



Bayesian Grammar Learning for Inverse Procedural Modeling

June 2013 | CVPR Anđelo Martinović¹, Luc Van Gool^{1,2}

¹ESAT/PSI/VISICS @ KU Leuven ²Computer Vision Lab @ ETH Zurich







Procedural modeling





























Inverse procedural modeling



. . .





Rules are known in advance

• Metropolis Procedural Modeling (Talton et al. 2011)



 Random Exploration of the Procedural Space for Single-View 3D Modeling of Buildings (Simon et al. 2009)

Shape Grammar Parsing via Reinforcement Learning (Teboul et al. 2011)







Rules are learned

- A Connection between Partial Symmetry and Inverse Procedural Modeling (Bokeloh et al. 2010)

 Inverse procedural modeling by automatic generation of L-systems (Št'ava et al. 2010)







Eidgenössische Technische Hochschule Zürich Swiss Federal Institute of Technology Zurich







- Grammar learning
 - Gold, 1967: No superfinite family of deterministic languages (including regular and context-free languages) can be identified in the limit
 - Horning, 1969: *Stochastic* context free grammars *can* be learned from only positive examples
 - Stolcke, 1994: Introducing Bayesian Model Merging (BMM) for HMM, SCFG and PAG induction
 - Hwang et al, 2011: Rediscovering BMM for induction of probabilistic programs
 - Talton et al, 2012: BMM is used to learn design patterns





- Grammar learning
 - Gold, 1967: No superfinite family of deterministic languages (including regular and context-free languages) can be identified in the limit
 - Horning, 1969: *Stochastic* context free grammars *can* be learned from only positive examples
 - Stolcke, 1994: Introducing Bayesian Model Merging (BMM) for HMM, SCFG and PAG induction
 - Hwang et al, 2011: Rediscovering BMM for induction of probabilistic programs
 - Talton et al, 2012: BMM is used to learn design patterns





- Grammar learning
 - Gold, 1967: No superfinite family of deterministic languages (including regular and context-free languages) can be identified in the limit
 - Horning, 1969: *Stochastic* context free grammars *can* be learned from only positive examples
 - Stolcke, 1994: Introducing Bayesian Model Merging (BMM) for HMM, SCFG and PAG induction
 - Hwang et al, 2011: Rediscovering BMM for induction of probabilistic programs
 - Talton et al, 2012: BMM is used to learn design patterns





- Grammar learning
 - Gold, 1967: No superfinite family of deterministic languages (including regular and context-free languages) can be identified in the limit
 - Horning, 1969: *Stochastic* context free grammars *can* be learned from only positive examples
 - Stolcke, 1994: Introducing Bayesian Model Merging (BMM) for HMM, SCFG and PAG induction
 - Hwang et al, 2011: Rediscovering BMM for induction of probabilistic programs
 - Talton et al, 2012: BMM is used to learn design patterns





- Grammar learning
 - Gold, 1967: No superfinite family of deterministic languages (including regular and context-free languages) can be identified in the limit
 - Horning, 1969: *Stochastic* context free grammars *can* be learned from only positive examples
 - Stolcke, 1994: Introducing Bayesian Model Merging (BMM) for HMM, SCFG and PAG induction
 - Hwang et al, 2011: Rediscovering BMM for induction of probabilistic programs
 - Talton et al, 2012: BMM is used to learn design patterns





- Grammar learning
 - Gold, 1967: No superfinite family of deterministic languages (including regular and context-free languages) can be identified in the limit
 - Horning, 1969: *Stochastic* context free grammars *can* be learned from only positive examples
 - Stolcke, 1994: Introducing Bayesian Model Merging (BMM) for HMM, SCFG and PAG induction
 - Hwang et al, 2011: Rediscovering BMM for induction of probabilistic programs
 - Talton et al, 2012: BMM is used to learn design patterns







- Grammar learning
 - Gold, 1967: No superfinite family of deterministic languages (including regular and context-free languages) can be identified in the limit
 - Horning, 1969: *Stochastic* context free grammars *can* be learned from only positive examples
 - Stolcke, 1994: Introducing Bayesian Model Merging (BMM) for HMM, SCFG and PAG induction
 - Hwang et al, 2011: Rediscovering BMM for induction of probabilistic programs
 - Talton et al, 2012: BMM is used to learn design patterns









Method overview







• Sampling novel buildings







• Sampling novel buildings







• Parsing existing facades







• Parsing existing facades







• Parsing existing facades













2D-ASCFG G = (N, T, S, R, P, A)







Nonterminal symbols





2D-ASCFG

G = (N, T, S, R, P, A)

Nonterminal symbols

Terminal symbols

























$$X \rightarrow \lambda_1 \lambda_2 ... \lambda_n$$














































What kind of a grammar can we learn?







