CODE EDITING FROM FEW EXEMPLARS BY ADAPTIVE MULTI-EXTENT COMPOSITION

Anonymous authors
Paper under double-blind review

ABSTRACT

This paper considers the computer source code editing with few exemplars. The editing exemplar, containing the original and modified support code snippets, showcases a certain editorial style and implies the edit intention for a query code snippet. To achieve this, we propose a machine learning approach to adapt the editorial style derived from few exemplars to a query code snippet. Our learning approach combines edit representations extracted from editing exemplars and compositionally generalizes them to the query code snippet editing via multi-extent similarities ensemble. Specifically, we parse the code snippets using language-specific grammar into abstract syntax trees. We apply the similarities measurement in multiple extents from individual nodes to collective tree representations, and ensemble them through a similarity-ranking error estimator. We evaluate the proposed method on two datasets in C# and Python languages and respectively show 8.0% and 10.9% absolute accuracy improvements compared to baselines.

1 INTRODUCTION

In recent years, a surge of interest has been witnessed in applying machine learning techniques to code editing (Zhao et al., 2019; Chen et al., 2019; Dinella et al., 2020; Chakraborty et al., 2020; Yasunaga & Liang, 2020). Code editing in software engineering intends to revise the design, structure, functionality, or implementation of existing programming codes into a desirable style. To maintain high-quality code projects or merge several repositories, the programmers typically fix one or few code snippets, and expect the same revision automatically applied to everywhere in demand over the whole project. Automating this process can facilitate a broad range of programming applications such as code migration, refactoring, version update, and bug repair. The problem of code editing with manually created exemplars has been studied by Yin et al. (2019); Brody et al. (2020); Yao et al. (2021). Their works simulate a demand of automatically generalizing similar edits that a programmer has made somewhere to other applicable code snippets. Technically, their methods parse the source code as an abstract syntax tree using language-specific grammar, and apply one-time or sequential edit actions over the tree, and always keep the modified code snippet satisfying grammatical rules. In particular, all these works focus on performing code editing with exact one editing exemplar, i.e., the *one-shot* setting. The given exemplar showcases a certain editorial style, and implies the underlying intent of making a kind of edit, e.g., update some keywords or reformat redundant brackets.

While the above studies have demonstrated promising results, simply adapting the editorial style from one single exemplar can lead to poor generalization or even incorrect code editing. Figure 1 presents two cases. In Case (a), when a model observes only the Support #1 exemplar, it may intuitively interpret the code refactoring rule as replacing the conditional expression ("`if (s==null)...`") by a coalesce one ("`...??...;`") only for a newly-declared variable ("`string s=...`"). Consequently, applying the model to a query input could mistakenly transform the newly-declared variable ("`string title=...`") rather than its re-assigned correspondence ("`title=ReformatTitle(...)`"; see Solution #1). However, Case (b) demonstrates a situation where two editing exemplars turn to be sufficient for an editing model to capture the logical transformation over code tokens. The foregoing articulates the necessity of leveraging few but not one exemplar to adapt an editorial style to a new code snippet. The *few* exemplars help the editing escape from the inductive bias learned from one exemplar, and enhance the generalized capacity for other snippets with various coding contexts.

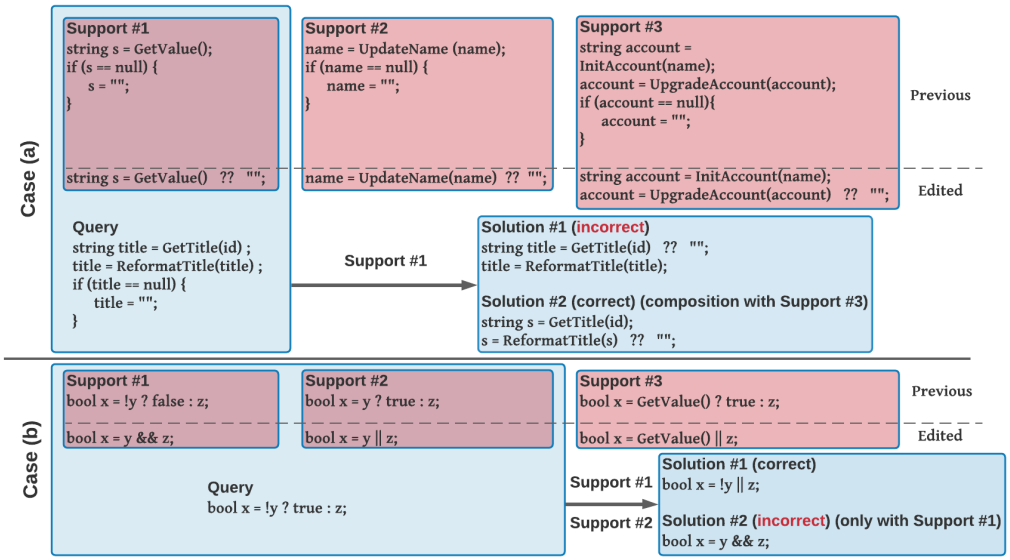

Figure 1: Two illustrative cases of our motivation for code editing learning with few exemplars, rather than with only one. The blue rectangles on the right part show the incorrect and correct solutions without or with multiple exemplars composition, respectively, in both Case (a) and (b). The editing with one exemplar may be dominated by the inductive bias learned from limited samples and thus lead to an infeasible solution. With few support exemplars, the desirable solution can be reached.

In light of this, in this paper, we consider the computer source code editing with few exemplars. The editing of a query code snippet is triggered by few editing exemplars that containing the original and modified support code snippets. Different from the above one-shot code editing scenario, few exemplars, even within the same editorial style, perform divergent coding contexts and edit actions. Therefore, it is not eligible to adapt the edit actions directly from an arbitrary exemplar. For this scenario, the major challenge lies in *how to identify and match helpful support exemplars for a query input and composite them to guide the editing*. To address this challenge, we propose a novel approach that combines edit representations extracted from exemplars and compositionally generalizes them to the query code snippet editing via multi-extent similarities ensemble. Specifically, we parse the support/query code snippets using language-specific grammar into abstract syntax trees, where editing actions are executed on nodes. Based on this, we model the similarity between the representations of the support and query abstract syntax trees. We notice that among few exemplars, some editing action only happens in a local region, while concurrently some editing action relies on a general coding context. This inspires us to consider multi-extent similarities between the representations of the support and query abstract syntax trees. We design a λ-softmax function by scaling the importance of nodes in an abstract syntax tree, i.e., $\lambda \to 0$ means the support/query code snippet is denoted by tree representation, while $\lambda \to \infty$ means the support/query code snippet is denoted by one single node representation in the abstract syntax tree. By controlling λ the intermediate positions between individual nodes and a collective tree, we provide multiple intermediate representational query-support matchings, then adaptively ensemble outcomes from these extents through a similarity-ranking error estimator for a robust composition. We use the term *multi-extent* to describe the scope of query-support matching at various levels from nodes to a collective tree representation.

