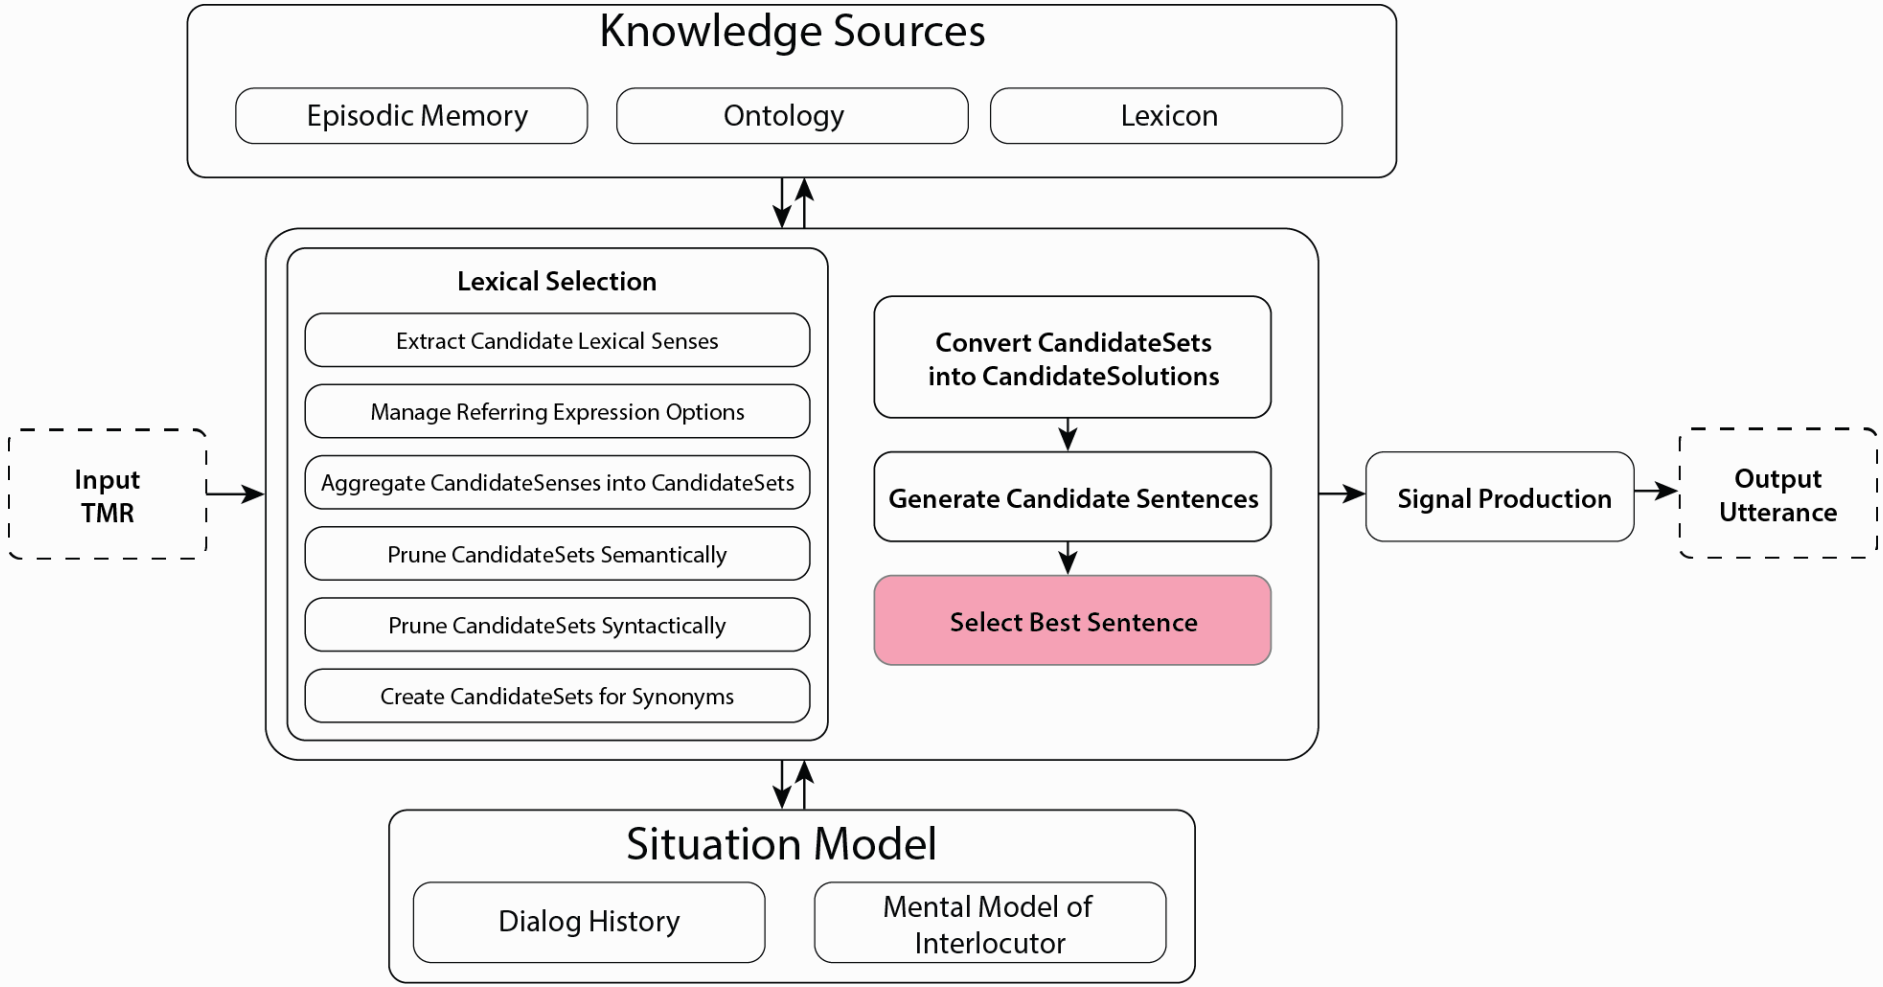

FASTEN-1

AGENT	HUMAN-1
THEME	PICTURE-1
DESTINATION	WALL-1
TIME	(< find-anchor-time)
from-sense	*fix-v2*

HUMAN-1

HAS-NAME	'Tom'
AGENT-OF	FASTEN-1
from-sense	*Tom-n1 ; from onomasticon*

PICTURE-1

THEME-OF	FASTEN-1
from-sense	*painting-n1*

WALL-2

DESTINATION-OF	FASTEN-1
COREFER	WALL-1
from-sense	*wall-n1*

FASTEN-1

AGENT	HUMAN-1
THEME	PICTURE-1
DESTINATION	WALL-1
TIME	(< find-anchor-time)
from-sense	*fix-v2*

HUMAN-1

HAS-NAME	'Tom'
AGENT-OF	FASTEN-1
from-sense	*Tom-n1 ; from onomasticon*

PICTURE-1

THEME-OF	FASTEN-1
from-sense	*painting-n1*

WALL-2

DESTINATION-OF	FASTEN-1
COREFER	WALL-1
from-sense	*wall-n1*

FASTEN-1	
AGENT	HUMAN-1
THEME	PICTURE-1
DESTINATION	WALL-1
TIME	(< find-anchor-time)
from-sense	*fix-v2*

HUMAN-1	
HAS-NAME	'Tom'
AGENT-OF	FASTEN-1
from-sense	*Tom-n1 ; from onomasticon*

PICTURE-1	
THEME-OF	FASTEN-1
from-sense	*painting-n1*

WALL-2	
DESTINATION-OF	FASTEN-1
COREFER	WALL-1
from-sense	*wall-n1*

FASTEN-1

AGENT	HUMAN-1
THEME	PICTURE-1
DESTINATION	WALL-1
TIME	(< find-anchor-time)
from-sense	*fix-v2*

HUMAN-1

HAS-NAME	'Tom'
AGENT-OF	FASTEN-1
from-sense	*Tom-n1 ; from onomasticon*

PICTURE-1

THEME-OF	FASTEN-1
from-sense	*painting-n1*

WALL-2

DESTINATION-OF	FASTEN-1
COREFER	WALL-1
from-sense	*wall-n1*

FASTEN-1

AGENT	HUMAN-1
THEME	PICTURE-1
DESTINATION	WALL-1
TIME	(< find-anchor-time)
from-sense	*fix-v2*

HUMAN-1

HAS-NAME	'Tom'
AGENT-OF	FASTEN-1
from-sense	*Tom-n1 ; from onomasticon*

PICTURE-1

THEME-OF	FASTEN-1
from-sense	*painting-n1*

WALL-2

DESTINATION-OF	FASTEN-1
COREFER	WALL-1
from-sense	*wall-n1*

Tom secures the painting to the wall.

Tom secures the picture to the wall.

Tom fastens the painting to the wall.

