

## Processing Natural Language About Ongoing Actions in ECG

**Steve Doubleday**

University of California, Irvine,  
Irvine, California 92697  
stevedoubleday@gmail.com

**Sean Trott**

International Computer Sciences Institute  
1947 Center Street  
Berkeley, California 94704

**Jerome Feldman**

International Computer Sciences Institute  
1947 Center Street  
Berkeley, California 94704

### Abstract

Actions may not proceed as planned; they may be interrupted, resumed or overridden. This is a challenge to handle in a natural language understanding system. We describe extensions to an existing implementation for the control of autonomous systems by natural language, which enables such systems to handle incoming language requests regarding actions. Language Communication with Autonomous Systems (LCAS) has been extended with support for X-nets, parameterized executable schemas representing actions. X-nets enable the system to control actions at a desired level of granularity, while providing a mechanism for language requests to be processed asynchronously. Standard semantics implemented include requests to stop, continue, or override the existing action. The specific domain demonstrated is the control of motion of a simulated robot, but the approach is general, and could be applied to other domains. (Demonstration video: <https://youtu.be/8PwHpng3Nj8>)

### Introduction

A challenge for natural language understanding systems that interface with robots is responding to changes in the situation communicated through language. Robots and other autonomous systems must deal with interrupts from the physical world, but to our knowledge none also incorporate a mechanism to interrupt an action through language. Actions may need to be interrupted, cancelled, stopped, resumed, or overridden. Underlying these responses are assumptions about the state of events. The temporal structure of events is expressed in linguistic constructions termed *aspect* (Narayanan 1997); an action can be about to start, or ongoing, or just completed. An action that is interrupted or stopped is presumed to be ongoing; one that is cancelled is presumed not to have started yet.

We have extended an existing ECG-based system, Language Communication with Autonomous Systems (LCAS) (Trott et al. 2015), to enable it to handle language requests regarding the actions of simulated robotic agents. LCAS is

the current implementation of a general approach to language understanding termed the Neural Theory of Language (NTL), which has been under development for many years. The NTL approach separates the computational problems of language understanding into two phases, *analysis* and *enactment* (Feldman, Bryant, and Dodge 2009). In the system architecture, these are termed the “language side” and the “action side”, respectively. During the analysis phase, an *Embodied Construction Grammar* (ECG) analyzer (Feldman, Bryant, and Dodge 2009) parses the text of the user’s request to the system, creating a formal semantic representation of the text that is further refined to produce inputs to the enactment phase. In the enactment phase, the focus of the current work, the actions referenced by the text are carried out. Potentially, these actions might also be simulated, to enable inference about future states; this is a goal for future work.

The extensions to LCAS center around the use of *X-nets* (Narayanan 1997), which implement parameterized executable schemas representing actions. An X-net is implemented as an extension to a *Petri net* (Murata 1989). The Petri net formalism is useful for the description and analysis of concurrent processing in distributed systems. An X-net implements many of the standard semantics we associate with action. Actions can be enabled, started, ongoing, or done, which expresses the normal execution path. Actions can also be suspended, resumed, or restarted, and handle exception conditions. In addition to having a shared standard action semantics, each X-net is tailored to meet the demands of a set of related tasks. In this case, a Move X-net serves as an interface to drive the motion of a simulated robot. Asynchronous communication links the language side of the system, handling new user requests as they arrive, to the action side of the system, which drives the motion of the simulated robot.


### X-Nets and Simulation Semantics

There is considerable evidence that language understanding proceeds in part by simulating the actions implied by the text (Bergen, Chang, and Narayan 2004; Bergen 2012; Schilling and Narayanan 2013). Some researchers have an-

alyzed the link between action and language by mapping from body parts to nouns, and mapping from body motions to verbs (Pastra and Aloimonos 2012). Focusing on the physical motions that constitute the elements of an action, it is possible to build a grammar of actions as sequences of body poses. This non-verbal grammar of action has been used to simplify the problem of recognizing the actions of other agents (Summers-Stay et al. 2012; Yang et al. 2015). A further extension is to interpret sequences of action primitives as action plans or fragments of action plans, which simplifies the creation of new action plans or the interpretation of the observed actions of other agents (Guha et al. 2013). Potentially, this approach might also be used as another way to simplify the problem of language acquisition, which relies on non-verbal information to narrow the range of possible interpretations of under-specified linguistic expressions (Chang and Gurevich 2004; Chang 2008).