We evaluate the algorithmic performance on two code editing datasets, one in C# (Yin et al., 2019) and one in Python (collected by us). We implement the proposed multi-extent exemplar composition mechanism on top of the state-of-the-art Graph2Edit model (Yao et al., 2021). On both datasets, our model outperforms baseline methods by 8.0-10.9% in terms of absolute accuracy. In addition, our experimental results show that (1) enabling few support exemplars can dramatically improve code editing; (2) compared with treating all support exemplars equally, precisely capturing the code snippet similarity and compositing support exemplars are crucial for code editing; and (3) the multi-extent approach offers better performance compared to its single-extent counterpart and other existing methods. Our collected *PyFixer* dataset and model implementation will be released.

2 PRELIMINARY

In what follows we shall use C_- and C_+ to represent the abstract syntax trees of the input (previous) and the output (edited) code snippets. We address the problem of code editing learning with few exemplars. Giving a set of K support exemplars with an identical editorial style $\{(C_-^{s_k}, C_+^{s_k})\}_{k=1}^K$, the goal is to adapt the editing from support set to a query snippet C_-^q and obtain its desirable C_+^q.

Here we introduce one of the state-of-the-arts edit learning models, Graph2Edit (Yao et al., 2021), that builds the one-shot code editing framework with three major components. (1) A tree encoder based on Gated Graph Neural Networks (Li et al., 2015) to embed the input abstract syntax tree C_- into node representations $Z \in \mathbb{R}^{N \times D}$, where N is the number of nodes in the tree and D is the latent feature dimension. The whole tree embedding t can be obtained by a graph pooling operation over all node representations; (2) An editing encoder to embed input C_- and corresponding output C_+ into a vector $f_\Delta(C_-, C_+) \in \mathbb{R}^D$ to represent the entire L sequential editing actions (four types of actions: add/delete a node, add a subtree, or stop editing) from C_- to C_+ as well as the procedural editorial style; (3) A decoder for editing action prediction working conditionally on the tree embedding t and editing representation $f_\Delta(C_-, C_+)$. Predictions include the operational type, the executive location on abstract syntax tree, and associated action values, e.g., to be replaced by which node. Denote $t_{1:l} = (t_1, \ldots, t_l)$ as the tree embedding history along the sequential editing, and $a_{1:l} = (a_1, \ldots, a_l)$ as the editing action history until step l. Formally, the whole predictions of actions can be represented as a conditional likelihood $\Pr(a_{1:L}|f_\Delta(C_-, C_+), t_1) = \prod_{l=1}^L \Pr(a_l|f_\Delta(C_-, C_+), t_{1:l})$. An earlier model Graph2Tree (Yin et al., 2019), where Graph2Edit derives from, shares the same methodological concept but edits the input in one-pass, instead of sequentially predict and apply the edit actions.

The Graph2Edit framework is trained in a self-reconstruction paradigm. The training feeds the tree encoder and the editing encoder with a pair of C_- and C_+, and maximizes the likelihood of the whole predictions of actions towards the shortest sequential actions produced by dynamic programming. The training details with the dynamic programming algorithm can be found in Yao et al. (2021). The training only sends the vanilla editing encoder with the ground-truth from one exemplar, while $\{(C_-^{s_k}, C_+^{s_k})\}_{k=1}^K$ with an identical editorial style are hard to collect from a wild code dataset.

Based on the training, Graph2Edit is mainly designed for code editing with only one support exemplar. It is struggled to directly tackle few supports since few exemplars, even within the same editorial style, perform divergent coding context and edit actions. Unfortunately, we argue the one-exemplar case lacks the capacity of compositional generalization on editorial styles as we discussed within the introductory part. We also empirically evidence this claim in our experimental section. Our method composites $f_\Delta(C_-^{s_k}, C_+^{s_k})$ from support set based on a multi-extent query-support matching to form the query edit representation f_Δ^q for decoding purpose which done by the code editing frameworks. In the next section, we elaborate on our proposed adaptive multi-extent similarity ensemble on exemplar composition to address the foregoing challenges in computer code editing with few exemplars.

3 ADAPTIVE MULTI-EXTENT COMPOSITION

We assume that if the input tree of a support exemplar enjoys more similarity with the query input tree, the query snippet is more likely to adopt the support exemplar's edit representation towards correct code editing. Based on this, we address compositional learning from a query-support matching perspective. Deriving from the basis of code editing reference, some editing can be adapted depending on individual node representation and regardless of its neighbor nodes and context. For instance, the removal of a redundant bracket should partially be independent of the inside context, or the change of some outdated function calling mode in a previous language version should be invariant to the host object. On the other hand, some editing highly relies on the contextual information among several objects and connections therein, leading the editing learning to be determined by a collective representation. Continuing with this analysis, we design a multi-extent measurement to perform multiple intermediate query-support matching between individual nodes and the collective tree. Specifically, we measure the multi-extent similarities between query and support snippets in the tree embedding space, and learn a convex combination of the editing representations from the support set to maximize the adaptability for query snippet editing. In the coming subsections, we introduce the query-support matching over abstract syntax tree, extend it with multi-extent representation in an ensemble mechanism, then illustrate the meta-learning paradigm and inference procedure.

3.1 QUERY-SUPPORT MATCHING: FROM COLLECTIVE TREE TO INDIVIDUAL NODE

Instead of treating each node equally, our model softly emphasizes some nodes in the query snippet when meet a similar node in a support exemplar, and vice versa. Considering one query snippet and a set with K support exemplars, let z_n^q and $z_n^{s_k}$ denote the n-th node representation in the abstract syntax tree of the query and k-th support snippet, respectively. Let $\varphi_\theta(\cdot,\cdot) : \mathcal{Z} \times \mathcal{Z} \to \mathbb{R}$ be a query-support node matching similarity function (e.g., a neural network) with a learnable parameter θ that takes the node representation from both sides as input. Note that $\varphi_\theta(\cdot,\cdot)$ does not have to be non-negative or symmetrical since the measurement across query and support can be deemed as directional. We compute the initial query-support node matching activation $m_n^{qs_k}$ and $m_n^{s_k}$ as follows:

$$m_n^{q,s_k} := \max\left\{\varphi_\theta\big(z_n^q, z_i^{s_k}\big)\right\}_{i=1}^{N_s} \text{ and } m_n^{s_k,q} := \max\left\{\varphi_\theta\big(z_i^q, z_n^{s_k}\big)\right\}_{i=1}^{N_q}, \tag{1}$$

where N_q and N_s are the numbers of nodes in the query and support abstract syntax trees. The maximum operation over the set of one-to-all node emphasizes the matching between individual nodes, and returns a high activation when there is at least one good match in the counterpart. The activation represents the likelihood that a query node appears in a support snippet and vice versa.