What kind of a grammar can we learn?







```
G = (N, T, S, R, P, A)
N = {S, X, Y}
T = {a}
```

S ^H →XX	[0.2]	{ {0.5,0.5} }
S ^V →XX	[0.8]	{ {0.5,0.5} }
X ^H →YY	[0.8]	{ {0.5,0.5} }
X ^V →YY	[0.2]	{ {0.5,0.5} }
Y <mark>^H</mark> →a	[1.0]	{ {1.0 } }





```
G = (N, T, S, R, P, A)
N = {S, X, Y}
T = {a}
```







```
G = (N, T, S, R, P, A)
N = {S, X, Y}
T = {a}
```







```
G = (N, T, S, R, P, A)N = \{S, X, Y\}T = \{a\}
```









Input lattice l:

Parsing:

44/112



G = (N, T, S, R, P, A)



A simple parsing example



Input lattice l:

Parsing:





$$G = (N, T, S, R, P, A)$$

N = {S, X, Y}
T = {a}



Input lattice l:



Parsing:

46/112





$$G = (N, T, S, R, P, A)$$

N = {S, X, Y}
T = {a}



Input lattice l:



Parsing:

 δ_1 :





$$G = (N, T, S, R, P, A)$$

N = {S, X, Y}
T = {a}











$$G = (N, T, S, R, P, A)$$

 $N = \{S, X, Y\}$
 $T = \{a\}$











$$G = (N, T, S, R, P, A)$$

N = {S, X, Y}
T = {a}











$$G = (N, T, S, R, P, A)$$

N = {S, X, Y}
T = {a}











$$G = (N, T, S, R, P, A)$$

N = {S, X, Y}
T = {a}











$$G = (N, T, S, R, P, A)$$

N = {S, X, Y}
T = {a}











$$G = (N, T, S, R, P, A)$$

N = {S, X, Y}
T = {a}











$$G = (N, T, S, R, P, A)$$

N = {S, X, Y}
T = {a}











$$G = (N, T, S, R, P, A)$$

N = {S, X, Y}
T = {a}



Input lattice l:



δ_2:







$$G = (N, T, S, R, P, A)$$

N = {S, X, Y}
T = {a}











$$G = (N, T, S, R, P, A)$$

N = {S, X, Y}
T = {a}











$$G = (N, T, S, R, P, A)$$

N = {S, X, Y}
T = {a}











$$G = (N, T, S, R, P, A)$$

N = {S, X, Y}
T = {a}











$$G = (N, T, S, R, P, A)$$

N = {S, X, Y}
T = {a}











$$G = (N, T, S, R, P, A)$$

N = {S, X, Y}
T = {a}











$$G = (N, T, S, R, P, A)$$

N = {S, X, Y}
T = {a}











$$G = (N, T, S, R, P, A)$$

N = {S, X, Y}
T = {a}











$$G = (N, T, S, R, P, A)$$

N = {S, X, Y}
T = {a}



Input lattice l:





Input likelihood: $L(1 | G) = \sum_{\delta \Rightarrow 1} P(\delta) = 0.52$





Ρ

[0.2]

[0.8]

[0.8]

[0.2]

[1.0]

{ {0.5,0.5} }

 $\{\{0.5, 0.5\}\}$

 $\{\{0.5, 0.5\}\}$

 $\{\{0.5, 0.5\}\}$

{ {1.0} }

$$G = (N, T, S, R, P, A)$$

N = {S, X, Y}
T = {a}

 $S \xrightarrow{H} X X$

 $S \xrightarrow{V} XX$

X^H→YY

X^V→YY

Y^H→a





 $\begin{array}{ll} \mbox{Input likelihood:} & L(l \mid G) = \sum_{\delta \Rightarrow l} P(\delta) = 0.52 \\ \mbox{Viterbi approximation:} & L(l \mid G) = max_{\delta \Rightarrow l} P(\delta) = 0.512 \end{array}$