The current research focuses on a different problem – the link between linguistic expressions and actions that an agent already “knows” how to do, and may be in some stage of performing. Actions and the simulation of actions can be modeled by executable schemas parameterized by language (Narayanan 1997; 2010; Schilling and Narayanan 2013). For example, verbs concerning motion can be organized in terms of an implicit ascending speed parameter: “crawl”, “amble”, “walk”, “run”, “dash”. Linguistic aspect involves the use of grammar to parameterize the state of actions in reference to a general event model. An action can be impending, under way, or completed.


An executable schema can be implemented as an X-net, a parameterized Petri net representation of an action. The Petri net formalism models distributed processing as a set of states and transitions among states, connected by directed arcs. Transitions fire when the pre-conditions defined by their input arcs are met, generating updates to the places connected to their output arcs. The pre- and post-conditions take the form of one or more tokens in a place. The vector of token counts for all places in the Petri net constitutes its current state, or *marking*. Figure 1 depicts a transition that will not fire; Figure 2 depicts an enabled transition before and after firing; after firing the transition is disabled. In this case, firing consumes two tokens and creates one token.



Inputs are not all marked; transition disabled.

Figure 1: Transition not yet enabled, as all input places are not marked.

X-nets define a standard action semantics, states and transitions that are common to all actions. Depicted as an event



Both inputs marked; transition enabled. Fires, creating one token.

Figure 2: When all input places are marked, transition is enabled and fires, marking the output place

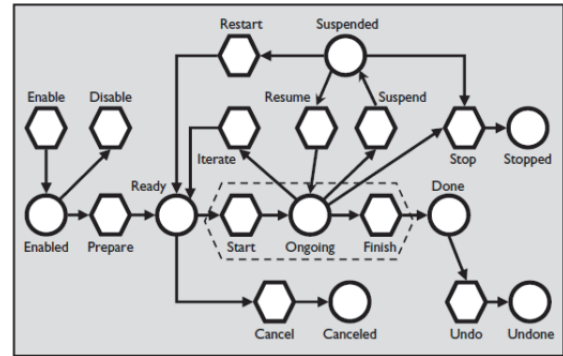


Figure 3: Event graph (controller) depicting the structure of events (Narayanan 2010)

graph in Figure 3, this serves as a useful intermediate-level abstraction for the representation of actions (Narayanan 2010; Doubleday 2017). Scalar parameters can be passed to an X-net, enabling it to perform computations to modify its processing, e.g., to move faster or slower. X-nets can also be composed of more granular X-nets, where each X-net has some ability to perform error recovery locally, with escalation to higher levels as needed.

These X-net capabilities enable a system to control actions. They also enable a system to reason in a consistent way about actions. X-nets provide a natural way to model linguistic aspect. By examining the current marking of an X-net, we can distinguish between actions that are about to begin, or are ongoing, or just completed. X-nets can also support reasoning about events in general (Sinha 2008), including hypothetical or counter-factual language. Finally, X-nets can support the use of simulation as a planning mechanism. Rather than viewing a sequence of actions as a vector in a pre-defined search space, e.g., (Guha et al. 2013), the agent might explore future alternatives by simulating the execution of one or more X-nets, and evaluating the results of the simulation (Hesslow 2002; Schilling and Narayanan 2013). For a formal analysis of the applicability of X-nets to the modeling of actions, including limitations and neural plausibility, see Narayanan (1997).

