

# Automatic Design of Semantic Similarity Ensembles Using Grammatical Evolution

Jorge Martinez-Gil

*Software Competence Center Hagenberg GmbH  
Softwarepark 32a, 4232 Hagenberg, Austria  
jorge.martinez-gil@scch.at*

---

## Abstract

Semantic similarity measures are widely used in natural language processing to catalyze various computer-related tasks. However, no single semantic similarity measure is the most appropriate for all tasks, and researchers often use ensemble strategies to ensure performance. This research work proposes a method for automatically designing semantic similarity ensembles. In fact, our proposed method uses grammatical evolution, for the first time, to automatically select and aggregate measures from a pool of candidates to create an ensemble that maximizes correlation to human judgment. The method is evaluated on several benchmark datasets and compared to state-of-the-art ensembles, showing that it can significantly improve similarity assessment accuracy and outperform existing methods in some cases. As a result, our research demonstrates the potential of using grammatical evolution to automatically compare text and prove the benefits of using ensembles for semantic similarity tasks. The source code that illustrates our approach can be downloaded from <https://github.com/jorge-martinez-gil/sesige>.

*Keywords:* Ensemble Learning, Grammatical Evolution, Semantic Similarity Measurement

---

## 1. Introduction

In recent times, ensemble learning has become a widely used technique to address the limitations of individual methods by combining them into a unified model. Aggregating the predictions of diverse methods aims to mitigate individual method shortcomings, such as outliers in response to specific inputs. Therefore, the fundamental premise behind ensemble learning is the expectation that a carefully chosen set of methods will yield superior results compared to any single method alone [29].

While ensemble learning has attracted considerable attention and received extensive research efforts [10], its application in semantic similarity measurement remains largely unexplored. This presents a compelling opportunity to showcase the potential of this approach to address the

challenge of automatically determining semantic similarity between pieces of textual information. The reason is that, despite advancements in semantic similarity measures, a lack of consensus persists among the individual suitability of these measures when assessing the semantic similarity between textual information [9].

The introduction of ensemble learning to semantic similarity measurement aims to bridge this gap and improve the assessments' reliability and accuracy. The motivation behind this approach comes from the idea that a diversified pool of semantic similarity measures can compensate for the inherent limitations of individual measures [1]. Through the aggregation of multiple measures, our proposed approach seeks to leverage the diversity of these measures to achieve a higher level of agreement and consistency.

Through this research, we aim to contribute to natural language processing (NLP) by providing a novel perspective on semantic similarity measurement. We propose adopting Grammatical Evolution (GE) [37] as an ensemble learning strategy to address the inherent misalignment among existing semantic similarity measures. Empirical evaluations conducted on benchmark datasets will demonstrate the effectiveness of GE ensemble-based approaches in significantly improving performance concerning most existing methods' capabilities.

The rationale behind this research is that GE can bring a new point of view to the semantic similarity measurement domain. The collective recommendation capability of various similarity measures allows for augmenting the quality and consistency of semantic similarity assessments, paving the way for more accurate and reliable real-world applications. Therefore, the significant contributions of this work can be summarized as follows:

- We propose, for the first time, the automatic learning of semantic similarity ensembles based on the notion of GE. This method offers advantages such as high accuracy, excellent interpretability, a platform-independent solution, and easy transferability to problems of analog nature.
- We implement and empirically evaluate our strategy to compare it with existing work and demonstrate its superiority in solving some of the most well-known benchmark dataset used by the research community.

The rest of this paper is organized as follows: Section 2 provides an overview of related work in ensemble learning using GE and other kinds of ensembles for semantic similarity. Section 3 introduces the problem statement. Section 4 presents the details of the proposed GE strategy to address the challenge. Section 5 describes the experimental setup and presents the evaluation results. Section 6 discusses the results obtained and future work directions. Finally, Section 7 concludes the paper.

## 2. State-of-the-art

GE is a particular form of genetic programming (GP) that uses formal grammar (FG) to generate computer programs [37]. GE is considered an evolutionary strategy that makes use of a genotype-to-phenotype strategy. To do that, GE uses an FG definition to describe the language that the model might produce. The most common approach uses the Backus-Naur Form (BNF) [13], a widely used notation to formulate an FG using production rules. These rules include terminals and non-terminals (which can be expanded into terminal and non-terminal symbols).

The BNF grammar allows defining the structure of the ensembles to be learned. Please note that in this work, the term ensemble is equivalent to a program aiming to aggregate an initial set of semantic similarity measures as effectively and efficiently as possible. The FG acts, therefore, as the guideline for the evolution of the ensembles, and it defines the set of valid ensembles that can be generated. This allows for a more structured and controlled evolution compared to rival techniques.

Furthermore, the evolution of the learning process is guided towards optimizing a fitness function, which measures the quality of the generated ensembles in the training phase. In our case, we can evaluate the quality based on the degree of correlation it presents concerning human judgment. Moreover, this fitness function also allows the selection of the ensembles that the parents will use in the next generation. This learning process is repeated until an optimal solution is discovered or a pre-defined number of iterations has been reached.

Furthermore, another advantage of GE is its ability to generate models that adhere to a given syntax and structure. Therefore, this approach is advantageous in domains where the capability of understanding the solution is essential.

### 2.1. Semantic Similarity

The challenge of semantic similarity measurement is a critical task in many computer-related fields [8, 15, 18, 19, 22, 32, 34, 36]. It aims to quantitatively capture the degree of likeness between two pieces of text based on their underlying meaning [16]. Significant progress has been made in recent years, leading to the development of novel techniques [23]. One prominent approach involves utilizing deep learning (DL) models, such as transformer-based architectures like BERT (Bidirectional Encoder Representations from Transformers) [6]. These models are pre-trained on vast amounts of text, enabling them to learn text representations. Fine-tuning these models has shown remarkable performance, outperforming traditional methods that rely on handcrafted features [30].

Another line of research focuses on leveraging distributional semantics, which captures meaning using distributional patterns of words in a large corpus. Methods such as word embeddings (e.g., Word2Vec [26]) represent words as dense vectors in a continuous vector space. The semantic resemblance between the textual pieces can then be estimated by comparing the vector

representations of these pieces using methods like cosine distance. Additionally, recent studies have explored incorporating contextual information using contextualized word embeddings, such as Embeddings from Language Models (ELMo) [33] and Universal Sentence Encoder (USE) [5]. Considering the surrounding words, these models generate context-dependent word representations, leading to improved semantic similarity estimation in a given context.

In recent times, ensembles have also emerged as a helpful technique in semantic similarity measurement, offering a reasonable solution to the challenges posed by the inherent complexity of human language [21]. The idea of aggregating multiple semantic similarity measures allows ensembles to mitigate the limitations of individual measures and capture a more comprehensive understanding of semantic similarity [3]. Ensembles exploit each measure’s inherent complementarity and different perspectives by leveraging the diversity of these existing measures [35]. This means improving the performance and transfer learning capabilities is usually possible [24]. With their ability to aggregate diverse perspectives and mitigate model biases, ensembles have proven helpful in semantic similarity measurement, pushing the boundaries of accuracy and offering promising lines of research [20].

In summary, state-of-the-art techniques for semantic similarity measurement have witnessed significant progress in the last years, driven by the use of DL models, the incorporation of contextual information, and the exploitation of ensembles. These approaches have demonstrated exemplary performance, being superior to traditional methods. As the field continues to evolve, further research and development are expected to improve the existing methods, facilitating many computer-related applications [25].

## *2.2. Grammatical evolution*

GE is a well-known technique in the domain of GP, combining the principles of genetic algorithms (GAs) and FGs. It has gained recognition as a state-of-the-art approach for evolving computer programs that exhibit complex behaviors [31]. It offers a framework to automatically generate programs (ensembles in our particular case) by evolving their syntax and semantics through a GA.

The ensembles can be represented through strings of symbols, which allows their manipulation and evolution using genetic operators through FGs. This facilitates the exploitation of a vast search space that allows the discovery of practical solutions to a wide range of computational problems [40]. Over time, GE has undergone remarkable advances, including knowledge integration, mutation process improvements, and new crossover operators. These advances have improved the accuracy and scalability of GE-based solutions, making it one of the most promising techniques in the GP landscape [39].