Method overview





Lattice creation





- Approach similar to Riemenschneider et al. (2012)
- Every rectangular region (tile) is labeled with the majority vote from the corresponding pixel labels





Bayesian model merging



- **Data incorporation**: given a body of data, build an initial grammar which generates only the input examples.
- **Model merging**: propose a candidate grammar by altering the structure of the currently best grammar.
- **Model evaluation**: evaluate the fitness of the candidate grammar compared to the currently best grammar.
- **Search**: use model merging to explore the grammar space, searching for the optimal grammar



'n_f



Nonterminal merging: Proposing candidate grammars



- Two nonterminals X_1 and X_2 are selected from the current grammar and replaced with a new nonterminal Y
 - RHS occurrences are replaced:

$$\begin{array}{ccc} Z_1 \rightarrow \mu_1 X_1 \lambda_1 \\ Z_2 \rightarrow \mu_2 X_2 \lambda_2 \end{array} \xrightarrow{\text{merge}} & \begin{array}{c} Z_1 \rightarrow \mu_1 Y \lambda_1 \\ \hline & \\ Z_2 \rightarrow \mu_2 Y \lambda_2 \end{array} \end{array}$$

- If $Z_1 = Z_2$, $\mu_1 = \mu_2$, $\lambda_1 = \lambda_2$, productions are merged and their attribute sets joined
- LHS occurrences are also replaced:

$$\begin{array}{ccc} X_1 \rightarrow \lambda_1 & \text{merge} & Y \rightarrow \lambda_1 \\ X_2 \rightarrow \lambda_2 & & & & Y \rightarrow \lambda_2 \end{array}$$

- Productions of the form $Y \rightarrow Y$ are removed
- Restriction: X_1 and X_2 must be 'label-compatible'



Evaluating candidate grammars



- Goal: find the best trade-off between fit to the input data D and a general preference for simpler models (MDL)
- From a Bayesian perspective:

$$\begin{array}{cccc} maximize & P(G \mid D) \sim P(G) \cdot P(D \mid G) \\ & & & & & \\ & & & & \\ posterior & prior & likelihood \end{array}$$


Prior term

 $P(G) = P(G_s) \cdot P(\theta_g \mid G_s)$ Parameter prior Structure prior



- Structure prior follows the Minimum Description Length (MDL) principle
 - $P(G_s) = e^{-DL(G_s)}$
 - Grammar description length (in bits):

 $DL(G_s) = \sum_{r \in \mathbb{R}} \sum_{n \in (r \cap N)} \log |N|$

• Parameter prior is a symmetrical Dirichlet





Likelihood term





Likelihood term



- To calculate the likelihood, one must integrate over the parameter prior...
- ...or, use the Viterbi assumption:
 - Every input sample is generated by a single derivation tree of the grammar
 - The likelihood of that sample is then the product of all rule probabilities used in the Viterbi derivation
 - However, both the Viterbi derivations and the optimal values of rule probabilities θ_{a} are unknown!

Eidgenössische Technische Hochschule Zürich Swiss Federal Institute of Technology Zurich

Likelihood term

KU LEUVEN



To calculate the likelihood, one must integrate over the parameter prior...

$$P(D|G_s) = \int_{\theta_g} P(\theta_g|G_s) P(D|G_s, \theta_g) d\theta_g$$

- ...or, use the Viterbi assumption:
 - Every input sample is generated by a **single** derivation tree of the grammar
 - The likelihood of that sample is then the product of all rule probabilities used in the Viterbi derivation
 - However, both the Viterbi derivations and the optimal values of rule probabilities θ_{a} are unknown!

Eidgenössische Technische Hochschule Zürich Swiss Federal Institute of Technology Zurich

Likelihood term

KU LEUVEN



To calculate the likelihood, one must integrate over the parameter prior...

$$P(D|G_s) = \int_{\theta_g} P(\theta_g|G_s) P(D|G_s, \theta_g) d\theta_g$$

- ...or, use the Viterbi assumption:
 - Every input sample is generated by a **single** derivation tree of the grammar
 - The likelihood of that sample is then the product of all rule probabilities used in the Viterbi derivation
 - However, both the Viterbi derivations and the optimal values of rule probabilities θ_{a} are unknown!