## System Architecture

### System for Natural Language Understanding

We have implemented an ECG-based system for natural language understanding (Trott et al. 2015; Khayrallah, Trott, and Feldman 2015; Epepe et al. 2016). The general architecture for this system is depicted in Figure 4. Most of our previous work has focused on extending the core system framework to the robotics domain, in applications such as Morse (Echeverria et al. 2011) and ROS (Epepe, Trott, and Feldman 2016), but the system was designed to facilitate simple re-targeting to new domains and applications. Ongoing work includes applications such as real-time strategy games (StarCraft), metaphor analysis, and mental space modeling.

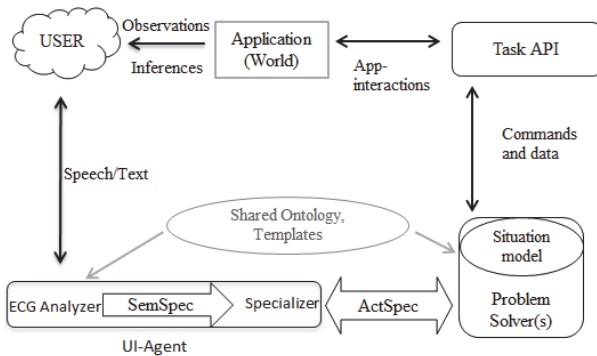


Figure 4: System Architecture for natural language understanding.

The system is modular; at the highest level, it can be divided into a *language side* and an *action side*. The language side, depicted on the left in Figure 4, receives speech or text as input, and produces a structured representation of the text’s meaning called an *Action Specification* (ActSpec). This ActSpec, which can describe a command, query, or assertion about the world, is communicated as a JSON object to the action side of the system. The action side, depicted on the right in Figure 4, consists of a *Problem Solver*, a *world model*, and a *Task API*. The Problem Solver unpacks the ActSpec to determine what task is to be accomplished and with what parameters. The Problem Solver refers to the world model to determine the constraints on the possible solutions to the problems posed by the target task, and computes a solution. Finally, the Problem Solver makes API calls to the underlying application, which constitute the actions necessary to accomplish the target task. Communication between the language and action sides is enabled by shared ActSpec templates, which guide ActSpec creation, and a shared ontology. The separation of the language and action sides allows them to be run on separate processes; thus, the action-side can solve ActSpecs while the language-side continues to receive text and speech input.

A crucial theoretical underpinning of the system design is previous work on mental simulation of language (Bergen 2012). The language side produces the parameters for the simulation of an event, which can then be simulated or executed by the action side.

More concretely, the ECG Analyzer (Bryant 2008), a cognitively plausible language parser, uses an *Embodied Construction Grammar* (Bergen, Chang, and Narayan 2004; Feldman, Bryant, and Dodge 2009) to produce a *Semantic Specification* (SemSpec) of the events described in language. The SemSpec maps constructional information to the meaning of the sentence. Task-relevant information from the SemSpec is extracted by the Specializer and formatted into an ActSpec. This ActSpec provides the parameters for the simulation of language.

In the system described in (Khayrallah, Trott, and Feldman 2015; Trott et al. 2015), the action side executed these incoming ActSpecs as they arrived. In the current work, X-nets are introduced to provide additional mechanisms for simulation and control.

### Integration of X-nets

The architecture depicted in Figure 4 has been modified to have the Problem Solver execute one or more X-nets to represent the action being undertaken through the task API. The Problem Solver examines the state of the X-net either during or after execution to make inferences about progress on the task. Integrating X-nets into the architecture enables the system to deal with the following types of tasks, events, or requests:

- suspension: “Robot1, stop moving!”
- resumption: “Robot1, continue moving!”
- interruption and redirection: while a move is in progress, e.g. “Robot1, move to the green box!”, a new target is specified: “Robot1, move to the blue box!”

For the application of controlling a simulated robot, the Problem Solver invokes the Move X-net (Figure 5). The Move X-net supports an additional parameterization, implementing different speeds for different verbs: amble (slow), move (normal), dash (fast).