To capture the multi-extent matching, we design a λ-softmax function by scaling the importance of nodes in an abstract syntax tree, and seek the query-support matching at the intermediate position between individual nodes and a collective tree:

$$\sigma_{n,\lambda}^{q,s_k} := \frac{N_q \exp\left(\lambda m_n^{qs_k}\right)}{\sum_{i=1}^{N_q} \exp\left(\lambda m_i^{qs_k}\right)} \text{ and } \sigma_{n,\lambda}^{s_k,q} := \frac{N_{s_k} \exp\left(\lambda m_n^{s_k}\right)}{\sum_{i=1}^{N_{s_k}} \exp\left(\lambda m_i^{s_k}\right)}. \tag{2}$$

$\sigma_{n,\lambda}^{qs_k}$ denotes the parameterized node activation from the query snippet to the k-th support snippet, and a similar expression can denote as $\sigma_{n,\lambda}^{s_k}$. Based on the above normalized activation for individual node representation, the collective tree representation of the query and support snippets can be calculated using a weighted average pooling as follows:

$$t_\lambda^q := \frac{1}{N_q K} \sum_{n=1}^{N_q} \sum_{k=1}^{K} \sigma_{n,\lambda}^{q,s_k} z_n^q \text{ and } t_\lambda^{s_k} := \frac{1}{N_{s_k}} \sum_{n=1}^{N_{s_k}} \sigma_{n,\lambda}^{s_k,q} z_n^{s_k}. \tag{3}$$

The representation of the query snippet has the summation of matching values on nodes from all support exemplars. Combining Eq. (2) and (3), an intuitive interpretation on variable λ raises: a larger λ denotes greater domination of matched individual nodes in the final tree representation, i.e., the sharpness of the outcomes after λ-softmax normalization. $\lambda \to 0$ preserves the final query/support representation as to their initial tree representation, and $\lambda \to \infty$ represents the final query/support representation approximately with only one single node representation at most of the time, where the node is selected upon the maximum activation from Eq. (1). This property holds due to the monotonous increase of the first-order gradient of the exponential function. We will tackle the λ value selection problem in the next subsection. We reach to the generation of edit representation for query sample editing and the updated tree representation yields the expression:

$$f_\Delta^q := \sum_{k=1}^{K} \phi_\theta(t_\lambda^q, t_\lambda^{s_k}) f_\Delta(C_-^{s_k}, C_+^{s_k}), \tag{4}$$

where $\phi_\theta(\cdot,\cdot)$ is a similarity measure over query and support tree representations and should satisfy the convexity in the above combination. Note that for a simplified expression, we do not distinguish the learnable parameter θ for each module but their parameters are isolated.

3.2 MULTI-EXTENT COMPOSITION AND ENSEMBLE

The λ-softmax in Eq. (2) controls the matching extent from individual node to collective tree. However, it is hard to access the optimal intermediate position, since code editing samples suffer from a huge variance. Therefore, a single extent may not be robust enough to help the model generalize in diverse editing scenarios. We hereby propose a multi-extent measurement to complement the above weakness by setting λ with different continuous values, and ensemble all the results $\phi_\theta(t_{\lambda_i}^q, t_{\lambda_i}^{s_k})$ for $\{\lambda_i\}_{i=1}^{N_\lambda}$ to enhance the generalization on complicated code snippets. However, setting λ arbitrarily

and ensemble all perspectives linearly without selection may corrupt the overall tree representation, and thus deteriorate the model. For a robust ensemble and aggregation, we consider assessing the quality of $\phi_\theta(t^q_{\lambda_i}, t^s_{\lambda_i})$ towards different λ_i by using a marginal ranking error on edit representation. Since we have the accessible ground-truth C^q_+ during training, we employ the ranking error e_{λ_i} to explicitly reflect how good $\phi_\theta(t^q_\lambda, t^{s_k}_\lambda)$ is under a certain λ with respect to the similarity between query and K support exemplars. In the inference phase, due to the missing of C^q_+, we employ a similarity-ranking error estimator $\hat{e}_{\lambda_i} = R_\mu(t^q_{\lambda_i}, t^{s_1}_{\lambda_i}, t^{s_2}_{\lambda_i}, \ldots, t^{s_K}_{\lambda_i}; \lambda_i) : \mathbb{R}^{(K+1)\times D} \to \mathbb{R}$ to predict the quality of each λ_i, where μ is the learnable parameter in R_μ, and ensemble the outcomes. Training details of R_μ will be elaborated in the next subsection.

To achieve multi-extent composition and ensemble, we first measure the similarity between query and support edit representations via the editing encoder:

$$w_k := \frac{\left\langle f_\Delta(C^q_-, C^q_+), f_\Delta(C^{s_k}_-, C^{s_k}_+)\right\rangle}{\left\|f_\Delta(C^q_-, C^q_+)\right\| \cdot \left\|f_\Delta(C^{s_k}_-, C^{s_k}_+)\right\|}. \tag{5}$$

Then we access the error for each λ by:

$$e_\lambda := \sum_{k=1}^{K} \left(w_{\Gamma(1)} - w_{\Gamma(k)}\right) \cdot \max\left\{0, \gamma - \left(\phi_\theta(t^q_\lambda, t^{s_{\Gamma(1)}}_\lambda) - \phi_\theta(t^q_\lambda, t^{s_{\Gamma(k)}}_\lambda)\right)\right\}, \tag{6}$$

where $\Gamma(\cdot)$ is the index mapping for K support exemplars that satisfy $w_{\Gamma(1)} \geq w_{\Gamma(2)} \geq \ldots \geq w_{\Gamma(k)}$, i.e., the mapping to sort s_k in a descending order, and γ is the margin set as a hyperparameter. The formulation indicates how well the nearest support exemplars can outperform other support exemplars, and the factor $(w_{\Gamma(1)} - w_{\Gamma(k)})$ reveals the error confidence. Note that the above error term depends on the edit model f_Δ, which is only an approximate estimation. Then we extend Eq. (4) by involving multiple extents and the reciprocal of its error for an adaptive ensemble:

$$f^q_\Delta := \sum_{k=1}^{K}\sum_{i=1}^{N_\lambda} \frac{1}{e_{\lambda_i} + \epsilon} \phi_\theta(t^q_{\lambda_i}, t^{s_k}_{\lambda_i}) f_\Delta(C^{s_k}_-, C^{s_k}_+). \tag{7}$$

where ϵ is a constant in denominator for numerical stability. We omit the normalization over $\phi_\theta(t^q_{\lambda_i}, t^{s_k}_{\lambda_i})/(e_{\lambda_i} + \epsilon)$ in the equation. In practice, we apply a softmax function on this term to keep the convexity of the combination. Note that in inference e_{λ_i} will be replaced by $\hat{e}_{\lambda_i}$.

