

Technical Implementation of Semantic Control in Eduba: An In-Depth Exploration

D. McDermott, J.E. Van Clief

September 22, 2024

Abstract

This sister paper provides a detailed technical exploration of the Eduba platform’s implementation, focusing on the integration of Large Language Models (LLMs) with Microsoft’s Semantic Kernel. Building upon the concepts presented in the previous paper, we delve into the specific code examples and architectural designs that enable the autonomous refinement of code and documents across various domains. We present practical implementations of Semantic Kernel plugins, elaborate on their functionalities, and demonstrate how they collaboratively support the system’s objectives. Challenges encountered during development are discussed, along with the solutions employed to overcome them.

Contents

1	Introduction	2
2	Overview of Eduba’s Architecture	2
3	Semantic Kernel Integration	2
3.1	Plugin Architecture	2
4	Detailed Code Examples	2
4.1	Code Editing Plugin	2
4.2	Build Execution Plugin	3
4.3	Test Execution Plugin	4
4.4	Merge Conflict Resolution Plugin	5
4.5	Documentation Synchronization Plugin	6
5	Implementation Details	7
5.1	Orchestrating Plugin Interactions	7
5.2	Handling Errors and Exceptions	7
6	Challenges and Solutions	7
6.1	Limitations of LLMs	7
6.2	Ensuring Code Quality and Security	7
6.3	Resource Management	8
7	Extending to Document Management	8
7.1	Document Editing Plugin	8
7.2	Knowledge Synthesis Plugin	9
8	Security and Ethical Considerations	10
8.1	Access Control	10
8.2	Data Privacy	10
8.3	Ethical AI Use	10

1 Introduction

The Eduba platform aims to revolutionize knowledge management by integrating AI-driven automation into collaborative environments. Central to this goal is the ability of the system to autonomously refine code and documents, ensuring consistency and alignment across multiple domains. This paper provides an in-depth technical examination of how Semantic Kernel facilitates this integration, with comprehensive code examples and explanations of the underlying mechanisms.

2 Overview of Eduba's Architecture

Eduba's architecture is designed to support seamless collaboration between human users and AI agents. The core components include:

- **LLM Integration:** Utilizing advanced language models to interpret and generate content.
- **Semantic Kernel Plugins:** Modular components that provide specific functionalities to the LLM.
- **Autonomous CI/CD Pipeline:** An automated pipeline that continuously integrates and deploys code changes.
- **Knowledge Management System:** A platform for managing documents, code, and other artifacts collaboratively.

The integration of these components allows Eduba to autonomously manage and refine knowledge assets while facilitating human-AI collaboration.

3 Semantic Kernel Integration

Semantic Kernel serves as the bridge between the LLM and the actionable components of the system. By developing custom plugins, we enable the LLM to perform tasks such as code editing, build execution, test automation, and documentation synchronization.

3.1 Plugin Architecture

Each plugin is implemented as a class with methods decorated by the `[SKFunction]` attribute, indicating that they are callable by the LLM. The plugins interact with the system's resources and provide the LLM with the necessary tools to execute actions.

4 Detailed Code Examples

In this section, we present code examples of key plugins, illustrating how they are implemented and how they contribute to the system's functionality.

4.1 Code Editing Plugin

The Code Editing Plugin allows the LLM to read, modify, and write code files.

Listing 1: Code Editing Plugin Implementation

```

1 using Microsoft.SemanticKernel;
2 using Microsoft.SemanticKernel.Orchestration;
3 using Microsoft.SemanticKernel.SkillDefinition;
4 using System.IO;
5

```

```

6 public class CodeEditingPlugin
7 {
8     [SKFunction("Modify code files based on instructions.")]
9     public async Task<SKContext> EditCodeAsync(
10         [SKName("filePath"), SKDescription("The path to the code file to edit.")] string filePath,
11         [SKName("instructions"), SKDescription("Instructions for modifying the code.")] string instructions,
12         SKContext context)
13     {
14         // Ensure the file exists
15         if (!File.Exists(filePath))
16         {
17             context.Fail($"File not found: {filePath}");
18             return context;
19         }
20
21         // Read the existing code
22         string code = await File.ReadAllTextAsync(filePath);
23
24         // Create a prompt for the LLM
25         var prompt = $"
26 You are a skilled C# developer. Modify the following code according to the instructions provided.
27
28 Instructions:
29 {instructions}
30
31 Code:
32 {code}
33
34 Modified Code:";
35
36         // Invoke the LLM to get the modified code
37         var result = await context.SemanticKernel.CompleteAsync(prompt);
38
39         // Write the modified code back to the file
40         await File.WriteAllTextAsync(filePath, result);
41
42         context.Variables["Status"] = "Code modified successfully.";
43         return context;
44     }
45 }

```

Explanation:

- The plugin defines an asynchronous function `EditCodeAsync` that takes a file path and modification instructions.
- It reads the existing code from the file and constructs a prompt for the LLM.
- The LLM generates the modified code, which is then written back to the file.
- Context variables are updated to indicate the operation's success.