### X-net Problem Solver

The incorporation of X-nets into the LCAS architecture allows a systematic decoupling of the language side and action side of the architecture, such that the user can interact without restriction with the system, often in response to the current system behavior. The user can interrupt, cancel, resume, or restart the action, as appropriate, without being constrained by blocking or synchronous control flows. As is seen in both biological and mechanical action loops, there is some inherent delay between a new request and the action response, but the asynchronous nature of Petri net execution enables this to be handled in a principled way, without resort to complex logic or ad hoc exception handling.

To support this decoupling, the execution of an X-net Problem Solver involves two threads, one that listens on a queue for new incoming requests from the language side, and a second thread that interacts with the particular X-net appropriate to the task at hand. In turn, the X-net interacts with the task API to accomplish the objective. For an ongoing action, a new incoming request will be translated by the Problem Solver into an update to one or more places in the

currently executing X-net, which in turn will cause a change in the X-net flow of control, corresponding to the meaning of the new request. Requests regarding actions that have not yet started or have completed may result in the invocation of a new X-net, or no change.

### Petri net Extensions to Support External Systems

X-net support is built on PIPE V5, an open source Java Petri net editor and debugging environment<sup>1</sup> (Bonet et al. 2007; Dingle, Knottenbelt, and Suto 2009; Tattersall and Knottenbelt 2014). To support the current work, PIPE was extended to provide support for interfaces to and from external systems, and for execution through an API<sup>2</sup>. The extensions include:

- **External Transition:** an extension to a Petri net transition, providing a mechanism to execute Java code whenever the transition fires. The External Transition has access to the state of the Petri net, and is optionally given a context passed from the external system, consisting of an instance of an arbitrary Java class.
- **External Input Place:** an extension to a Petri net place, providing a mechanism for an external system to update the marking of the place with one or more tokens.
- **External Output Place:** an extension to a Petri net place, providing a mechanism for an external system to be notified whenever the marking of the place changes.
- **PNML extensions:** extensions to the Petri Net Markup Language, which defines the specification of a Petri net in XML format, to support external transitions and places.
- **Merge Place:** an extension to a Petri net place, enabling a Petri net to be composed of multiple other Petri nets for purposes of design and coding. Merge places serve to connect the various Petri nets together, and are then collapsed to single places at execution time, creating a single executable Petri net.
- **Runner:** the Runner interface enables an external system to interact with a Petri net. The interface supports the following functions:
  - load a Petri net from a PNML file
  - start Petri net execution
  - mark a place in a Petri net
  - subscribe to notifications of a change in marking of a place in a Petri net
  - pass an object to be made available to a Petri net transition when it fires
  - subscribe to notifications of global Petri net events, including starting or stopping execution, and firing a transition

<sup>1</sup><https://github.com/sarahtattersall/PIPE>

<sup>2</sup>Current code: <https://github.com/sjdayday/PIPECore/tree/hierarchical-nets>. Documentation: <https://github.com/sjdayday/PIPECore/wiki>. In the future, the extensions will be integrated as a release of PIPE.

### Standard Action Semantics and the Move X-net

Figure 5 shows the Move X-net, which controls the motion of a simulated robot. The upper portion of the X-net is an implementation of a subset of the standard action semantics depicted in Figure 3. Updating or reading places in this X-net provides a simple, standard interface to many of the operations that we assume should be available in any action – starting, stopping, resuming or restarting an action.

The lower portion of the Move X-net depicts the logic specific to controlling motion. This is a simplified interface that defers specifics of trajectory planning to the Problem Solver and to the simulator (Morse, in the current example), but additional logic could be added as needed for other domains, without affecting the operation of the standard action semantics in the upper portion of the X-net.