3.3 Learning and Inference

In the learning phase, we have a set of K support exemplars $\{(C^{s_k}_-, C^{s_k}_+)\}_{k=1}^K$ and one complete query snippet $\{C^q_-, C^q_+\}$ as inputs. Following the Graph2Edit framework (Yao et al., 2021), we obtain the z^q_n and $z^{s_k}_n$ of the n-th node representation of the query and k-th support snippet and the edit representation $f_\Delta(C^{s_k}_-, C^{s_k}_+)$ for each support exemplar. L^q is the number of steps towards the shortest sequential actions from C^q_- to C^q_+, and $a_{1:L^q}$ denotes the corresponding action sequence. Let $\mathbf{e} = (e_{\lambda_1}, e_{\lambda_2}, \ldots, e_{\lambda_N})$ calculated by Eq. (6) and $\hat{\mathbf{e}} = (\hat{e}_{\lambda_1}, \hat{e}_{\lambda_2}, \ldots, \hat{e}_{\lambda_N})$ is its estimation from the estimator R_μ. With N_λ extents, our objective function can be written as:

$$\min_{\theta,\mu} -\Pr\left(a_{1:L^q} | f^q_\Delta, t^q\right) + ||\mathbf{e}||_1 + \mathrm{D}_{\mathrm{KL}}(\hat{\mathbf{e}} \| \mathbf{e}), \tag{8}$$

where $||\cdot||_1$ is the vector l_1-norm to sum up the absolute values of all elements, $\mathrm{D}_{\mathrm{KL}}(\cdot\|\cdot)$ is the KL-divergence, f^q_Δ is the composition of edit representations in Eq. (7), $t^q = \sum_{n=1}^{N_q} z^q_n / N_q$ is the initial tree representation of C^q_-, θ and μ are the learnable parameters in the similarity measurement ϕ_θ, φ_θ and the estimator R_μ. We omit the balanced factor for each term for a simplified expression. The above objective function consists of three parts. The first term maximizes the likelihood between the learned convex combination of edit representations and optimal editing actions, while the rest two terms aim to minimize the ranking error between the query and support snippets at different extent levels and train an estimator to accurately estimate the error for inference usage. The parameters for tree encoder, editing encoder, and decoder are obtained via a pre-training stage on another data split following the self-reconstruction paradigm from Graph2Edit, while their parameters are not tuned in the meta-learning stage for optimization, since we observe a serve overfitting if doing so.

Table 1: Experimental results of code editing with 5 support exemplars on *C#Fixer* dataset

Macro Accuracy	Split #1	Split #2	Split #3	Split #4	Split #5	Avg.
Graph2Tree-RS	0.270 $\pm$ 0.027	0.360 $\pm$ 0.011	0.341 $\pm$ 0.029	0.243 $\pm$ 0.009	0.384 $\pm$ 0.017	0.320
Graph2Tree-AE	0.381 $\pm$ 0.022	0.363 $\pm$ 0.008	0.372 $\pm$ 0.011	0.275 $\pm$ 0.014	0.349 $\pm$ 0.028	0.348
Graph2Edit-RS	0.279 $\pm$ 0.022	0.423 $\pm$ 0.008	0.413 $\pm$ 0.018	0.225 $\pm$ 0.007	0.396 $\pm$ 0.017	0.347
Graph2Edit-AE	0.336 $\pm$ 0.025	0.471 $\pm$ 0.012	0.465 $\pm$ 0.022	0.267 $\pm$ 0.018	0.402 $\pm$ 0.023	0.388
Ours (GED)	0.363 $\pm$ 0.019	0.479 $\pm$ 0.012	0.487 $\pm$ 0.021	0.302 $\pm$ 0.012	0.457 $\pm$ 0.026	0.418
Ours (CS)	0.387 $\pm$ 0.016	0.500 $\pm$ 0.004	0.501 $\pm$ 0.016	0.337 $\pm$ 0.008	0.507 $\pm$ 0.009	0.447
Ours	**0.416** $\pm$ 0.015	**0.514** $\pm$ 0.021	**0.522** $\pm$ 0.015	**0.352** $\pm$ 0.018	**0.539** $\pm$ 0.023	**0.468**
Micro Accuracy	**Split #1**	**Split #2**	**Split #3**	**Split #4**	**Split #5**	**Avg.**
Graph2Tree-RS	0.290 $\pm$ 0.024	0.490 $\pm$ 0.008	0.332 $\pm$ 0.032	0.166 $\pm$ 0.008	0.619 $\pm$ 0.010	0.380
Graph2Tree-AE	0.391 $\pm$ 0.017	0.546 $\pm$ 0.006	0.353 $\pm$ 0.013	0.177 $\pm$ 0.005	**0.673** $\pm$ 0.010	0.428
Graph2Edit-RS	0.281 $\pm$ 0.018	0.538 $\pm$ 0.010	0.417 $\pm$ 0.022	0.167 $\pm$ 0.010	0.599 $\pm$ 0.010	0.400
Graph2Edit-AE	0.343 $\pm$ 0.019	0.603 $\pm$ 0.011	0.460 $\pm$ 0.025	0.184 $\pm$ 0.012	0.590 $\pm$ 0.020	0.436
Ours (GED)	0.367 $\pm$ 0.014	0.607 $\pm$ 0.014	0.485 $\pm$ 0.024	0.200 $\pm$ 0.008	0.606 $\pm$ 0.022	0.453
Ours (CS)	0.388 $\pm$ 0.010	0.616 $\pm$ 0.008	0.500 $\pm$ 0.018	0.209 $\pm$ 0.003	0.625 $\pm$ 0.015	0.467
Ours	**0.411** $\pm$ 0.011	**0.636** $\pm$ 0.030	**0.524** $\pm$ 0.017	**0.218** $\pm$ 0.011	0.653 $\pm$ 0.039	**0.488**

We follow an episode training mechanism (Snell et al., 2017) that is widely used in few-shot learning to simulate the encounter of new editorial style in testing. For each training forward, we sample a fixed number of editorial styles with K support exemplars and one query snippet per class and backpropagate the loss. In the inference phase, given a set of K support exemplars $\{(C_-^{s_k}, C_+^{s_k})\}_{k=1}^K$ and one input query snippet C_-^q, we replace e_{λ_i} with $\hat{e}_{\lambda_i}$ in Eq. (7) and yields the formulation.