4.2 Build Execution Plugin

The Build Execution Plugin allows the LLM to initiate a build process and retrieve any errors.

Listing 2: Build Execution Plugin Implementation

```

1 using Microsoft.SemanticKernel;
2 using Microsoft.SemanticKernel.Orchestration;
3 using Microsoft.SemanticKernel.SkillDefinition;
4 using System.Diagnostics;
5
6 public class BuildExecutionPlugin
7 {
8     [SKFunction("Execute the build process and return errors.")]
9     public async Task<SKContext> ExecuteBuildAsync(

```

```

10     [SKName("solutionPath"), SKDescription("The path to the solution file.")] string solutionPath,
11     SKContext context)
12 {
13     // Configure the build process
14     var startInfo = new ProcessStartInfo
15     {
16         FileName = "dotnet",
17         Arguments = $"build \"{solutionPath}\"",
18         RedirectStandardOutput = true,
19         RedirectStandardError = true,
20         UseShellExecute = false,
21         CreateNoWindow = true
22     };
23
24     // Start the build process
25     using var process = new Process { StartInfo = startInfo };
26     process.Start();
27
28     // Read output and errors
29     string output = await process.StandardOutput.ReadToEndAsync();
30     string errors = await process.StandardError.ReadToEndAsync();
31
32     process.WaitForExit();
33
34     // Update context with build results
35     context.Variables["BuildOutput"] = output;
36     context.Variables["BuildErrors"] = errors;
37
38     return context;
39 }
40 }

```

Explanation:

- The plugin runs the `dotnet build` command on the specified solution.
- Standard output and error streams are captured.
- Build results are stored in context variables for the LLM to analyze.

4.3 Test Execution Plugin

The Test Execution Plugin enables the LLM to run tests and analyze the results.

Listing 3: Test Execution Plugin Implementation

```

1 using Microsoft.SemanticKernel;
2 using Microsoft.SemanticKernel.Orchestration;
3 using Microsoft.SemanticKernel.SkillDefinition;
4 using System.Diagnostics;
5
6 public class TestExecutionPlugin
7 {
8     [SKFunction("Run tests and return results.")]
9     public async Task<SKContext> ExecuteTestsAsync(
10         [SKName("testProjectPath"), SKDescription("The path to the test project.")] string testProjectPath,
11         SKContext context)
12     {
13         // Configure the test process
14         var startInfo = new ProcessStartInfo
15         {
16             FileName = "dotnet",
17             Arguments = $"test \"{testProjectPath}\"",
18             RedirectStandardOutput = true,
19             RedirectStandardError = true,
20             UseShellExecute = false,
21             CreateNoWindow = true

```

```

22     };
23
24     // Start the test process
25     using var process = new Process { StartInfo = startInfo };
26     process.Start();
27
28     // Read output and errors
29     string output = await process.StandardOutput.ReadToEndAsync();
30     string errors = await process.StandardError.ReadToEndAsync();
31
32     process.WaitForExit();
33
34     // Update context with test results
35     context.Variables["TestOutput"] = output;
36     context.Variables["TestErrors"] = errors;
37
38     return context;
39 }
40 }

```

Explanation:

- The plugin runs the `dotnet test` command on the specified project.
- Captures the test output and any errors.
- Results are stored for further analysis by the LLM.

4.4 Merge Conflict Resolution Plugin

This plugin allows the LLM to resolve merge conflicts during code integration.

Listing 4: Merge Conflict Resolution Plugin Implementation

```

1 using Microsoft.SemanticKernel;
2 using Microsoft.SemanticKernel.Orchestration;
3 using Microsoft.SemanticKernel.SkillDefinition;
4 using System.IO;
5
6 public class MergeConflictResolutionPlugin
7 {
8     [SKFunction("Resolve merge conflicts in a code file.")]
9     public async Task<SKContext> ResolveConflictsAsync(
10         [SKName("filePath"), SKDescription("The path to the conflicted code file.")] string filePath,
11         SKContext context)
12     {
13         // Read the conflicting code
14         string code = await File.ReadAllTextAsync(filePath);
15
16         // Check for merge conflict markers
17         if (!code.Contains("<<<<<<") || !code.Contains(">>>>>>"))
18         {
19             context.Fail("No merge conflicts found in the file.");
20             return context;
21         }
22
23         // Create a prompt for the LLM
24         var prompt = $"
25 You are an expert in resolving code merge conflicts. The following code contains merge conflicts. Resolve
26 Conflicted Code:
27 {code}
28 Resolved Code:";
29
30         // Invoke the LLM to resolve the conflicts
31         var result = await context.SemanticKernel.CompleteAsync(prompt);