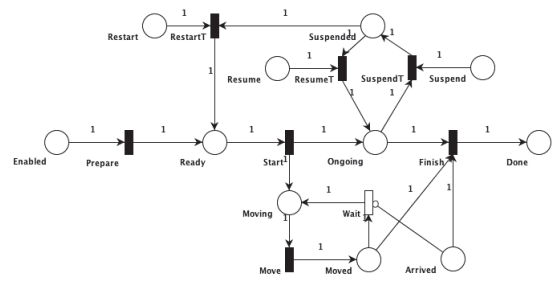


Figure 5: Move X-net

### X-net Interface to an External System

The mechanics of the interface between the LCAS system and the Petri net implementation of the Move X-net is depicted in Figure 6. This serves as an example of the general interaction between an X-net and any external system.

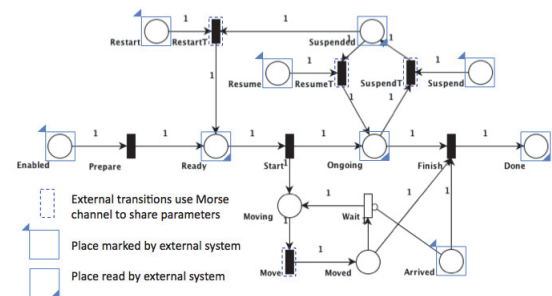


Figure 6: Move X-net interface to an external system

From the perspective of the X-net, the external system here is the Problem Solver component of the LCAS system, which controls execution of the motion by populating some subset of the standard external input places: Enabled, Suspend, Resume and Restart. The Problem Solver monitors progress of the motion by listening for updates to standard external output places: Ready, Ongoing, Done, and

Suspended. The interaction with the task API, in this case the Morse simulator, happens through bi-directional updates to an instance of the task-specific MotionChannel Java class. These updates happen during the external transitions: Move, SuspendT, ResumeT, and RestartT. Each transition communicates an appropriate target operation, as well as the target position of the simulated robot, from the perspective of the X-net. When any transition fires, control passes from the X-net back to the Problem Solver, which performs any appropriate pending action and updates the current position of the robot from the perspective of the Morse simulator, and then returns control back to the X-net. Finally, when the Morse simulator detects that the simulated robot has arrived at the destination, it updates the Arrived external output place to trigger transition of the Move X-net to the Done state.

The mechanics of the integration between the Python implementation of LCAS and the Java implementation of the Petri net is handled through the JPy Python package. JPy enables a Python script to start and attach a new JVM to the current Python thread, and then to create instances of Java classes. The methods and public fields of the Java instances are then available for inspection and update by the script through syntax that is consistent with Python. The first Problem Solver script creates the JVM, and any subsequent Problem Solver instances use the same JVM. Each Problem Solver then creates an X-net as needed for its current task, using the Runner API, and interacts with that X-net until its task is completed.

The interaction between the Problem Solver and the X-net can be configured to be effectively synchronous, such that every state change in the X-net enables the Problem Solver to regain control. Depending on the task, however, the Problem Solver may not need such granular control, and may only register for asynchronous notifications. Regardless of the degree of coupling between the Problem Solver and the X-net, this does not alter the overall asynchronous relationship between the language side and the action side in the LCAS architecture. The processing of requests from the end user proceeds asynchronously from the combined activity of the Problem Solver, X-net, and supporting task API. The standard action semantics provide natural support for the interruption of ongoing action, the preemption of pending action, or no operation when the action is already complete.

## Robot Motion Example

### Normal Motion

The integration of the Problem Solver, Move X-net, and the Morse simulator to accomplish motion of the simulated robot, has been demonstrated in the context of various language commands<sup>3</sup>. In response to the command “Robot1, move to the blue box!”, the language side generates an Act-Spec to the action side, adding it to the queue of requests to be processed by the X-net Morse Problem Solver. The Problem Solver listens for each firing event in an executing X-net, which gives it an opportunity to check the queue for any new incoming requests. This ensures that latency involved in