The state-of-the-art in GE involves developing hybrid approaches that combine GE with other techniques like particle swarm optimization [12]. These hybrid approaches can leverage the

strengths of multiple techniques to overcome limitations and improve search capability. Additionally, increased focus is on improving scalability through parallel and distributed computing paradigms. Researchers have achieved efficacy in solving computationally intensive problems using these paradigms. Furthermore, advancements in fitness approximation techniques have significantly improved efficiency by reducing computational overhead. Continually exploring novel techniques aims to improve this GP approach’s performance, scalability, and applicability.

### *2.3. Differences between Genetic Programming and Genetic Algorithm*

The main distinction between GP and GAs is their optimization approaches. GAs optimize a given function by searching for optimal parameter values, while GP generates programs (ensembles in this case) that perform well on a specific task. GP uses a higher-level representation to capture complex relationships among variables, enabling the encoding of complex solutions within the population. It incorporates a refined selection process to maintain population diversity and avoid premature convergence. GP’s advanced crossover operator generates novel solutions, while its mutation operator maintains diversity by introducing variations. Additionally, GP utilizes a complex fitness function that ensures a more accurate and thorough assessment of ensembles during the evolutionary process.

### *2.4. Contribution over the state-of-the-art*

We propose exploring GE as a suitable approach for learning ensembles within the domain of semantic similarity measurement. The primary goal is to identify a program (or ensemble, as applicable in our specific case) that attains a near-optimal fitness value for a given objective function, which involves emulating human judgment. While the usual way to proceed is to use a tree-structured expression that can be directly manipulated [21], our method uses genetic operators on an integer string that is subsequently transformed into an ensemble using a BNF grammar. The advantages offered by this strategy include enhanced accuracy, improved interpretability of the resulting model, and a simplified process of converting the model into the most widely used programming languages.

## **3. Problem Statement**

Let us assume a set of measures  $\mathcal{M} = M_1, M_2, \dots, M_n$ , where  $n$  is the number of measures acting as candidates. Let us assume that each  $M_i$  takes two texts  $X$  and  $Y$  as input and produces a single score  $S_i$ .

The aim is to automatically select a subset from  $\mathcal{M}$  and aggregate the corresponding  $M_i$  into an ensemble  $E$ , such that  $E(X, Y)$  provides an accurate score.

Furthemore, let us also assume a vector  $\mathbf{w} = [w_1, w_2, \dots, w_n]$ , where  $w_i \in 0, 1$  states the inclusion of  $M_i$  in the ensemble so that if  $w_i = 0$ , then  $M_i$  is excluded; and if  $w_i = 1$ , then  $M_i$  is included.

The ensemble  $E(X, Y)$  is defined as an aggregation function of a subset from  $\mathcal{M}$  where each  $M_i$  is weighted by its inclusion value as shown in Equation 1.

$$E(X, Y) = \sum_{i=1}^n w_i \cdot M_i(X, Y) \quad (1)$$

In this research, we use GE to build the ensemble function. Please note that GE provides a framework for generating and evolving an ensemble based on BNF grammar. In this case, the BNF grammar defines the rules for constructing the aggregation strategies.

Therefore, the problem consists of finding the  $\mathbf{w}$  that maximizes the ensemble’s performance. Examples of performance can be measures such as precision and recall. Nevertheless, in the case of semantic similarity measurement, the challenge is to emulate human judgment [1]. This means that we need to use statistical methods such as correlation coefficients. Therefore, we aim to optimize the correlation between the ensemble results and a human-curated ground truth dataset.

To do that, given a gold standard  $\mathcal{G}$ , i.e., a dataset created and curated by human experts, the goal is to maximize the correlation between the  $\mathcal{G}$  and the results from the proposed strategy  $\mathcal{S}$  as shown in Equation 2.

$$S = \arg \max_{\mathcal{S}} \text{correl}(\vec{\mathcal{G}}, \vec{\mathcal{S}}) \quad (2)$$

$\mathcal{S}$  can take different semantic similarity measures as input. These measures will function as weak estimators to obtain intermediate semantic similarity scores to learn a higher-level yet robust strategy able to work over unseen data. In short, the goal is to identify an ensemble capable of adapting to training data that can perform well on data never seen before.

GE can evolve candidate ensembles, evaluating their correlation to a human-curated training set. The fitness function guides the search process by assigning fitness to each candidate ensemble based on performance. The process iteratively evolves the population of candidate ensembles, using genetic operators, until a termination condition is met, such as reaching a maximum number of generations (previously defined by the operator) or achieving a satisfactory fitness level for the problem at hand, since the ideal result will be challenging to achieve.

In this way, a computer language’s syntax and semantics can be created following the rules described within GE. These criteria are applied to produce a population of computer programs, or ensembles in our specific case, capable of evolving. The approach generates new strings of symbols equivalent to the most successful ensembles in the population.

The degree to which an ensemble successfully correlates to the ground truth is critical in determining that success. The reason is that ensembles more successful at completing have a greater chance of being picked for reproduction and mutation. In contrast, the less successful ones will not be passed on to the next generation. In this way, the more successful ensembles will become prevalent.

One of the most significant benefits is that GE facilitates building ensembles that can solve complex issues automatically. During the evolutionary process, the approach automatically explores the space of possible ensembles and selects the one that maximizes the performance. Our hypothesis is that the resulting ensemble can estimate semantic similarity for unseen textual inputs. This hypothesis will be empirically tested later in this paper.

#### 4. Methods

We have seen that GE is a powerful evolutionary computation technique that combines GAs with an FG. We can automatically learn complex similarity models capable of capturing the nuances of natural language by leveraging the adaption capability of GE.

The process begins by defining a BNF grammar that represents the structure of the possible semantic similarity models. This BNF grammar serves as a guideline for generating diverse candidate solutions. Each candidate solution represents a unique ensemble of semantic similarity measures. Algorithm 1 shows us how, through an iterative process, the approach explores the space of potential solutions, gradually improving their performance through fitness evaluation and selection.

---

**Algorithm 1** Grammatical Evolution using Genetic Programming

---

```
1: Input: Grammar  $G$ , Population size  $N$ , Termination condition
2: Output: Best individual
3: Initialize population  $P$  with  $N$  random individuals
4: Evaluate fitness for each individual in  $P$ 
5: while termination condition not met do
6:   Select parents for reproduction based on fitness
7:   Initialize empty offspring population  $O$ 
8:   for each pair of parents do
9:     Apply crossover to create two offspring
10:    Apply mutation to each offspring
11:    Add the offspring to  $O$ 
12:   end for
13:   Evaluate fitness for the offspring in  $O$ 
14:   Select individuals for the new population based on fitness
15:   Replace the current population  $P$  with the new population
16: end while
17: return Best individual
```

---

The fitness evaluation is based on an objective function measuring the quality of the generated ensembles. This function could consider factors such as the ensemble’s output’s accuracy, robustness, and diversity, although this research focuses on accuracy (understood as correlation to human judgment). The idea behind aggregating multiple semantic similarity measures allows the ensembles to capture different aspects of the problem being addressed. Furthermore, the adaptive nature of the process enables the ensembles to learn and evolve, continuously refining their performance over time.

GE not only automates the ensemble learning process but also pushes the boundaries of semantic similarity modeling. Allowing the ensembles to adapt to data eliminates the stage of manual selection of semantic similarity measures, which is usually time-consuming and error-prone. Instead, the ensembles learn from the training data, discovering hidden patterns that may not be apparent to the users.

#### 4.1. Mathematical foundation

GE works by implementing a genotype-phenotype mapping mechanism that can be defined as follows:

- Let  $G$  be the genotype, coded as a string representing an individual’s genetic code. The genotype holds several genes, each consisting of a fixed number of bits.
- The phenotype  $P$  is the resulting program generated from the genotype  $G$ . It represents the executable form of the genetic code.
- The mapping from the genotype to the phenotype is a function  $f : \{0, 1\}^n \rightarrow P$ , where  $n$  is the length of the genotype. This function inputs the binary string  $G$  and produces the corresponding phenotype  $P$ .
- The mapping is defined by a grammar  $C$  and a mapping table  $M$ . The grammar  $C$  specifies the production rules that define the syntax of the phenotype so that each rule in the grammar corresponds to a gene in the genotype. The mapping table  $M$  maps each gene in the genotype to a production rule from the grammar.