```

```

34
35     // Write the resolved code back to the file
36     await File.WriteAllTextAsync(filePath, result);
37
38     context.Variables["Status"] = "Merge conflicts resolved successfully.";
39     return context;
40 }
41 }

```

Explanation:

- The plugin detects merge conflict markers in the code file.
- Constructs a prompt instructing the LLM to resolve the conflicts.
- The LLM generates the resolved code, which replaces the conflicted file.

4.5 Documentation Synchronization Plugin

This plugin ensures that the documentation stays consistent with the codebase.

Listing 5: Documentation Synchronization Plugin Implementation

```

1 using Microsoft.SemanticKernel;
2 using Microsoft.SemanticKernel.Orchestration;
3 using Microsoft.SemanticKernel.SkillDefinition;
4 using System.IO;
5
6 public class DocumentationSyncPlugin
7 {
8     [SKFunction("Update documentation to align with code changes.")]
9     public async Task<SKContext> SyncDocumentationAsync(
10         [SKName("codeFilePath"), SKDescription("The path to the code file.")] string codeFilePath,
11         [SKName("docFilePath"), SKDescription("The path to the documentation file.")] string docFilePath,
12         SKContext context)
13     {
14         // Read the code and documentation
15         string code = await File.ReadAllTextAsync(codeFilePath);
16         string documentation = await File.ReadAllTextAsync(docFilePath);
17
18         // Create a prompt for the LLM
19         var prompt = $"
20 You are a documentation specialist. Update the following documentation to reflect the current state of the
21 Code:
22 {code}
23 Current Documentation:
24 {documentation}
25 Updated Documentation:";
26
27         // Invoke the LLM to update the documentation
28         var result = await context.SemanticKernel.CompleteAsync(prompt);
29
30         // Write the updated documentation back to the file
31         await File.WriteAllTextAsync(docFilePath, result);
32
33         context.Variables["Status"] = "Documentation synchronized successfully.";
34         return context;
35     }
36 }
37 }
38 }
39 }

```

Explanation:

- The plugin reads both the code and existing documentation.
- Uses the LLM to generate updated documentation that reflects code changes.

- Updates the documentation file with the new content.

5 Implementation Details

5.1 Orchestrating Plugin Interactions

The plugins work collaboratively to achieve the system's objectives. The orchestration involves:

- **Sequential Execution:** Plugins are executed in a specific order based on dependencies (e.g., code edits before builds).
- **Context Sharing:** Plugins share context variables, allowing them to access results from previous steps.
- **Error Handling:** If a plugin encounters an error, the system can decide to retry, skip, or escalate the issue.

5.2 Handling Errors and Exceptions

Robust error handling is crucial for an autonomous system. Strategies include:

- **Validation Checks:** Before executing actions, plugins validate inputs and preconditions.
- **Retries:** Certain operations may be retried with modified parameters if they fail initially.
- **Fallback Mechanisms:** If the LLM fails to produce a valid output, the system can revert to default behaviors or notify human operators.

6 Challenges and Solutions

6.1 Limitations of LLMs

Challenge:

LLMs may generate incorrect or suboptimal code, particularly in complex scenarios.

Solution:

- **Prompt Engineering:** Crafting precise and unambiguous prompts to guide the LLM effectively.
- **Post-Generation Validation:** Implementing automated checks to verify the correctness of the LLM's outputs.
- **Human-in-the-Loop:** Allowing human reviewers to oversee critical changes, especially in production environments.

6.2 Ensuring Code Quality and Security

Challenge:

Automatically generated code may not adhere to best practices or may introduce security vulnerabilities.

Solution:

- **Static Analysis Tools:** Integrating tools like SonarQube or Roslyn analyzers to assess code quality.
- **Security Scanners:** Employing security-focused analysis to detect potential vulnerabilities.
- **Coding Standards Enforcement:** Defining and enforcing coding standards through automated checks.

6.3 Resource Management

Challenge:

Running multiple processes (builds, tests, LLM computations) can strain system resources.

Solution:

- **Asynchronous Execution:** Utilizing async programming patterns to optimize resource usage.
- **Task Scheduling:** Prioritizing tasks based on urgency and resource availability.
- **Scalable Infrastructure:** Leveraging cloud resources to scale computational power as needed.

7 Extending to Document Management

While the examples above focus on code, similar principles apply to document management.