processing a new request is short. In this case, the incoming request causes a new instance of a PetriNetRunner to be created, which implements the Runner API. Next, an instance of the Move X-net is instantiated through the Runner API, and the Problem Solver subscribes for updates to the various external output places of interest in the X-net, primarily those defined in the standard action semantics. A new instance of a MotionChannel is created, to act as a shared context for all of the External Transitions in the Move X-net, and as the communication channel between the X-net and Morse. The Problem Solver updates the MotionChannel with the location of the blue box, and the default speed. The Problem Solver requests that the Enabled place in the X-net be updated with a single token, and then requests that the X-net be run. The PetriNetRunner starts the execution of the Petri net, and notifies the listening Problem Solver of every event to which it has subscribed. As noted earlier, these notifications include the firing of every transition, so that any new incoming language requests can be processed promptly. Flow of control passes through the Ready place and then through the Start transition, causing both the Ongoing and Moving places to be populated, as depicted in Figure 7.

When the Move External Transition fires, the MotionChannel is updated with the target operation of “move”. The target location having been previously populated by the Problem Solver, control passes back to the Problem Solver, which passes the MotionChannel information to the Morse simulator, which begins the simulated robot motion. Control then returns to the X-net, which loops through the Moved place and the Wait transition, before returning to Moving.

The Wait transition populates the Moving place. The transition is designated as a timed transition, to wait a configurable amount of time before firing. Currently, however, the transition fires immediately, pending addition of the necessary support to the underlying PIPE implementation. Implementation as a timed transition is intended as a performance optimization, to prevent the Moving / Moved loop from consuming system resources unproductively, and as a better model of the delays involved in physical motion.

Each time the MotionChannel information is passed to Morse, Morse updates the current location of the robot; this information is currently unused by the X-net, but would be available if more advanced functions were required. Finally, when the simulated robot reaches the target location, the MotionChannel is updated, and the Arrived place is marked, driving the X-net through the Finish transition to the Done place. The Problem Solver is notified that the X-net is Done, indicating normal completion of the motion.

### Processing Other Requests

The power of the standard action semantics is seen when the system processes other incoming requests. Requests of very different syntactic form and meaning are handled with little additional complexity. “Robot1, stop moving!” causes the Problem Solver to populate the Suspend place. If the X-net action is currently Ongoing, this causes the SuspendT External Transition to fire. This transition updates the MotionChannel target operation to be “suspend”, which in turn causes Morse to stop the motion of the robot simulator. The

<sup>3</sup>Demonstration video: <https://youtu.be/8PwHpng3Nj8>






Figure 7: Motion in progress




Figure 8: Motion to be interrupted and re-directed

transition’s firing also updates the state of the X-net to Suspended, causing the Problem Solver to be notified that motion has stopped.

In response to the request “Robot1, continue moving!”, the Problem Solver populates the Resume place. If the current state of the X-net is Suspended, this drives the ResumeT External Transition, which updates the MotionChannel target operation to “resume”. This causes Morse to start the simulated robot moving again, resuming its trajectory toward the current goal. Race conditions are handled simply, with no additional logic. If the X-net is not in a Suspended state, the update of the Resume place has no impact on processing; it is effectively ignored.

Finally, the more complicated case of interruption and re-direction is also handled simply. The user might wish to override a request that is in progress, e.g., “Robot1, move to the blue box!” The new request might be “Robot1, dash to the green box!” This new request causes the Problem Solver to calculate a trajectory towards the green box, and to update the MotionChannel both with the location of the green box, and with a higher speed value. As well, both the Suspend and Restart places are marked as in Figure 8. X-net execution first transitions to the Suspended state, as in the “stop moving!” example, causing Morse to stop motion of the simulated robot. The population of both the Restart and the Suspended places then drives the firing of the RestartT External Transition, which updates the MotionChannel target operation to “restart”. This causes Morse to start the motion of the simulated robot again, following the new trajectory and speed. The X-net flow of control then passes back through the Ready place and the Start transition, re-entering the Ongoing state.