The mapping function  $f$  could be formulated as follows:

$$f(G) = \begin{cases} \text{Base case :} & \text{If } G \text{ is a terminal gene,} \\ & \text{return the corresponding terminal symbol.} \\ \text{Recursive case :} & \text{If } G \text{ is a non-terminal gene,} \\ & \text{apply the corresponding production rule from } M \\ & \text{to the genes that follow in the genotype.} \end{cases}$$



In this way, the genotype-phenotype mapping transforms the binary representation of the genotype into an executable phenotype.

#### 4.2. Fitness Function

Let  $F(\mathbf{w})$  be the fitness function that evaluates the performance of an ensemble represented by  $\mathbf{w}$ . The fitness function is defined based on the performance regarding the semantic similarity task, i.e., the correlation coefficient. The fitness function, therefore, assesses the quality of the ensemble.

#### 4.3. Genetic Operators

Genetic operators are used in the evolution process of the ensemble configurations. Two commonly used genetic operators are crossover and mutation:

##### 4.3.1. Crossover

Given two parent ensembles  $\mathbf{w}^1$  and  $\mathbf{w}^2$ , crossover produces two offspring ensembles  $\mathbf{w}^1$  and  $\mathbf{w}^2$  by exchanging genetic information between the parents. There are several strategies to do that [31].

##### 4.3.2. Mutation

Mutation introduces random changes to the ensemble. A mutation operator alters one or more elements of the binary vector  $\mathbf{w}$  to explore new configurations. Our mutation process randomly selects positions in  $\mathbf{w}$  and flips the corresponding values.

#### 4.4. Grammar Rules

GE uses grammar rules to generate and interpret promising ensembles. These rules define the structure and constraints for building ensembles using production rules. Each rule specifies how to expand non-terminal symbols into terminal or non-terminal symbols. The grammar rules ensure that the generated configurations follow the required syntax and semantics.

Generally speaking, the exact formulation of the fitness function, genetic operators, and grammar rules can vary depending on the specific implementation and requirements of the task. These aspects must be carefully designed and tailored to the problem domain to effectively guide the search process and generate high-performing ensemble configurations using GE. However, it is usually assumed that a set of points is common to all aggregation strategies. For example, when tackling a problem with GE, a suitable BNF grammar definition must initially be defined. The BNF can be either the specification of an entire language or, perhaps more feasible from the point of view of resource consumption, a subset of a language appropriate for the problem at hand. An example of FG inspired in Python can be seen in Example 1, which defines a language for arithmetic expressions and some additional mathematical functions.

#### EXAMPLE 1

```
<expr> ::= <expr>+<expr>|
          <expr>-<expr>|
          <expr>*<expr>|
          pdiv(<expr>,<expr>)|
          psqrt(<expr>)|
          np.sin(<expr>)|
          np.tanh(<expr>)|
          np.exp(<expr>)|
          plog(<expr>)|
          x[:, 0]|x[:, 1]|x[:, 2]|x[:, 3]|x[:, 4]|
          <c><c>.<c><c>
<c> ::= 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9
```

The ability to easily customize the output structures by manipulating the BNF grammar is an advantage distinguishing GE from rival techniques. Furthermore, the genotype-phenotype mapping enables search operators to operate not only on solution trees, as is the case in standard GP, but also on the genotypes (i.e., integer or binary lists), partially derived phenotypes or the fully-formed phenotypic derivation trees. This capability expands the range of entities on which search operators can act and contributes to the efficacy of the GE methodology.

Our work is implemented using PonyGE2 [7], a state-of-the-art open-source library for GP that offers comprehensive features and tools for evolutionary computation. It combines versatility, efficiency, and usability to explore and optimize complex problems. The library supports GP paradigms, including GE [38].

## 5. Results

In this section, we present the findings of our experiments focused on semantic similarity measurement. We will also explore two ways to build ensembles using the Python language. From now on, we will call one GE, which only searches for accuracy. Furthermore, the other, which we will call GE-i from now on, will look for a Python style that facilitates interpretability. We will see examples later and conduct a comparative analysis of the outcomes produced by our proposed strategies concerning state-of-the-art GP techniques to assess effectiveness and implications.

### 5.1. Baseline Selection

Our baseline is one of the top-performing methods for aggregating similarity scores, i.e., linear regression [17]. Linear regression aims to establish a functional relationship between the previously considered semantic similarity measures. This relationship can be represented using a mathematical equation, which connects the output with multiple semantic similarity measures, as depicted in Eq. 3.

$$\vec{\alpha} = \arg \min (D, \vec{\alpha}) = \arg \min \sum_{i=1}^n (\vec{\alpha} \cdot \vec{a}_i - b_i)^2 \quad (3)$$

Eq. 3 represents the minimization problem involved in linear regression, aiming to find the optimal vector  $\vec{\alpha}$  that minimizes the discrepancy  $D$  between the predicted values and the actual values. The optimization process seeks to minimize the sum of squared differences between the dot product of the vector  $\vec{\alpha}$  and the vector  $\vec{a}_i$ , representing the semantic similarity measures, and the corresponding target values  $b_i$ . The symbol *arg min* denotes the argument that minimizes the expression within the parentheses, and the index  $i$  ranges from 1 to  $n$ , representing the number of instances. In this way, linear regression is a foundational approach for building ensembles by quantifying the association between the semantic similarity measures and the desired output, allowing for the derivation of predictive models.

### 5.2. Datasets

The first dataset used in our experiments is the so-called **Miller & Charles** dataset [27], from now **MC30**. This is the standard dataset community members use when evaluating research methodologies that concentrate on general cases. It includes 30 use cases on the comparison of words of daily use. Therefore, this dataset aims to evaluate the semantic similarity between words that are components of a general-purpose scenario.

The second dataset is the so-called **GeReSiD50** dataset [2] and is drawn from geospatial research. It covers a pool of textual phrases, each of which has been grouped into one of 50 unique pairings. This pool of sentences includes over 100 different geographical expressions. On each of the 50 pairings, human opinions about the degree of semantic similarity were solicited and recorded individually. These 50 pairings include samples that are in no way comparable to one another and others that, in human view, are virtually indistinguishable.

### 5.3. Evaluation Criteria

When exploring correlation coefficients, researchers commonly consider two extensively utilized ones: the Pearson Correlation Coefficient (PCC) and the Spearman Rank Correlation Coefficient (SRCC). On the one hand, the PCC serves as a measure to assess the linear correlation between a gold standard and the resulting outcomes of the proposed strategy. On the other hand,

the SRCC is primarily used to investigate the ordinal correlation between the anticipated values derived from the proposed strategy and the actual observed values. This study aims to closely examine the ensemble’s accuracy concerning these two correlation coefficients, as discussed in [11].

#### 5.4. Empirical results

We provide an overview of the outcomes derived from our empirical assessment of the above benchmarks. Tables 1 and 2 show the reference data for the semantic similarity measures that will be part of the ensemble for solving the **MC30** and the **GeReSiD50** benchmark datasets, respectively. Our primary pool of measures will be based on different variants over BERT [6] since there is a broad consensus about their superiority in tackling this task. **Truth** represents the ground truth values, ranging from 0 to 1, as a reference for comparison. **Bert-Cos.** displays the results obtained by encoding the text pieces using BERT and calculating similarity based on the cosine formula. **Bert-Man.** presents results obtained using the Manhattan distance. **Bert-Euc.** shows results based on the Euclidean distance. **Bert-Inn.** reflects results obtained using the Inner Product similarity measure. Lastly, **Bert-Ang.** illustrates results obtained by calculating similarity using the cosine of the angle.