4 EXPERIMENT

4.1 EXPERIMENTAL SETTING

Dataset We include two code editing datasets for performance evaluation. (1) *C#Fixer* (Yin et al., 2019) is a dataset containing 2,878 editing pairs, generated by applying two C# fixing tools Roslyn[1] and Roslynator[2] on six projects. A fixer is building on top of the C# compiler used to perform refactoring and modernization on codes. One fixer is designed for one type of refactoring and 16 fixers are selected in this dataset. A pre-trained model for *C#Fixer* is trained on *GitHubEdits* dataset using self-reconstruction. The dataset collects 54 C# projects on GitHub which extract the source code before and after the commits from users. We use the data of *GitHubEdits* only for pre-training and follow the training protocol in Yao et al. (2021) but not for few-shot experiments. (2) Benefiting from the flexibility and generalized capacity of abstract syntax description language, beyond C#, we collect another dataset *PyFixer* written in Python 2. We use Python-Future[3], a tool that edits the codes over 5,959 projects written in Python 2 to satisfy the compatibility of Python 3. We obtain the code pairs by taking codes before and after applying the tool and inspecting the changes. To simulate the pre-training and few-shot learning as in *GitHubEdits* and *C#Fixer*, we select 12 types of fixers and extract a fixed number of each fixer from the main data to form the dataset for few-shot learning, while the rest are preserved for pre-training. After data cleaning, we have 14,941 code pairs for (pre) training, 1,664 for validation, and 2,319 for few-shot experiments. The two datasets only contain source code but without any author's name, contact information, or offensive content.

Baseline We employ modifications on two abstract syntax tree-based code editing frameworks, Graph2Tree (Yin et al., 2019) and Graph2Edit (Yao et al., 2021), to fit them in the few exemplars editing setting and involve the following baselines. RS (Random Selection): we randomly select one support exemplar from the support set and use its edit representation for standard decoding on query snippet; AE (Average Edit Representation): we infer the edit representation of K support

[1]https://github.com/dotnet/roslyn

[2]https://github.com/josefpihrt/roslynator

[3]https://python-future.org/index.html

Table 2: Experimental results of code editing with 5 support exemplars on *PyFixer* dataset. Macro and micro accuracy are integrated since these two are nearly identical on this dataset.

Macro/Micro Accuracy	Split #1	Split #2	Split #3	Split #4	Split #5	Avg.
Graph2Edit-RS	0.408 $\pm$ 0.017	0.381 $\pm$ 0.015	0.498 $\pm$ 0.033	0.293 $\pm$ 0.007	0.206 $\pm$ 0.006	0.357
Graph2Edit-AE	0.297 $\pm$ 0.012	0.347 $\pm$ 0.008	0.488 $\pm$ 0.010	0.276 $\pm$ 0.008	0.194 $\pm$ 0.009	0.320
Ours (GED)	0.337 $\pm$ 0.014	0.368 $\pm$ 0.007	0.509 $\pm$ 0.012	0.297 $\pm$ 0.013	0.219 $\pm$ 0.012	0.346
Ours (CS)	0.441 $\pm$ 0.020	0.409 $\pm$ 0.011	0.543 $\pm$ 0.011	0.347 $\pm$ 0.010	0.276 $\pm$ 0.014	0.403
Ours	**0.510** $\pm$ 0.036	**0.429** $\pm$ 0.021	**0.544** $\pm$ 0.022	**0.369** $\pm$ 0.021	**0.294** $\pm$ 0.018	**0.429**

exemplars using editing encoder respectively, and take the weighted mean of the K representations for query decoding. Moreover, we include two other baselines within our proposed composition framework. Our (GDE) means that we employ the distance metric Graph Edit Distance (Sanfeliu & Fu, 1983; Abu-Aisheh et al., 2015) among every query-support snippet pair, take the reciprocal proportion, and form a convex combination over K edit representations, and Ours (CS) means that we employ mean pooling over query and support trees, compute the similarity by cosine similarity for every query-support snippet pair, normalize, and form a convex combination over entire available edit representations from support set based on the obtained similarity. Note that we find Graph2Tree contains numerous C# specific implementations and currently not applicable for Python, so we do not include this baseline on *PyFixer*. All the baselines are under MIT license for public usage.

Protocol We randomly split fixers for few-shot experiments into meta-train and meta-test sets. For *C#Fixer*, the meta-train set contains 12 types of fixers and the meta-test set has the rest 4 fixers. The meta-train set of *PyFixer* contains 8 fixers and the meta-test set takes the rest 4. To perform a comprehensive empirical evaluation, we repeat the random split procedure and obtain 5 different splits denoted from Split #1 to Split #5. For meta-training, to simulate the low-resource scenario and mitigate the imbalance on the number of samples per fixer, we take 10 samples for each fixer and follow a standard episode training strategy. For meta-testing, we iterate through all samples and each of them as a query snippet at a time, then randomly sample from the rest of its fixer peer to build its support set. All numerical results are run with 5 different random seeds. Mean and standard deviation are reported. Note that the variance not only stems from the training per se, but also is affected by the random sampling of support set in testing. Since we fix the random seed every time, all methods are guaranteed to encounter identical query-support combinations in the evaluation under the same seed. All configurable hyperparameters are tuned on a validation split outside the five evaluative splits. We defer the model architecture and configurable hyperparameters to supplementary material. The accuracy per pair is a binary value obtained by whether the model edits C_- to meet exact ground-truth C_+. We report both Macro and Micro accuracy since there is a significant fixer imbalance in *C#Fixer*. Experiments are run on a single NVIDIA RTX 3090 graphical card using Pytorch.

4.2 Result and Analysis

Main numerical results of code editing with 5 support exemplars on *C#Fixer* and *PyFixer* are presented in Table 1 and Table 2. Comparing to baseline methods, we nearly achieve the best results on all splits of two datasets, across macro and micro accuracy. We outperform the second-best baseline method by bringing a absolute improvement range from 8.0% (Macro accuracy on *C#Fixer*) to 10.9% (*PyFixer*). With the same base architecture (exclude Graph2Tree), comparing to GED and CS, our improvements show the significance of multi-extent composition learning of edit representations from support set. Quantitative results demonstrate the effectiveness of treating nodes in abstract syntax tree differently and ensemble learning by the meta-learning strategy. In the following, we present two extra experiments and a visualization case to further investigate our approach.