Tom fastens the picture to the wall.

Tom fixes the painting to the wall.

Tom fixes the picture to the wall.

…

So, what about explainability?

- This line of work enables agents to generate natural language from content specifications that explain agent behavior

- The nature of our knowledge bases and processing allows an agent to point directly to the knowledge and internal processes which generated a particular behavior; in this case the behavior of generating a natural language sentence(s)

- How does this work differ from the knowledge-based work on NLG that was carried out half a century ago and proved to be so difficult that people gave up on it?

- None of the problems have gone away, but our understanding of both the content (linguistics, cognitive modeling, ontology) and strategy have evolved

- As regards strategy, we focus on ***lifelong learning*** (to overcome the so-called ”knowledge bottleneck”) and integrating language processing with other channels of perception and general reasoning

- It has become clear at this point that AI will not reach its potential without explainability, so no matter how difficult it is to build explainable systems, it must remain on agenda

- The goal of this work is to develop natural language generation capabilities for a content-centric cognitive architecture called OntoAgent

- R&D is still in the early stages; however, this paper describes some of the core design features of our natural language generation system

- Our goal is to solve some of the problems of broad coverage NLG, including many of the micro-theories covered by our NLU system, and to bootstrap development of our NLG system with meaning representations from open-domain texts

Thank you!