Table 1: Results obtained for the **MC30** benchmark dataset by different methods in isolation

<b>UC</b>	<b>Truth</b>	<b>Bert-Cos.</b>	<b>Bert-Man.</b>	<b>Bert-Euc.</b>	<b>Bert-Inn.</b>	<b>Bert-Ang.</b>
UC1	1.000	0.921	0.642	0.642	0.993	0.873
UC2	0.980	0.818	0.462	0.462	0.863	0.805
UC3	0.980	0.899	0.607	0.605	0.922	0.856
UC4	0.959	0.936	0.678	0.680	1.000	0.886
UC5	0.944	0.860	0.525	0.526	0.916	0.830
UC6	0.921	0.558	0.165	0.170	0.577	0.688
UC7	0.893	0.839	0.488	0.491	0.893	0.817
UC8	0.872	0.855	0.507	0.512	0.926	0.826
UC9	0.793	0.824	0.471	0.465	0.886	0.808
UC10	0.786	0.615	0.216	0.208	0.665	0.711
UC11	0.778	0.512	0.124	0.125	0.533	0.671
UC12	0.758	0.679	0.296	0.291	0.704	0.738
UC13	0.753	0.842	0.494	0.492	0.907	0.818
UC14	0.719	0.621	0.230	0.222	0.658	0.713
UC15	0.423	0.685	0.294	0.291	0.725	0.740
UC16	0.429	0.641	0.238	0.242	0.680	0.721
UC17	0.296	0.530	0.141	0.138	0.556	0.678
UC18	0.281	0.523	0.120	0.127	0.554	0.675
UC19	0.242	0.712	0.310	0.313	0.776	0.752
UC20	0.227	0.479	0.079	0.084	0.512	0.659
UC21	0.222	0.693	0.307	0.307	0.719	0.744
UC22	0.214	0.672	0.285	0.270	0.724	0.735
UC23	0.161	0.626	0.241	0.219	0.677	0.715
UC24	0.140	0.487	0.079	0.100	0.509	0.662
UC25	0.107	0.476	0.089	0.085	0.504	0.658
UC26	0.107	0.560	0.166	0.161	0.595	0.689
UC27	0.033	0.534	0.147	0.131	0.573	0.679
UC28	0.028	0.492	0.134	0.106	0.512	0.664
UC29	0.020	0.645	0.246	0.254	0.670	0.723
UC30	0.002	0.384	0.000	0.000	0.413	0.625

Table 2: Results obtained for the **GeReSiD50** benchmark dataset by different methods in isolation

UC	Truth	Bert-Cos.	Bert-Man.	Bert-Euc.	Bert-Inn.	Bert-Ang.
UC1	0.017	0.320	0.046	0.133	0.373	0.604
UC2	0.021	0.275	0.054	0.109	0.316	0.589
UC3	0.031	0.391	0.139	0.193	0.440	0.628
UC4	0.050	0.450	0.160	0.220	0.525	0.649
UC5	0.052	0.174	0.000	0.050	0.200	0.556
UC6	0.058	0.544	0.238	0.300	0.616	0.683
UC7	0.072	0.354	0.089	0.160	0.408	0.615
UC8	0.081	0.563	0.260	0.310	0.646	0.690
UC9	0.085	0.240	0.015	0.080	0.281	0.577
UC10	0.094	0.233	0.025	0.088	0.267	0.575
UC11	0.109	0.152	0.000	0.023	0.181	0.549
UC12	0.124	0.377	0.098	0.164	0.446	0.623
UC13	0.139	0.394	0.133	0.181	0.460	0.629
UC14	0.149	0.477	0.163	0.228	0.571	0.658
UC15	0.154	0.497	0.207	0.258	0.574	0.666
UC16	0.161	0.683	0.374	0.428	0.742	0.739
UC17	0.204	0.368	0.099	0.164	0.428	0.620
UC18	0.210	0.606	0.299	0.354	0.677	0.707
UC19	0.217	0.456	0.185	0.234	0.519	0.651
UC20	0.235	0.366	0.121	0.176	0.414	0.619
UC21	0.269	0.634	0.310	0.359	0.749	0.719
UC22	0.273	0.319	0.095	0.139	0.365	0.603
UC23	0.290	0.510	0.204	0.271	0.582	0.670
UC24	0.328	0.603	0.279	0.339	0.700	0.706
UC25	0.369	0.413	0.122	0.184	0.493	0.635
UC26	0.389	0.506	0.200	0.256	0.597	0.669
UC27	0.391	0.768	0.456	0.497	0.883	0.779
UC28	0.399	0.676	0.356	0.404	0.782	0.736
UC29	0.417	0.669	0.348	0.395	0.776	0.733
UC30	0.438	0.501	0.197	0.255	0.587	0.667
UC31	0.490	0.639	0.315	0.369	0.740	0.720
UC32	0.514	0.427	0.136	0.206	0.495	0.640
UC33	0.535	0.497	0.191	0.259	0.571	0.666
UC34	0.557	0.492	0.174	0.236	0.594	0.664
UC35	0.594	0.800	0.500	0.534	0.915	0.795
UC36	0.611	0.561	0.243	0.309	0.641	0.689
UC37	0.617	0.753	0.444	0.480	0.868	0.771
UC38	0.621	0.713	0.400	0.442	0.815	0.753
UC39	0.645	0.532	0.230	0.284	0.614	0.679
UC40	0.650	0.665	0.354	0.400	0.750	0.731
UC41	0.668	0.574	0.256	0.317	0.665	0.695
UC42	0.748	0.920	0.682	0.706	1.053	0.872
UC43	0.762	0.704	0.385	0.426	0.826	0.749
UC44	0.764	0.631	0.330	0.372	0.710	0.717
UC45	0.764	0.726	0.399	0.449	0.848	0.759
UC46	0.769	0.658	0.333	0.391	0.751	0.729
UC47	0.781	0.572	0.248	0.312	0.666	0.694
UC48	0.811	0.651	0.322	0.382	0.750	0.726
UC49	0.873	0.751	0.425	0.475	0.876	0.770
UC50	0.904	0.866	0.588	0.617	1.000	0.834

It is important to remark that the outcomes of our reported experiments are based on 30 independent runs, owing to the inherent non-deterministic characteristics of the methods. Therefore, we aim to report a snapshot of the values achieved.

### 5.5. Assessing Semantic Similarity in a General-purpose Context

Figure 1 displays the results for two evaluation criteria, PCC and SRCC, over the **MC30** benchmark dataset. The x-axis represents different strategies used for evaluation. At the same time, Linear Regression (LR) is the baseline, as discussed earlier. A dotted horizontal line represents it.

The state-of-the-art genetic ensembles are Tree-based Genetic Programming (TGP) [14], Linear Genetic Programming (LGP) [4], and Cartesian Genetic Programming (CGP) [28] precisely as in [21]. GE is the approach proposed in this work, and GE-i is the interpretable variant of GE discussed earlier. It is essential to note that all the ensembles are trained on the same training dataset to facilitate the fairness of the comparisons.

In the first subplot (a), the LGP achieves relatively high performance compared to the other methods. The boxplot displays the distribution of PCC values obtained from 30 experimental runs. The box represents the interquartile range (IQR), where the central box spans from the lower quartile (Q1) to the upper quartile (Q3). The line within the box corresponds to the median value. The whiskers extend to the minimum and maximum values.

In the second subplot (b), the GE method (first blue boxplot) demonstrates the best performance regarding SRCC. The boxplot characteristics are the same as in the previous subplot but now represent the distribution of SRCC values.

Overall, both subplots suggest that the LGP outperforms the other evaluated methods regarding PCC, and GE is superior regarding SRCC on the **MC30** benchmark dataset. GE-i, although interpretable, achieves the worst performance.




Figure 1: Results for the a) **PCC** and b) **SRCC** over the **MC30** benchmark dataset

As a matter of curiosity, we can see in Example 2 the code generated for both PCC and SRCC over the MC30 dataset. This given source code is represented in Python and uses the Numpy library, which supports mathematical operations on arrays and matrices. The result is computed using various mathematical functions and operators. The reason is that we are using the FG seen in Example 1. It is important to note that the expressions within parentheses are evaluated and combined using the specified operators.