More Support Exemplars Figure 2 shows the performance trends when enlarging the size of support set, i.e., offering more resources for code editing. An estimated regression with uncertainty is plotted from $K = 3$ to $K = 7$. It is relatively straightforward that the performance on random selection is not affected by the scale of support set. The average edit representation performs differently on two datasets. Adding more exemplars into support set can help to improve the robustness of the weighted mean representation, or lead it to deviate the representational manifold if high variances exist among exemplars. Our approach consistently benefits from more support

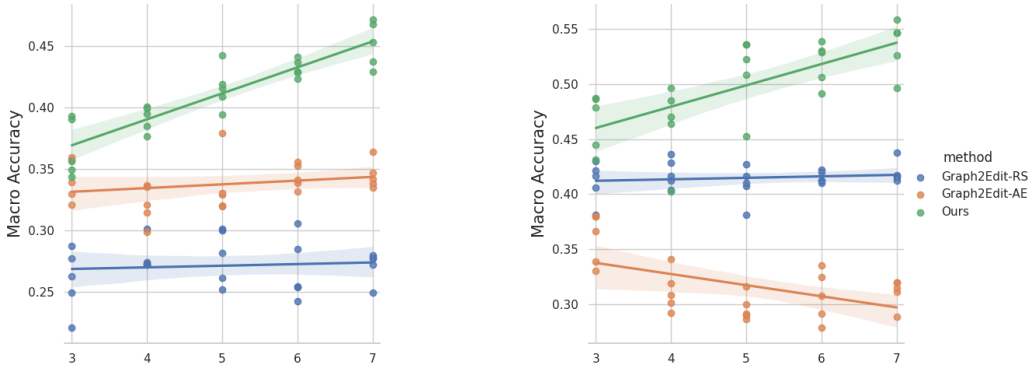

Figure 2: Performance of code editing with various K support exemplars, where the integer $K \in [3, 7]$ along the x-axis. Left is on *C#Fixer* dataset Split #1, and right is on *PyFixer* dataset Split #1. An estimated regression model is plotted for every group of data by the seaborn package, showing as the straight line with deviations. Our approach consistently benefit from adding more support exemplars and is capable to find a better composition, while the other two baselines do not enjoy that.

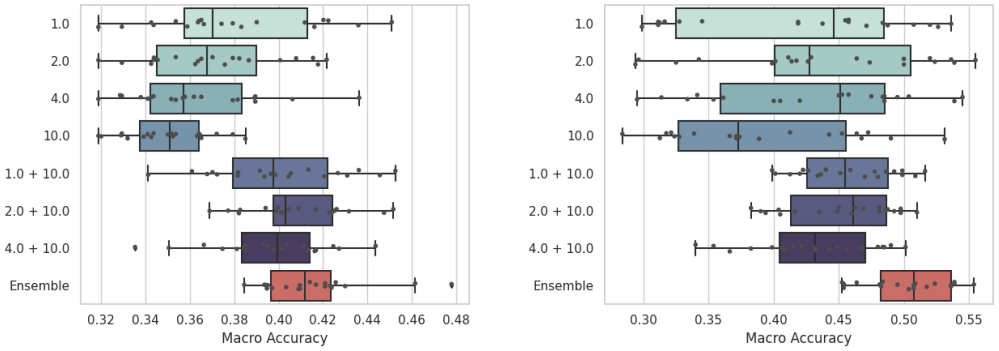

Figure 3: Experiments on multi-extent parameter λ and its ensemble. The y-axis represents the instantiated value of λ, where the plus symbol denotes the ensemble of the two extents, and 'Ensemble' includes all these four. Left: *C#Fixer* dataset Split #1; Right: *PyFixer* dataset Split #1.

exemplars and gradually enlarging the margin towards the two baselines. This demonstrates our model is capable to find a better composition when there are more options, therefore preserving the possibility of capability enhancement when more resources are available.

Investigation of Multi-Extent Ensemble Figure 3 shows the functional analyses on multi-extent measurement and its ensemble, where we investigate seven possible options on extent parameter λ, four in a single extent, three in a double extent ensemble, and compare to the full model with four levels of extents plotting in red. A single extent in the first four rows underperforms the ensemble when using two or four extents. This confirms our methodological hypothesis that a single extent may not be enough to express the similarity between the query-support snippet pair due to the high variances and that ensembling multiple extents helps the model robustness. The full model with four extents also performs better than the double extent ensemble shown by the median line of the box plot. However, adding more extents implies increasing computational costs and hence a trade-off.

Visual Diagram for Node Activation on Query-Support Matching We visualize the node activation in Figure 4. The value of an activation is normalized within the query or support snippet to show the relative magnitude of the value. Matching on individual nodes facilitates query snippet to find the valuable reference in Support #1 under 'BinaryExpression' regardless of the disturbing content under 'ConditionalExpression.' Also, there is a faint activation on the multiple operators in Support #2 since the objects on this operation is under 'MemberAccessExpression.'

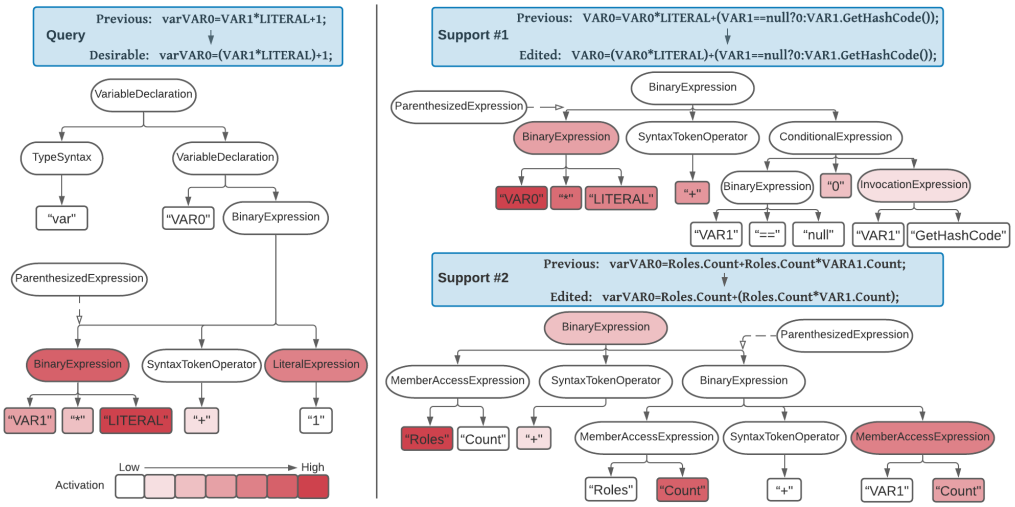

Figure 4: Visualization of normalized node activation with $\lambda = 1.0$ on query snippet and two support exemplars. We plot the top six nodes with the highest activation for each tree and show them in decreasing color saturation. Abstract syntax trees are plotted after pruning on insignificant parts.