$$\delta_{i}: \begin{array}{c|c} s \end{array} \xrightarrow{?} \\ x \end{array} \xrightarrow{(?)^{2}} \\ \hline x \end{array} \xrightarrow{(?)^{2}} \\ \hline Y \\ Y \end{array} \xrightarrow{(?)^{2}} \\ \hline a \\ a \end{array} \xrightarrow{a} \\ P(\delta_{i})=? \end{array}$$



Likelihood term







- Solution: Expectation-Maximization
- **E-step**: starting from an estimate of θ_{g} , calculate the expected usage counts of each rule
 - Find Viterbi derivations of all input data, count the number of times each rule is used



- Solution: Expectation-Maximization
- **E-step**: starting from an estimate of θ_g , calculate the expected usage counts $\hat{c}(X \to \lambda)$ of each rule
 - Find Viterbi derivations of all input data, count the number of times each rule is used



- Solution: Expectation-Maximization
- **E-step**: starting from an estimate of θ_g , calculate the expected usage counts $\hat{c}(X \to \lambda)$ of each rule
 - Find Viterbi derivations of all input data, count the number of times each rule is used



- Solution: Expectation-Maximization
- **E-step**: starting from an estimate of θ_g , calculate the expected usage counts $\hat{c}(X \to \lambda)$ of each rule
 - Find Viterbi derivations of all input data, count the number of times each rule is used

$$\hat{P}(X \to \lambda) = \frac{\hat{c}(X \to \lambda)}{\sum_{\mu} \hat{c}(X \to \mu)}$$



2D Earley parsing







- How to find the Viterbi derivations in the E-step?
- Generalizing the Earley's string parser (Earley, 1970) to 2D
 - Top-down parsing algorithm
 - Worst-case complexity is O(n³), but usually performs better
 - e.g. O(n²) for unambiguous grammars
 - Does not require the grammar to be in CNF
 - Calculation of Viterbi probabilities with the extension of (Stolcke, 1994)



- How to find the Viterbi derivations in the E-step?
- Generalizing the Earley's string parser (Earley, 1970) to 2D
 - Top-down parsing algorithm
 - Worst-case complexity is O(n³), but usually performs better
 - e.g. O(n²) for unambiguous grammars
 - Does not require the grammar to be in CNF
 - Calculation of Viterbi probabilities with the extension of (Stolcke, 1994)

A. Martinović and L. Van Gool. Earley parsing for 2D stochastic context free grammars. Technical Report KUL/ESAT/PSI/1301, KU Leuven, 2013.



Search in model space



- Posterior calculation is modified using a global prior weight *w*
- Minimizing the energy

 $E(G \mid D) = -w \log P(G) - \log P(D \mid G)$

- w small: low generalization, search procedure stops earlier
- *w* large: increasing the tendency to generalize beyond data
- Greedy best-first search
 - In each iteration, every pair of nonterminals is considered for merging
 - All candidate grammars are evaluated
 - The candidate grammar with minimum energy is accepted if this energy is lower than the current grammar's
 - The rule probabilities are learned in each step using the EM procedure



Obtaining the final model

- Postprocessing Induced lattice grammar
- Casting the induced grammar back to image space
 - 1. Sequences of the same nonterminal symbol are collapsed, e.g.

 $X \rightarrow \lambda Y Y \mu \{\{s_1, y_1, y_2, s_2\}\} \qquad \xrightarrow{\text{collapse}} X \rightarrow \lambda Y \mu \{\{s_1, y_1 + y_2, s_2\}\}$

2. For every rule $p = (X \to \lambda_1 \dots \lambda_k)$, we fit a (k-1)-variate Gaussian distribution $\phi(A) = \mathcal{N}(\bar{\mu}, \hat{\Sigma})$ to the set of its attributes $A(p) = \{\alpha_1 \dots \alpha_n\}$

$$\bar{\mu} = \frac{1}{n} \sum_{j=1}^{n} \alpha_j$$
$$\hat{\Sigma} = \frac{1}{n-1} \sum_{j=1}^{n} (\alpha_j - \bar{\mu}) (\alpha_j - \bar{\mu})^T$$

• This enables us to sample productions with continuous attributes in image space