7.1 Document Editing Plugin

A plugin for editing documents based on collaborative inputs.

Listing 6: Document Editing Plugin Implementation

```
1 using Microsoft.SemanticKernel;
2 using Microsoft.SemanticKernel.Orchestration;
3 using Microsoft.SemanticKernel.SkillDefinition;
4 using System.IO;
5
6 public class DocumentEditingPlugin
7 {
8     [SKFunction("Edit a document based on provided instructions.")]
9     public async Task<SKContext> EditDocumentAsync(
10         [SKName("docFilePath"), SKDescription("The path to the document file.")] string docFilePath,
11         [SKName("instructions"), SKDescription("Instructions for editing the document.")] string instructions,
12         SKContext context)
13     {
14         // Ensure the file exists
15         if (!File.Exists(docFilePath))
16         {
17             context.Fail($"Document not found: {docFilePath}");
18             return context;
19         }
20
21         // Read the existing document
22         string document = await File.ReadAllTextAsync(docFilePath);
23
24         // Create a prompt for the LLM
25         var prompt = $"
26 You are an expert editor. Modify the following document according to the instructions provided.
27
28 Instructions:
29 {instructions}
30
31 Document:
32 {document}
33
34 Edited Document:";
35
36         // Invoke the LLM to edit the document
37         var result = await context.SemanticKernel.CompleteAsync(prompt);
38
39         // Write the edited document back to the file
40         await File.WriteAllTextAsync(docFilePath, result);
41
42         context.Variables["Status"] = "Document edited successfully.";
43         return context;
```



```

44     }
45 }

```

Explanation:

- Similar to the Code Editing Plugin but tailored for document content.
- Useful for incorporating contributions from multiple collaborators.

7.2 Knowledge Synthesis Plugin

A plugin that synthesizes information from multiple documents.

Listing 7: Knowledge Synthesis Plugin Implementation

```

1 using Microsoft.SemanticKernel;
2 using Microsoft.SemanticKernel.Orchestration;
3 using Microsoft.SemanticKernel.SkillDefinition;
4 using System.Collections.Generic;
5 using System.IO;
6
7 public class KnowledgeSynthesisPlugin
8 {
9     [SKFunction("Synthesize information from multiple documents.")]
10    public async Task<SKContext> SynthesizeKnowledgeAsync(
11        [SKName("docFilePaths"), SKDescription("List of document file paths.")] List<string> docFilePaths,
12        SKContext context)
13    {
14        // Read all documents
15        var documents = new List<string>();
16        foreach (var path in docFilePaths)
17        {
18            if (File.Exists(path))
19            {
20                documents.Add(await File.ReadAllTextAsync(path));
21            }
22            else
23            {
24                context.Fail($"Document not found: {path}");
25                return context;
26            }
27        }
28
29        // Create a prompt for the LLM
30        var prompt = $"
31 You are an expert researcher. Synthesize the key information from the following documents into a cohesive
32
33 Documents:
34 {string.Join("\n\n--\n\n", documents)}
35
36 Summary:";
37
38        // Invoke the LLM to generate the synthesis
39        var result = await context.SemanticKernel.CompleteAsync(prompt);
40
41        // Store the summary in context
42        context.Variables["Summary"] = result;
43
44        return context;
45    }
46 }

```

Explanation:

- Reads multiple documents and constructs a prompt for synthesis.
- The LLM produces a summary that combines the key points.
- Useful for knowledge management and ensuring consistency across documents.

8 Security and Ethical Considerations

8.1 Access Control

Implementation:

- Plugins check user permissions before performing actions.
- Sensitive operations require elevated privileges.

8.2 Data Privacy

Implementation:

- Sensitive data is encrypted at rest and in transit.
- The LLM is restricted from accessing confidential information unless necessary.

8.3 Ethical AI Use

Considerations:

- Ensuring the LLM does not generate biased or inappropriate content.
- Implementing oversight mechanisms to monitor and correct undesirable outputs.

9 Conclusion

The Eduba platform's integration of LLMs with Semantic Kernel enables a powerful, autonomous system capable of refining code and documents across various domains. By implementing specialized plugins and robust orchestration mechanisms, the system can perform complex tasks that traditionally require human intervention. While challenges exist, thoughtful design and continuous improvement ensure that the system operates effectively and ethically.

References

- Microsoft Semantic Kernel: <https://github.com/microsoft/semantic-kernel>
- OpenAI GPT-4: <https://openai.com/product/gpt-4>
- Eduba Project Documentation: <https://eduba.example.com/docs>
- C# Documentation: <https://docs.microsoft.com/en-us/dotnet/csharp/>
- .NET Command-Line Tools: <https://docs.microsoft.com/en-us/dotnet/core/tools/>