5 Related Work

Machine Learning for Code Editing Code editing using machine learning is receiving increasing attention in recent years. Depending on how a code snippet is represented, the research can be summarized into two categories, i.e., performing code editing over code token sequences or abstract syntax tree. In the former category, Yin et al. (2019); Chen et al. (2019); Yasunaga & Liang (2020) directly generate the expected edited code token sequence, while Shin et al. (2018); Vu & Haffari (2018); Dong et al. (2019); Zhao et al. (2019) predict the editing operational sequence, which has been demonstrated to bear more sample efficiency. On the other hand, works in the second category (Yin et al., 2019; Chakraborty et al., 2020; Tarlow et al., 2020; Dinella et al., 2020; Brody et al., 2020; Yao et al., 2021) attempt to edit the trees of programs similarly by either directly generating the edited trees or predicting the tree edit operations. Our work follows the scenario of Yin et al. (2019); Yao et al. (2021). Nonetheless, both of the two works have focused on code editing from one shot of editing exemplar, which may not be sufficient for learning generalizable editing, as we elaborated in Section 1. This has inspired us to investigate the few-shot code editing problem.

Few-shot Learning Few-shot learning has been developed in many low-resource scenes (Vinyals et al., 2016; Snell et al., 2017; Finn et al., 2017), and with advance in meta seq2seq learning considering the concept of generalized composition (Gu et al., 2018; Lake, 2019; Nye et al., 2020). Our work conceptually relates to the matching principle in few-shot learning (Vinyals et al., 2016; Snell et al., 2017) but with a different compositional task and a specific design for robust multi-extent abstract syntax tree matching. To our best knowledge, we are the first to consider few-shot compositional generalization in code editing, which we believe can inspire technological development in programming-related applications like code transfer, refactoring, and migration.

6 Conclusion

In this work, we considered a novel problem in software engineering: code editing with few exemplars. Based on previous frameworks on code editing using abstract syntax tree to present the computer code snippets, we proposed an adaptive multi-extent composition method to perform varying intermediate representations between a collective tree and individual nodes. We leveraged an ensemble approach to gather the query-support matching from multiple extents and delivered a robust composition over support editing representations for the query snippet decoding. Evaluations on two code editing datasets demonstrated the effectiveness of our method over baselines by a large margin.

7 ETHICS STATEMENT

We discuss the potential ethical considerations in this section. We do not discover a noteworthy negative societal impact originate from our work. One minor concern is the research in machine learning for source code may encourage and facilitate the code data collection from open-source projects. The data may unintentionally involve user privacy like the author's name, organization, or file date, and the collection process may violate the usage license.

REFERENCES

Zeina Abu-Aisheh, Romain Raveaux, Jean-Yves Ramel, and Patrick Martineau. An exact graph edit distance algorithm for solving pattern recognition problems. In *Proceedings of the International Conference on Pattern Recognition Applications and Methods*, 2015.

Shaked Brody, Uri Alon, and Eran Yahav. A structural model for contextual code changes. *Proceedings of the ACM on Programming Languages*, 2020.

Saikat Chakraborty, Yangruibo Ding, Miltiadis Allamanis, and Baishakhi Ray. Codit: Code editing with tree-based neural models. *IEEE Transactions on Software Engineering*, 2020.

Zimin Chen, Steve James Kommrusch, Michele Tufano, Louis-Noël Pouchet, Denys Poshyvanyk, and Martin Monperrus. Sequencer: Sequence-to-sequence learning for end-to-end program repair. *IEEE Transactions on Software Engineering*, 2019.

Elizabeth Dinella, Hanjun Dai, Ziyang Li, Mayur Naik, Le Song, and Ke Wang. Hoppity: Learning graph transformations to detect and fix bugs in programs. In *International Conference on Learning Representations*, 2020.

Yue Dong, Zichao Li, Mehdi Rezagholizadeh, and Jackie Chi Kit Cheung. Editnts: An neural programmer-interpreter model for sentence simplification through explicit editing. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, 2019.

Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In *International Conference on Machine Learning*, 2017.

Jiatao Gu, Yong Wang, Yun Chen, Victor OK Li, and Kyunghyun Cho. Meta-learning for low-resource neural machine translation. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, 2018.

Brenden M Lake. Compositional generalization through meta sequence-to-sequence learning. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett (eds.), *Advances in Neural Information Processing Systems*, 2019.

Yujia Li, Daniel Tarlow, Marc Brockschmidt, and Richard Zemel. Gated graph sequence neural networks. *arXiv preprint arXiv:1511.05493*, 2015.

Maxwell I Nye, Armando Solar-Lezama, Joshua B Tenenbaum, and Brenden M Lake. Learning compositional rules via neural program synthesis. *arXiv preprint arXiv:2003.05562*, 2020.

Alberto Sanfeliu and King-Sun Fu. A distance measure between attributed relational graphs for pattern recognition. *IEEE Transactions on Systems, Man, and Cybernetics*, 1983.

Richard Shin, Illia Polosukhin, and Dawn Song. Towards specification-directed program repair, 2018.

Jake Snell, Kevin Swersky, and Richard Zemel. Prototypical networks for few-shot learning. In *Advances in Neural Information Processing Systems*, 2017.

Daniel Tarlow, Subhodeep Moitra, Andrew Rice, Zimin Chen, Pierre-Antoine Manzagol, Charles Sutton, and Edward Aftandilian. Learning to fix build errors with graph2diff neural networks. In *Proceedings of the IEEE/ACM 42nd International Conference on Software Engineering Workshops*, 2020.

Oriol Vinyals, Charles Blundell, Timothy Lillicrap, Koray Kavukcuoglu, and Daan Wierstra. Matching networks for one shot learning. *arXiv preprint arXiv:1606.04080*, 2016.

Thuy Vu and Gholamreza Haffari. Automatic post-editing of machine translation: A neural programmer-interpreter approach. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, 2018.