EXAMPLE 2

Ensemble optimized for PCC over MC30

```
import numpy as np

result = (
    BERT-Euc - BERT-Inn + pdiv(BERT-Euc, np.sin(BERT-Ang)) - BERT-Cos +
    np.exp(psqrtpdiv(np.tanh(BERT-Man), BERT-Ang))) + BERT-Euc +
    psqrtpdiv(BERT-Cos, np.sin(pdiv(BERT-Inn, pdiv(np.sin(BERT-Ang),
    BERT-Inn) * BERT-Man - BERT-Euc) * pdiv(BERT-Man, BERT-Inn))))
) / (BERT-Inn - pdiv(71.24, BERT-Cos * plog(76.12)))
```

Ensemble optimized for SRCC over MC30

```
import numpy as np

result = (
    BERT-Euc - BERT-Inn + pdiv(BERT-Euc, np.sin(BERT-Ang)) - BERT-Cos +
    np.exp(psqrtpdiv(np.tanh(BERT-Man), BERT-Ang))) + BERT-Euc +
    psqrtpdiv(BERT-Cos, np.sin(pdiv(BERT-Inn, pdiv(np.sin(BERT-Ang),
    BERT-Inn) * BERT-Man - BERT-Euc) * pdiv(BERT-Man, BERT-Inn))))
) / (BERT-Inn - pdiv(71.24, BERT-Cos * plog(76.12)))
```

We also show the changes over time in essential variables during the GE process. Figure 2 shows the progression of these parameters. Specifically, we focus on four key parameters: Average Fitness, Average Genome Length, Average Tree Nodes, and Best Fitness.

- The **Average Fitness** provides insights into the overall performance of the evolving population. It reflects the average fitness value of individuals in each generation, indicating the progress achieved by the GE strategy.
- The **Average Genome Length** tracks the average length of individual genomes within the population at different training stages. The goal of monitoring this variable is to understand how the complexity of GE-generated solutions changes over time.



- The **Average Tree Nodes** measures the average number of nodes in the evolved solutions. It offers valuable information about the complexity and intricacy of evolved ensembles, shedding light on the strategy’s search for space exploration.
- Lastly, the **Best Fitness** represents the fitness value of the best individual in each generation. Observing this variable helps to assess the progress in finding optimal solutions as training is performed. Please remember that these values are for the training phase, and then it remains to test the generated ensemble on previously unseen data.

Analyzing the evolution of these variables allows us to obtain insights into how they contribute to PCC optimization and interact during the GE process over the MC30 benchmark dataset. This analysis provides a valuable view into the behavior of GE and the overall performance of the approach.

The examination of Figure 3 reports us a comprehensive visualization concerning the progressive evolution of the aforementioned important variables when optimizing SRCC over the **MC30** benchmark dataset.

### 5.6. Assessing Semantic Similarity in a Domain-Specific Context

Figure 4 displays the results for both PCC and SRCC over the **GeReSiD50** benchmark dataset. As in the previous case, the x-axis represents different strategies used for evaluation. Linear Regression (LR) is again the baseline, as discussed earlier, and is represented by a dotted horizontal line. The state-of-the-art genetic ensembles are again TGP [14], LGP [4], and CGP [28]. GE is again the approach proposed in this work, and GE-i is the interpretable variant of GE, precisely as we discussed in the previous case.

In the first subplot (a), the LGP achieves relatively high performance compared to the other methods. The boxplot displays the distribution of PCC values obtained from 30 experimental runs. The box again represents the IQR, where the central box spans from the lower to the upper quartile. The line within the box corresponds to the median value. The whiskers extend to the minimum and maximum values.

In the second subplot (b), the GE method demonstrates the best performance regarding SRCC. The boxplot characteristics are the same as in the previous subplot but now represent the distribution of SRCC values.

As with the general purpose use case, both subplots suggest again that the LGP outperforms the other evaluated methods regarding PCC, and GE is superior regarding SRCC on the **GeReSiD50** benchmark dataset. GE-i, although interpretable, achieves the worst performance once again.

We provide the automatically-generated Python source code in Example 3. It is an ensemble optimized for assessing the PCC over MC30 that mainly consists of two functions:




Figure 2: Evolution of different variables during the ensemble learning process for **PCC** over the **MC30** benchmark dataset

$my\_pearson(x, y)$  and  $p()$ . The  $my\_pearson(x, y)$  function calculates the PCC between two float arrays, while the  $p()$  function aims to maximize through an algebraic formula that needs to be learned. In order to do that, the source code reads data from training and test CSV files, extracts the relevant columns, and performs the proper calculations to generate a new column with the expression to be learned.

The PCC coefficient between the *response* column and the new column is then computed, which serves as the goal (PCC over unseen data) to be maximized. In this way, the code also includes functionality to apply the learned expression to a separate test dataset, calculating the resulting PCC and enabling the performance assessment on unseen data. This approach ensures that the optimized ensemble learns from the training data and shows its generalization ability on previously unseen data.

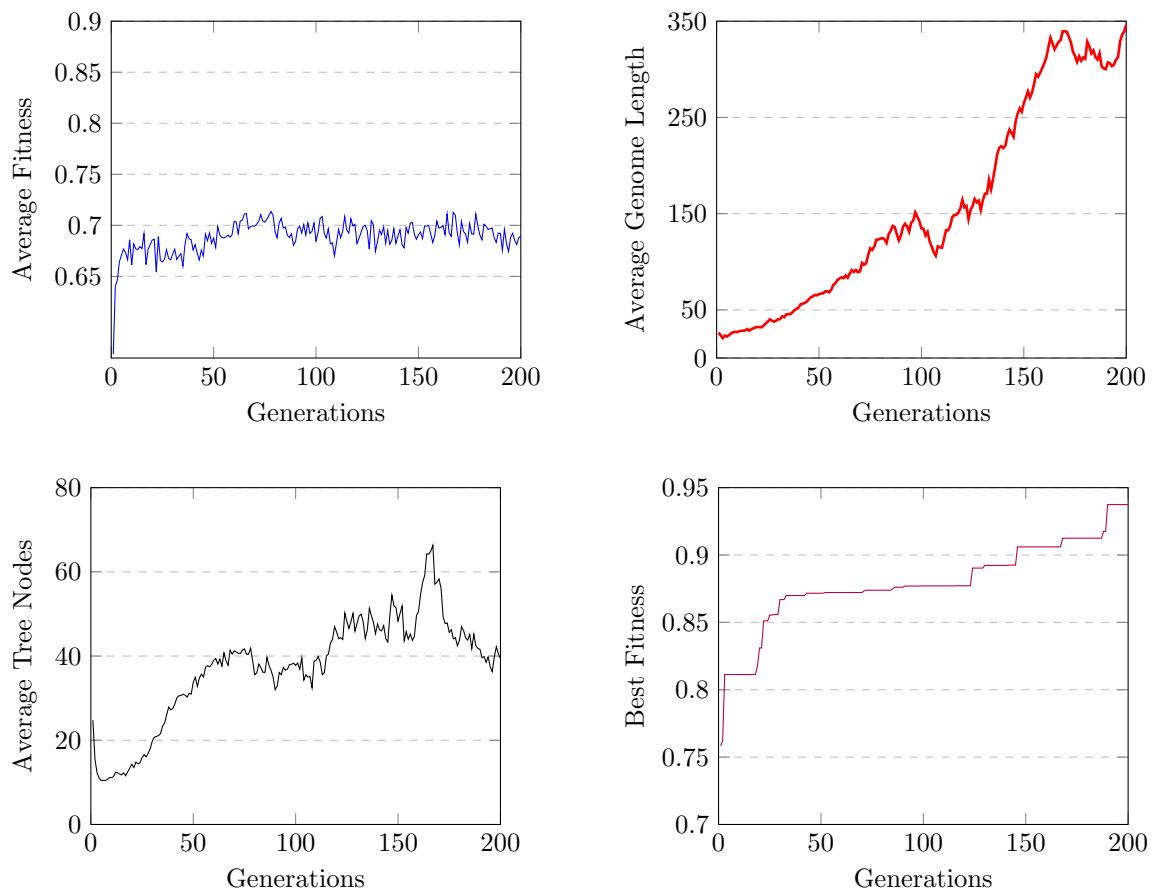


Figure 3: Evolution of different variables during the ensemble learning process for **SRCC** over the **MC30** benchmark dataset