Parsing in image space

- Find the grammar derivation δ which optimally describes the image
- Optimization is difficult:
 - 1. Rule attributes can have continuous values
 - Solution: Markov Chain Monte Carlo (MCMC) reduces optimization to sampling
 - 2. Grammar is stochastic, the number of rules per derivation can change
 - Solution: Reversible jump MCMC (rjMCMC)
- Posterior of a derivation δ given an image I:











- Rule term: sum of negative log probabilities of all rules in derivation δ
- Attribute term: discrepancy between proposed and expected attributes
- Image term: discrepancy between Random Forest pixel classifier and labeling induced by terminal nodes in δ



$$\begin{split} E(\delta|I) &= -\log P(I|\delta) - \sum_{s \in \delta} \log P(r_s) - \sum_{s \in \delta} \log \phi(A(r_s)) \\ E_{\delta} &= E_{\delta}^{image} + E_{\delta}^{rule} + E_{\delta}^{attribute} \end{split}$$

- Rule term: sum of negative log probabilities of all rules in derivation δ
- Attribute term: discrepancy between proposed and expected attributes
- Image term: discrepancy between Random Forest pixel classifier and labeling induced by terminal nodes in δ



$$E(\delta|I) = -\log P(I|\delta) - \sum_{s \in \delta} \log P(r_s) - \sum_{s \in \delta} \log \phi(A(r_s))$$
$$E_{\delta} = E_{\delta}^{image} + \underbrace{E_{\delta}^{rule}}_{\delta} + E_{\delta}^{attribute}$$

- Rule term: sum of negative log probabilities of all rules in derivation δ
- Attribute term: discrepancy between proposed and expected attributes
- Image term: discrepancy between Random Forest pixel classifier and labeling induced by terminal nodes in δ



$$E(\delta|I) = -\log P(I|\delta) - \sum_{s \in \delta} \log P(r_s) - \sum_{s \in \delta} \log \phi(A(r_s))$$
$$E_{\delta} = E_{\delta}^{image} + \underbrace{E_{\delta}^{rule}}_{\delta} + E_{\delta}^{attribute}$$

- Rule term: sum of negative log probabilities of all rules in derivation δ
- Attribute term: discrepancy between proposed and expected attributes
- Image term: discrepancy between Random Forest pixel classifier and labeling induced by terminal nodes in δ



$$E(\delta|I) = -\log P(I|\delta) - \sum_{s \in \delta} \log P(r_s) - \sum_{s \in \delta} \log \phi(A(r_s))$$
$$E_{\delta} = E_{\delta}^{image} + \underbrace{E_{\delta}^{rule}}_{\delta} + \underbrace{E_{\delta}^{attribute}}_{\delta}$$

- Rule term: sum of negative log probabilities of all rules in derivation δ
- Attribute term: discrepancy between proposed and expected attributes
- Image term: discrepancy between Random Forest pixel classifier and labeling induced by terminal nodes in δ



$$E(\delta|I) = -\log P(I|\delta) - \sum_{s \in \delta} \log P(r_s) - \sum_{s \in \delta} \log \phi(A(r_s))$$
$$E_{\delta} = E_{\delta}^{image} + \underbrace{E_{\delta}^{rule}}_{\delta} + \underbrace{E_{\delta}^{attribute}}_{\delta}$$

- Rule term: sum of negative log probabilities of all rules in derivation δ
- Attribute term: discrepancy between proposed and expected attributes
- Image term: discrepancy between Random Forest pixel classifier and labeling induced by terminal nodes in δ



$$E(\delta|I) = -\log P(I|\delta) - \sum_{s \in \delta} \log P(r_s) - \sum_{s \in \delta} \log \phi(A(r_s))$$
$$E_{\delta} = E_{\delta}^{image} + E_{\delta}^{rule} + E_{\delta}^{attribute}$$

- Rule term: sum of negative log probabilities of all rules in derivation δ
- Attribute term: discrepancy between proposed and expected attributes
- Image term: discrepancy between Random Forest pixel classifier and labeling induced by terminal nodes in δ



$$E(\delta|I) = -\log P(I|\delta) - \sum_{s \in \delta} \log P(r_s) - \sum_{s \in \delta} \log \phi(A(r_s))$$
$$E_{\delta} = E_{\delta}^{image} + E_{\delta}^{rule} + E_{\delta}^{attribute}$$