Ziyu Yao, Frank F. Xu, Pengcheng Yin, Huan Sun, and Graham Neubig. Learning structural edits via incremental tree transformations. In *International Conference on Learning Representations*, 2021.

Michihiro Yasunaga and Percy Liang. Graph-based, self-supervised program repair from diagnostic feedback. In *International Conference on Machine Learning*, 2020.

Pengcheng Yin, Graham Neubig, Miltiadis Allamanis, Marc Brockschmidt, and Alexander L. Gaunt. Learning to represent edits. In *International Conference on Learning Representations*, 2019.

Rui Zhao, David Bieber, Kevin Swersky, and Daniel Tarlow. Neural networks for modeling source code edits. *arXiv preprint arXiv:1904.02818*, 2019.

A EXPERIMENTAL DETAILS

We collect the *PyFixer* dataset by separately applying 51 fixers in Python-Future[4] over 5,959 projects written in Python 2 to satisfy the compatibility of Python 3. We obtain the code pairs by inspecting the changed lines and get the source code before and after the refactoring. We select 12 types of fixers and extract 200 samples for each fixer from the main data to form the dataset for few-shot learning, while the rest are preserved for model pre-training.

For training, we set the batch size to 16, epochs to 15, learning rate to 1e-4, with a gradient accumulation for every 2 optimized steps. The extent parameter λ for main experiments are set as an ensemble for 10, 5, 2, and 1. The architectural parameters are set the same as Graph2Edit (Yao et al., 2021). The predictor R is implemented with a multilayer perceptron. Other detailed implementation or configurations are in our source code since it is not possible to exhasutively include every detail with plain text.

The *PyFixer* are working under Python 3.8, having the abstract syntax tree grammar from https://docs.python.org/3.8/library/ast.html. We list the choice of splits for the two datasets below.

Table 3: Fixer Splits on *C#Fixer*

Split #1	CA2007, IDE0004, RCS1015, RCS1021, RCS1032, RCS1058, RCS1077, RCS1097, RCS1118, RCS1123, RCS1197, RCS1206	RCS1146, RCS1207, RCS1202, RCS1089
Split #2	IDE0004, RCS1015, RCS1032, RCS1058, RCS1077, CA2007, RCS1089, RCS1146, RCS1202, RCS1206, RCS1207, RCS1097	RCS1118, RCS1123, RCS1021, RCS1197
Split #3	RCS1015, RCS1021, RCS1032, RCS1058, RCS1077, RCS1097, CA2007, IDE0004, RCS1118, RCS1146, RCS1202, RCS1207	RCS1123, RCS1197, RCS1206, RCS1089
Split #4	RCS1123, RCS1021, RCS1032, RCS1058, RCS1206, RCS1097, CA2007, IDE0004, RCS1118, RCS1146, RCS1202, RCS1207	RCS1015, RCS1197, RCS1077, RCS1089
Split #5	RCS1077, RCS1021, RCS1032, RCS1058, RCS1206, RCS1097, CA2007, IDE0004, RCS1197, RCS1146, RCS1089, RCS1207	RCS1118, RCS1015, RCS1123, RCS1202

Table 4: Fixer Split on *PyFixer*

Split #1	08, 15, 18, 23, 34, 42, 47, 49	05, 41, 09, 35
Split #2	08, 09, 35, 23, 34, 42, 47, 05	41, 49, 15, 18
Split #3	41, 09, 15, 23, 49, 42, 47, 05	08, 18, 34, 35
Split #4	41, 09, 34, 23, 49, 42, 08, 05	47, 15, 18, 35
Split #5	15, 09, 34, 23, 49, 35, 08, 47	05, 41, 18, 42

We use the fixer's name in Table 3 and the fixer's index in Table 4 corresponding to the order in the Python-Future package. For examples on fixers, C# fixers are included in (Yao et al., 2021), and we provide examples from *PyFixer* in what follows.

B EDITING EXAMPLES FOR FIXERS IN *PyFixer*

We list the 12 kinds of fixers with names in Table 5.

[4]https://python-future.org/index.html

Table 5: Fixer examples in *PyFixer*

	Fixer 05: lib2to3.fixes.fix_filter
C_-	chunk = filter(lambda x: x.feature == feature, ichunk)
C_+	chunk = [x for x in ichunk if x.feature == feature]
	Fixer 08: lib2to3.fixes.fix_has_key
C_-	key = request.matchdict["key"] if request.matchdict.has_key("key") else ""
C_+	key = request.matchdict["key"] if "key" in request.matchdict else ""
	Fixer 09: lib2to3.fixes.fix_idioms
C_-	return type(self) == type(other) and self.__dict__ == other.__dict__
C_+	return isinstance(self, type(other)) and self.__dict__ == other.__dict__
	Fixer 15: lib2to3.fixes.fix_map
C_-	map(int, [a for a in kwargs.values() if a != kwargs['releaselevel']])
C_+	list(map(int, [a for a in kwargs.values() if a != kwargs['releaselevel']]))
	Fixer 18: lib2to3.fixes.fix_next
C_-	second_gff_chunk = second_gff.next()
C_+	second_gff_chunk = next(second_gff)
	Fixer 23: lib2to3.fixes.fix_raw_input
C_-	raw_input("Start the server on 0:1 and press enter.".format(str(parsed.ip), str(parsed.p)))
C_+	input("Start the server on 0:1 and press enter.".format(str(parsed.ip), str(parsed.p)))
	Fixer 34: lib2to3.fixes.fix_zip
C_-	context.update(zip(inlines_names, kwargs.get('inlines', [])))
C_+	context.update(list(zip(inlines_names, kwargs.get('inlines', []))))
	Fixer 34: lib2to3.fixes.fix_absolute_import
C_-	from elks import Elks
C_+	from .elks import Elks
	Fixer 41: lib2to3.fixes.fix_future_standard_library
C_-	import urlparse
C_+	import urllib.parse
	Fixer 42: lib2to3.fixes.fix_future_standard_library_urllib
C_-	from urllib2 import HTTPError, urlopen, Request
C_+	from urllib.request import urlopen, Request
	Fixer 42: lib2to3.fixes.fix_print_with_import
C_-	print('site_data_dir', app.locations.site_data_dir)
C_+	print(('site_data_dir', app.locations.site_data_dir))
	Fixer 49: lib2to3.fixes.fix_unicode_keep_u
C_-	username = unicode(origin.getFrom()).split('/')[1].replace(" ","")
C_+	username = str(origin.getFrom()).split('/')[1].replace(" ","")