### EXAMPLE 3

Ensemble optimized for PCC over MC30

```
import pandas as pd
import numpy as np

def my_pearson(x, y):
    return np.abs(np.corrcoef(x, y)[0,1])

def p():

    df = pd.read_csv('c:/mc-training.txt')
    df2 = pd.read_csv('c:/mc-validation.txt')

    x, x0, x1, x2, x3, x4 = df['response'].to_numpy(), \
        df['x0'].to_numpy(), df['x1'].to_numpy(), df['x2'].to_numpy(), \
        df['x3'].to_numpy(), df['x4'].to_numpy()

    y, y0, y1, y2, y3, y4 = df2['response'].to_numpy(), \
        df2['y0'].to_numpy(), df2['y1'].to_numpy(), df2['y2'].to_numpy(), \
        df2['y3'].to_numpy(), df2['y4'].to_numpy()
```

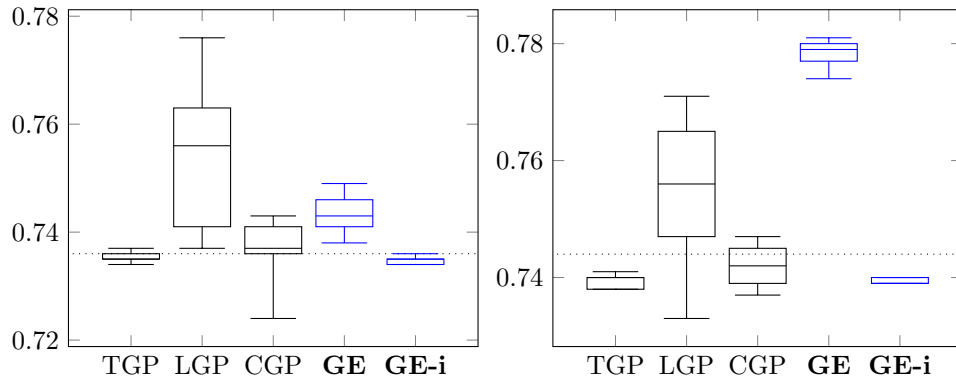


Figure 4: Results for the a) **PCC** and b) **SRCC** over the **GeReSiD50** benchmark dataset

We also provide the generated code for SRCC in Example 4. It is an ensemble optimized for SRCC over MC30 that implements two functions,  $my\_spearman(x, y)$  and  $p()$ , to maximize SRCC. The  $my\_spearman(x, y)$  function calculates the SRCC between two arrays. The  $p()$  function loads training and validation datasets, extracts relevant columns, and performs calculations on the data. It defines an auxiliary expression involving variables, replaces one set of variables with another, evaluates the expression, and assigns the results to a new column in the validation dataset. Finally, the SRCC is computed between the *response* and newly created columns. The objective is maximizing the value returned by  $p()$ , representing the SRCC over unseen data.

EXAMPLE 4

Ensemble optimized for SRCC over MC30

```

import pandas as pd
import numpy as np
from scipy.stats import spearmanr

def my_spearman(x, y):
    return np.abs(spearmanr(x, y)[0])

def p():

    df = pd.read_csv('c:/geresid-training.txt')
    df2 = pd.read_csv('c:/geresid-validation.txt')

    x, x0, x1, x2, x3, x4 = df['response'].to_numpy(), \
        df['x0'].to_numpy(), df['x1'].to_numpy(), df['x2'].to_numpy(), \
        df['x3'].to_numpy(), df['x4'].to_numpy()

    y, y0, y1, y2, y3, y4 = df2['response'].to_numpy(), \
        df2['y0'].to_numpy(), df2['y1'].to_numpy(), df2['y2'].to_numpy(), \
        df2['y3'].to_numpy(), df2['y4'].to_numpy()

    aux = 'x3 * x3 * x4'
    aux2 = aux.replace('x', 'y')
    df2['new'] = eval(aux2)

    return my_spearman(y, df2['new'].to_numpy())

* The goal is to maximize p()

```

Once again, examining Figure 5 deepens our understanding of the progressive evolution of critical variables in generating the ensemble using **PCC** over the **GeReSiD50** benchmark dataset. This analysis sheds light on the optimization process's behavior, the approach's effectiveness, and the interplay between critical parameters.

At the same time, and once again, the examination of Figure 6 shows us a comprehensive understanding of the progressive evolution of these critical variables. However, this time is intended to understand better the process of optimizing **SRCC** over the **GeReSiD50** benchmark dataset.

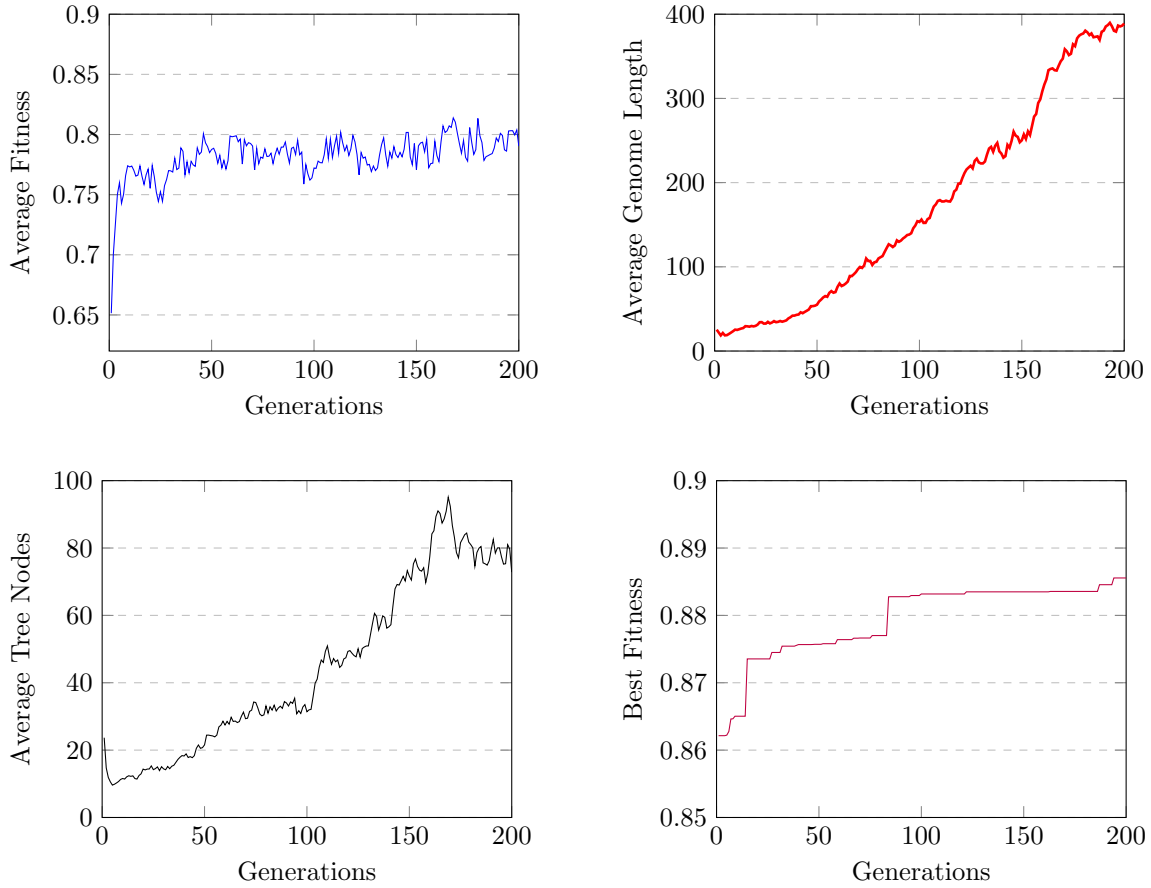


Figure 5: Evolution of different variables during the ensemble learning process for **PCC** over the **GeReSiD50** benchmark dataset

This strategy is based on the iterative nature of the optimization process. It is not a linear progression but a continuous analysis and refinement cycle. Each iteration builds upon the lessons learned from the previous ones, gradually converging towards a more effective setting.