### Non-Linguistic Events

In addition to linguistic requests, a robot frequently encounters other events as it solves various tasks. Many of these events involve changes to the environment, such as a moving object suddenly crossing paths with the robot’s trajectory, or encountering a previously unknown obstacle. Environmental interrupts are challenging because the robot must dynamically adjust its world model to reflect these changes, and then develop a new plan to continue solving the original task.

We implemented a general protocol for how the Problem

Solver should behave when experiencing world-based interrupts, and extended this to the Morse domain. In the Morse demo (Trott et al. 2015), we addressed the problem of encountering objects that were not previously in the Problem Solver’s world model. Novel objects could potentially disrupt the Problem Solver’s original plan for solving a task; for example, if the robot is asked to push a box north, it might find that a previously unknown box is in the way of its original path.

On the Morse side, our solution involved attaching a simulated proximity sensor<sup>4</sup> to the instance of the robot model. Whenever the robot passes within a certain threshold of distance from an object, the proximity sensor relays information back to the Problem Solver about the object’s location, color, and size. If the Problem Solver already knows about the object, the world model is simply verified to make sure the object location is correct. If the object is unknown, however, the Problem Solver performs three functions, in order:

1. The Problem Solver updates its world model with the new information.
2. The Problem Solver sends an ActSpec back to the UI-Agent, to notify the human user that a new object has been discovered.
3. The Problem Solver develops a new plan for the original task, now taking into account the updated world model.

Previous work also addressed the problem of disparate world models among multiple agents (Trott et al. 2015); in that case, (2) also involves sending notification ActSpecs to the other robots, so that they can also update their world models.

### Conclusion

We have demonstrated the ability to handle language requests regarding the current state of actions in a graceful fashion, enabling actions to be interrupted, resumed and overridden. This is a useful advance in the natural language control of autonomous systems. The separation of language inputs from the control of actions enables both to proceed asynchronously. This enables the overall system to be responsive to new requests, while maintaining control of ac-

<sup>4</sup><http://www.openrobots.org/morse/doc/1.2/user/sensors/proximity.html>

tions at an appropriate level of granularity. Although demonstrated in the limited domain of movement of a simulated robot, the approach can be extended to other domains in a straightforward manner.

## Limitations

There are several limitations of the system. The system does not address the problem of perceiving and interpreting the actions of other agents (Summers-Stay et al. 2012; Yang et al. 2015). From a technical perspective, support in PIPE for external interfaces is limited to the API described earlier; interaction through a graphical interface is not yet supported, but should be available in the future. The standard semantics of X-nets invites some programmatic support; this has been discussed<sup>5</sup> but work has not yet begun. Finally, no effort has been made to optimize the system for performance and scale.

## Future Work

The system supports current action. A next step is to have the system support simulated action, where the execution of X-nets can be used to reason about the consequences of possible future actions (Sinha 2008). This would enable the evaluation of counterfactual or hypothetical statements, such as “Robot1, if you moved north of the green box, could you push it south?”

Future work should also examine the problem of integrating our work on non-linguistic events with the mechanism for X-net control.

## Acknowledgements

We thank Luca Gilardi for his help in designing the system framework, and for the creation of the GUI for the ECG Workbench Editor. Professor William Knottenbelt and his students at Imperial College, London, created and have maintained PIPE through many releases. Sarah Tattersall performed a thorough refactoring of the PIPE code base, without which the external system extensions would not have been possible.

## References

Bergen, B.; Chang, N.; and Narayan, S. 2004. Simulated action in an embodied construction grammar. In *Proceedings of the Twenty-Sixth Annual Conference of the Cognitive Science Society*, 108–113.

Bergen, B. K. 2012. *Louder than words: The new science of how the mind makes meaning*. Basic Books.

Bonet, P.; Lladó, C. M.; Puijaner, R.; and Knottenbelt, W. J. 2007. PIPE v2. 5: A Petri net tool for performance modelling. In *Proc. 23rd Latin American Conference on Informatics (CLEI 2007)*.

Bryant, J. E. 2008. *Best-Fit Constructional Analysis*. Ph.D. Dissertation, University of California at Berkeley.