### 5.7. Summary of results

Table 3 summarizes the results obtained for the **MC30** benchmark dataset. It is divided into two sections, and each of them features two columns: the first denoting the method or ensemble used and the second representing the performance, i.e., the PCC in the initial section and SRCC in the subsequent section. These scores assess the degree of correlation between the predicted and ground truth values. Values are reported as the median of the results of the 30 independent runs.




Figure 6: Evolution of different variables during the ensemble learning process for **SRCC** over the **GeReSiD50** benchmark dataset

Table 3: Summary of results obtained for the **MC30** benchmark dataset

Method/Ensemble	PCC	Method/Ensemble	SRCC
Bert-Cos.	0.740	Bert-Cos.	0.701
Bert-Man.	0.744	Bert-Man.	0.689
Bert-Euc.	0.751	Bert-Euc.	0.718
Bert-Inn.	0.728	Bert-Inn.	0.711
Bert-Ang.	0.746	Bert-Ang.	0.701
LR	0.757	LR	0.770
TGP	0.757	TGP	0.758
LGP	<u>0.845</u>	LGP	0.822
CGP	0.777	CGP	0.766
<b>GE</b>	0.794	<b>GE</b>	<u>0.859</u>
<b>GE-i</b>	0.752	<b>GE-i</b>	0.719

The tabular presentation of the results enables comparisons of the effectiveness of various methods or ensembles, thus facilitating the identification of optimal approaches for the specific task. We can see that LGP is giving better results for **PCC** and GE for **SRCC**.

Table 4 summarizes the results obtained for the **GeReSiD50** benchmark dataset. The table also consists of two sections, each containing two columns. The first column displays the method or ensemble used in the study, while the second column represents the performance denoted as the **PCC** and **SRCC**, respectively. Values are again reported as the median result of the 30 independent runs.

Table 4: Summary of results obtained for the **GeReSiD50** benchmark dataset

Method/Ensemble	PCC	Method/Ensemble	SRCC
Bert-Cos.	0.725	Bert-Cos.	0.724
Bert-Man.	0.706	Bert-Man.	0.715
Bert-Euc.	0.711	Bert-Euc.	0.727
Bert-Inn.	0.735	Bert-Inn.	0.740
Bert-Ang.	0.722	Bert-Ang.	0.724
LR	0.736	LR	0.744
TGP	0.735	TGP	0.740
LGP	<u>0.756</u>	LGP	0.752
CGP	0.738	CGP	0.745
<b>GE</b>	0.743	<b>GE</b>	<u>0.779</u>
<b>GE-i</b>	0.735	<b>GE-i</b>	0.740

It is possible to observe that when operating over the **GeReSiD50** dataset, LGP performs better in terms of **PCC**, and GE presents better results in terms of **SRCC**, as in the previous case.

## 6. Discussion

Semantic similarity ensembles are advantageous over other methods as they can effectively leverage the capabilities of a broad spectrum of established semantic similarity measures. As a result, these models often yield predictions of superior accuracy compared to utilizing individual methods in isolation. Our work has shown the advantages of using GE techniques to build ensembles in this context. Our research on the use of GE to address this specific challenge allows us to identify several advantages over traditional GP-based strategies:

1. It presents greater flexibility, allowing for evolving solutions with diverse structures. This flexibility enables GE to handle the problem effectively.
2. It demonstrates good efficiency compared to other GP-based methods due to generating directly executable code for each solution. This process reduces computational overhead and accelerates the evolution of solutions.



3. It establishes an appropriate trade-off between accuracy and interpretability. Achieving very accurate or interpretable ensembles in extreme cases or a balance between both features when necessary.

However, GE also has its drawbacks. One area for improvement lies in interpreting the evolved solutions. Our approach might improve efficiency but sacrifice transparency, making it challenging for people to comprehend the evolved solutions’ inner workings and underlying mechanisms. Moreover, this kind of encoding also poses difficulties in debugging GE, as understanding the complex relationships within the evolved solutions can take time and effort. Lastly, GE is usually blind in finding good starting points for the search process.

## 7. Conclusion

In this work, we have presented a novel approach for the automated design of ensembles of semantic similarity measures using GE. Through empirical evaluations on several benchmark datasets, we have demonstrated the superior performance of our method compared to existing state-of-the-art GP-based ensembles in some cases. These findings show the potential of GE for automatic semantic similarity measure selection and aggregation when the goal is to achieve superior performance compared to individual semantic similarity measures.

Furthermore, our proposed strategy offers several notable advantages over traditional methods: It enables handling a large pool of candidate semantic similarity measures without requiring manual feature selection or parameter tuning, alleviating the process’s time-consuming and knowledge-intensive aspects. Moreover, our approach demonstrates flexibility, allowing human operators to easily add or remove semantic similarity measures from the initial pool per their requirements.

In conclusion, our proposed strategy exhibits promising potential for automated NLP system design. Moreover, extending our approach to other machine learning domains beyond semantic similarity measures is possible. Future research endeavors could explore the application of GE to other NLP tasks and investigate alternative search strategies and fitness functions. Overall, our work sheds light on the significance of ensemble methods and showcases the potential of evolutionary algorithms in facilitating semantic similarity measurement.

## Acknowledgments

The research reported in this paper has been funded by the Federal Ministry for Climate Action, Environment, Energy, Mobility, Innovation, and Technology (BMK), the Federal Ministry for Digital and Economic Affairs (BMDW), and the State of Upper Austria in the frame of SCCH, a center in the COMET - Competence Centers for Excellent Technologies Programme managed by Austrian Research Promotion Agency FFG.

## References

- [1] Ballatore, A., Bertolotto, M., & Wilson, D. C. (2013). The semantic similarity ensemble. *J. Spatial Inf. Sci.*, 7, 27–44. doi:10.5311/JOSIS.2013.7.128.
- [2] Ballatore, A., Bertolotto, M., & Wilson, D. C. (2014). An evaluative baseline for geo-semantic relatedness and similarity. *GeoInformatica*, 18, 747–767. doi:10.1007/s10707-013-0197-8.
- [3] Bär, D., Biemann, C., Gurevych, I., & Zesch, T. (2012). UKP: computing semantic textual similarity by combining multiple content similarity measures. In E. Agirre, J. Bos, & M. T. Diab (Eds.), *Proceedings of the 6th International Workshop on Semantic Evaluation, SemEval-NAACL-HLT 2012, Montréal, Canada, June 7-8, 2012* (pp. 435–440). The Association for Computer Linguistics.
- [4] Brameier, M., Banzhaf, W., & Banzhaf, W. (2007). *Linear genetic programming* volume 1. Springer.
- [5] Cer, D., Yang, Y., Kong, S., Hua, N., Limtiaco, N., John, R. S., Constant, N., Guajardo-Cespedes, M., Yuan, S., Tar, C., Strophe, B., & Kurzweil, R. (2018). Universal sentence encoder for english. In E. Blanco, & W. Lu (Eds.), *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, EMNLP 2018: System Demonstrations, Brussels, Belgium, October 31 - November 4, 2018* (pp. 169–174). Association for Computational Linguistics. doi:10.18653/v1/d18-2029.
- [6] Devlin, J., Chang, M., Lee, K., & Toutanova, K. (2019). BERT: pre-training of deep bidirectional transformers for language understanding. In J. Burstein, C. Doran, & T. Solorio (Eds.), *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers)* (pp. 4171–4186). Association for Computational Linguistics. doi:10.18653/v1/n19-1423.
- [7] Fenton, M., McDermott, J., Fagan, D., Forstenlechner, S., Hemberg, E., & O’Neill, M. (2017). Ponyge2: Grammatical evolution in python. In *Proceedings of the Genetic and Evolutionary Computation Conference Companion* (pp. 1194–1201).
- [8] Han, L., Kashyap, A. L., Finin, T., Mayfield, J., & Weese, J. (2013). Umbc\_ebiquity-core: Semantic textual similarity systems. In M. T. Diab, T. Baldwin, & M. Baroni (Eds.), *Proceedings of the Second Joint Conference on Lexical and Computational Semantics, \*SEM 2013, June 13-14, 2013, Atlanta, Georgia, USA* (pp. 44–52). Association for Computational Linguistics.

- [9] Harispe, S., Ranwez, S., Janaqi, S., & Montmain, J. (2015). *Semantic Similarity from Natural Language and Ontology Analysis*. Synthesis Lectures on Human Language Technologies. Morgan & Claypool Publishers. doi:10.2200/S00639ED1V01Y201504HLT027.
- [10] He, X., Zhao, K., & Chu, X. (2021). Automl: A survey of the state-of-the-art. *Knowl. Based Syst.*, 212, 106622. doi:10.1016/j.knosys.2020.106622.
- [11] Hill, F., Reichart, R., & Korhonen, A. (2015). Simlex-999: Evaluating semantic models with (genuine) similarity estimation. *Comput. Linguistics*, 41, 665–695. doi:10.1162/COLI\_a\_00237.
- [12] Kennedy, J., & Eberhart, R. (1995). Particle swarm optimization. In *Proceedings of ICNN'95-international conference on neural networks* (pp. 1942–1948). IEEE volume 4.
- [13] Knuth, D. E. (1964). Backus normal form vs. backus naur form. *Communications of the ACM*, 7, 735–736.
- [14] Koza, J. R. (1992). *Genetic programming: on the programming of computers by means of natural selection* volume 1. MIT press.
- [15] Lastra-Díaz, J. J., & García-Serrano, A. (2015). A new family of information content models with an experimental survey on wordnet. *Knowl.-Based Syst.*, 89, 509–526. doi:10.1016/j.knosys.2015.08.019.
- [16] Lastra-Díaz, J. J., García-Serrano, A., Batet, M., Fernández, M., & Chirigati, F. (2017). HESML: A scalable ontology-based semantic similarity measures library with a set of reproducible experiments and a replication dataset. *Inf. Syst.*, 66, 97–118. doi:10.1016/j.is.2017.02.002.
- [17] Lastra-Díaz, J. J., Goikoetxea, J., Taieb, M. A. H., García-Serrano, A., Aouicha, M. B., & Agirre, E. (2019). A reproducible survey on word embeddings and ontology-based methods for word similarity: Linear combinations outperform the state of the art. *Eng. Appl. Artif. Intell.*, 85, 645–665. doi:10.1016/j.engappai.2019.07.010.
- [18] Martínez-Gil, J. (2014). An overview of textual semantic similarity measures based on web intelligence. *Artif. Intell. Rev.*, 42, 935–943. URL: <https://doi.org/10.1007/s10462-012-9349-8>. doi:10.1007/s10462-012-9349-8.
- [19] Martínez-Gil, J. (2016). Coto: A novel approach for fuzzy aggregation of semantic similarity measures. *Cogn. Syst. Res.*, 40, 8–17. doi:10.1016/j.cogsys.2016.01.001.

- [20] Martinez-Gil, J. (2022). A comprehensive review of stacking methods for semantic similarity measurement. *Machine Learning with Applications*, *10*, 100423. doi:10.1016/j.mlwa.2022.100423.
- [21] Martinez-Gil, J. (2023). A comparative study of ensemble techniques based on genetic programming: A case study in semantic similarity assessment. *Int. J. Softw. Eng. Knowl. Eng.*, *33*, 289–312. doi:10.1142/S0218194022500772.
- [22] Martinez-Gil, J., & Chaves-Gonzalez, J. M. (2019). Automatic design of semantic similarity controllers based on fuzzy logics. *Expert Syst. Appl.*, *131*, 45–59. doi:10.1016/j.eswa.2019.04.046.
- [23] Martinez-Gil, J., & Chaves-Gonzalez, J. M. (2020). A novel method based on symbolic regression for interpretable semantic similarity measurement. *Expert Syst. Appl.*, *160*, 113663. doi:10.1016/j.eswa.2020.113663.
- [24] Martinez-Gil, J., & Chaves-Gonzalez, J. M. (2023). Transfer learning for semantic similarity measures based on symbolic regression. *Journal of Intelligent & Fuzzy Systems*, (pp. 1–13).
- [25] Martinez-Gil, J., Mokadem, R., Küng, J., & Hameurlain, A. (2023). Neurofuzzy semantic similarity measurement. (p. 102155). volume 145. URL: <https://doi.org/10.1016/j.datak.2023.102155>. doi:10.1016/j.datak.2023.102155.
- [26] Mikolov, T., Sutskever, I., Chen, K., Corrado, G. S., & Dean, J. (2013). Distributed representations of words and phrases and their compositionality. In *Advances in Neural Information Processing Systems 26: 27th Annual Conference on Neural Information Processing Systems 2013. Proceedings of a meeting held December 5-8, 2013, Lake Tahoe, Nevada, United States*. (pp. 3111–3119).
- [27] Miller, G., & Charles, W. (1991). Contextual correlates of semantic similarity. *Language and Cognitive Processes*, *6*, 1–28.
- [28] Miller, J. F. (2020). Cartesian genetic programming: its status and future. *Genet. Program. Evolvable Mach.*, *21*, 129–168. doi:10.1007/s10710-019-09360-6.
- [29] Naimi, A. I., & Balzer, L. B. (2018). Stacked generalization: an introduction to super learning. *European journal of epidemiology*, *33*, 459–464.
- [30] Navigli, R., & Martelli, F. (2019). An overview of word and sense similarity. *Nat. Lang. Eng.*, *25*, 693–714. doi:10.1017/S1351324919000305.
- [31] O’Neill, M., & Ryan, C. (2001). Grammatical evolution. *IEEE Transactions on Evolutionary Computation*, *5*, 349–358.

- [32] Pedersen, T., Patwardhan, S., & Michelizzi, J. (2004). Wordnet: : Similarity - measuring the relatedness of concepts. In D. L. McGuinness, & G. Ferguson (Eds.), *Proceedings of the Nineteenth National Conference on Artificial Intelligence, Sixteenth Conference on Innovative Applications of Artificial Intelligence, July 25-29, 2004, San Jose, California, USA* (pp. 1024–1025). AAAI Press / The MIT Press.
- [33] Peters, M. E., Neumann, M., Iyyer, M., Gardner, M., Clark, C., Lee, K., & Zettlemoyer, L. (2018). Deep contextualized word representations. In M. A. Walker, H. Ji, & A. Stent (Eds.), *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2018, New Orleans, Louisiana, USA, June 1-6, 2018, Volume 1 (Long Papers)* (pp. 2227–2237). Association for Computational Linguistics. doi:10.18653/v1/n18-1202.
- [34] Pirrò, G. (2009). A semantic similarity metric combining features and intrinsic information content. *Data Knowl. Eng.*, 68, 1289–1308. doi:10.1016/j.datak.2009.06.008.
- [35] Potash, P., Boag, W., Romanov, A., Ramanishka, V., & Rumshisky, A. (2016). Simihawk at semeval-2016 task 1: A deep ensemble system for semantic textual similarity. In *Proceedings of the 10th International Workshop on Semantic Evaluation, SemEval-NAACL-HLT 2016, San Diego, CA, USA, June 16-17, 2016* (pp. 741–748).
- [36] Rus, V., Lintean, M. C., Banjade, R., Niraula, N. B., & Stefanescu, D. (2013). SEMILAR: the semantic similarity toolkit. In *51st Annual Meeting of the Association for Computational Linguistics, ACL 2013, Proceedings of the Conference System Demonstrations, 4-9 August 2013, Sofia, Bulgaria* (pp. 163–168).
- [37] Ryan, C., Collins, J. J., & Neill, M. O. (1998). Grammatical evolution: Evolving programs for an arbitrary language. In *Genetic Programming: First European Workshop, EuroGP'98 Paris, France, April 14-15, 1998 Proceedings 1* (pp. 83–96). Springer.
- [38] Vu, T. M. (2021). Software review: Pony ge2. *Genetic programming and evolvable machines*, 22, 383–385.
- [39] Wang, H., Lou, Y., & Bäck, T. (2019). Hyper-parameter optimization for improving the performance of grammatical evolution. In *2019 IEEE Congress on Evolutionary Computation (CEC)* (pp. 2649–2656). IEEE.
- [40] Whigham, P. A. et al. (1995). Grammatically-based genetic programming. In *Proceedings of the workshop on genetic programming: from theory to real-world applications* (pp. 33–41). Citeseer volume 16.