Chang, N., and Gurevich, O. 2004. Context-driven construction learning. In *Proceedings of the 26th Annual Meeting of the Cognitive Science Society*, 204–209.

<sup>5</sup><https://github.com/sjdayday/xschema/wiki>

Chang, N. 2008. *Constructing Grammar: A Computational Model of the Emergence of Early Constructions*. University of California, Berkeley.

Dingle, N. J.; Knottenbelt, W. J.; and Suto, T. 2009. PIPE2: a tool for the performance evaluation of generalised stochastic Petri Nets. *ACM SIGMETRICS Performance Evaluation Review* 36(4):34–39.

Doubleday, S. 2017. *A Computational Model of Spatial Navigation During Simulated Motion in Rats*. Ph.D. Dissertation, University of California at Irvine.

Echeverria, G.; Lassabe, N.; Degroote, A.; and Lemaignan, S. 2011. Modular Open Robots Simulation Engine: MORSE. In *2011 IEEE International Conference Robotics and Automation*, 46–51.

Eppe, M.; Trott, S.; Raghuram, V.; Feldman, J.; and Janin, A. 2016. Application-Independent and Integration-Friendly Natural Language Understanding. *Forthcoming (under review)*.

Eppe, M.; Trott, S.; and Feldman, J. 2016. Exploiting Deep Semantics and Compositionality of Natural Language for Human-Robot-Interaction. In *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems in Daejeon, Korea*.

Feldman, J.; Bryant, J. E.; and Dodge, E. 2009. A Neural Theory of Language and Embodied Construction Grammar. In *The Oxford Handbook of Computational Linguistics*. Oxford University Press. 38 – 111.

Guha, A.; Yang, Y.; Fermu, C.; Aloimonos, Y.; et al. 2013. Minimalist plans for interpreting manipulation actions. In *2013 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 5908–5914. IEEE.

Hesslow, G. 2002. Conscious thought as simulation of behaviour and perception. *Trends in cognitive sciences* 6(6):242–247.

Khayrallah, H.; Trott, S.; and Feldman, J. 2015. Natural Language For Human Robot Interaction. In *International Conference on Human-Robot Interaction (HRI)*.

Murata, T. 1989. Petri nets: Properties, analysis and applications. In *Proceedings of the IEEE 77.4*, volume 77, 541–580. IEEE.

Narayanan, S. 1997. *Knowledge-based Action representations for metaphor and aspect (KARMA)*. Ph.D. Dissertation, University of California at Berkeley.

Narayanan, S. 2010. Mind changes: A simulation semantics account of counterfactuals. In *Cognitive Science*. Working Paper, University of California, Berkeley.

Pastra, K., and Aloimonos, Y. 2012. The minimalist grammar of action. *Philosophical Transactions of the Royal Society of London B: Biological Sciences* 367(1585):103–117.

Schilling, M., and Narayanan, S. 2013. Communicating with Executable Action Representations. In *AAAI Spring Symposium: Designing Intelligent Robots*.

Sinha, S. 2008. *Answering Questions About Complex Events*. Ph.D. Dissertation, University of California at Berkeley.

Summers-Stay, D.; Teo, C. L.; Yang, Y.; Fermüller, C.; and Aloimonos, Y. 2012. Using a minimal action grammar for activity understanding in the real world. In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 4104–4111. IEEE.

Tattersall, S., and Knottenbelt, W. 2014. PIPE—The Great Re-Plumbing. Master’s thesis, Imperial College London.

Trott, S.; Appriou, A.; Feldman, J.; and Janin, A. 2015. Natural Language Understanding and Communication for Multi-Agent Systems. In *AAAI Fall Symposium*, 137–141.

Yang, Y.; Li, Y.; Fermüller, C.; and Aloimonos, Y. 2015. Robot learning manipulation action plans by “watching” unconstrained videos from the world wide web. In *AAAI*, 3686–3693.