- Rule term: sum of negative log probabilities of all rules in derivation δ
- Attribute term: discrepancy between proposed and expected attributes
- Image term: discrepancy between Random Forest pixel classifier and labeling induced by terminal nodes in δ



$$E(\delta|I) = -\log P(I|\delta) - \sum_{s \in \delta} \log P(r_s) - \sum_{s \in \delta} \log \phi(A(r_s))$$
$$E_{\delta} = E_{\delta}^{image} + E_{\delta}^{rule} + E_{\delta}^{attribute}$$

- Rule term: sum of negative log probabilities of all rules in derivation δ
- Attribute term: discrepancy between proposed and expected attributes
- Image term: discrepancy between Random Forest pixel classifier and labeling induced by terminal nodes in δ







- rjMCMC with Metropolis-Hastings update
- Markov chain is initialized with a random derivation



- Two possible moves in every iteration:
 - **Diffusion** move parse tree doesn't change, dimensionality of α is preserved
 - **Jump** move a rule in the tree is resampled, dimensionality of α changes
- Improvements:
 - **Parallel tempering** 8 chains are run in parallel on different cooling schedules
 - Delayed rejection a jump move is always followed by a diffusion move, and the moves are accepted or rejected in unison



- rjMCMC with Metropolis-Hastings update
- Markov chain is initialized with a random derivation



- Two possible moves in every iteration:
 - **Diffusion** move parse tree doesn't change, dimensionality of α is preserved
 - **Jump** move a rule in the tree is resampled, dimensionality of α changes
- Improvements:
 - **Parallel tempering** 8 chains are run in parallel on different cooling schedules
 - Delayed rejection a jump move is always followed by a diffusion move, and the moves are accepted or rejected in unison



- rjMCMC with Metropolis-Hastings update
- Markov chain is initialized with a random derivation



Parse tree

"Parameter vector" Concatenated attributes in a pre-order traversal of τ

- Two possible moves in every iteration:
 - **Diffusion** move parse tree doesn't change, dimensionality of α is preserved
 - **Jump** move a rule in the tree is resampled, dimensionality of α changes
- Improvements:
 - **Parallel tempering** 8 chains are run in parallel on different cooling schedules
 - Delayed rejection a jump move is always followed by a diffusion move, and the moves are accepted or rejected in unison





- rjMCMC with Metropolis-Hastings update
- Markov chain is initialized with a random derivation



Parse tree

"Parameter vector" Concatenated attributes in a pre-order traversal of τ

- Two possible moves in every iteration:
 - **Diffusion** move parse tree doesn't change, dimensionality of α is preserved
 - **Jump** move a rule in the tree is resampled, dimensionality of α changes
- Improvements:
 - **Parallel tempering** 8 chains are run in parallel on different cooling schedules
 - Delayed rejection a jump move is always followed by a diffusion move, and the moves are accepted or rejected in unison





- rjMCMC with Metropolis-Hastings update
- Markov chain is initialized with a random derivation



Parse tree

"Parameter vector" Concatenated attributes in a pre-order traversal of τ

- Two possible moves in every iteration:
 - **Diffusion** move parse tree doesn't change, dimensionality of α is preserved
 - **Jump** move a rule in the tree is resampled, dimensionality of α changes
- Improvements:
 - **Parallel tempering** 8 chains are run in parallel on different cooling schedules
 - Delayed rejection a jump move is always followed by a diffusion move, and the moves are accepted or rejected in unison





rjMCMC diffusion move

- Select a random node h_k in the derivation tree
- Sample from a Gaussian proposal distribution centered on the current parameters

 $\delta = (\tau, \alpha)$





Acceptance probability of the move

$$\rho_{\delta \to \delta'} = \min\{1, \frac{p(\delta'|I)}{p(\delta|I)}\} = \min\{1, e^{-(E_{\delta'} - E_{\delta})}\}$$



rjMCMC jump move

- A random node *h* is selected in the derivation tree
- A new rule is sampled from all rules applicable to selected LHS
- If RHS has different length, the whole subtree rooted on h is re-derived
 - Topology of τ changes
 - Dimensionality of α changes



$\delta = (\tau, \alpha)$



Acceptance probability of the move

$$\begin{split} \rho_{\delta \to \delta'} &= \\ \min\{1, \frac{q_{\tau'}(h)}{q_{\tau}(h)} e^{-\left[\left(E_{\delta'}^{img} + E_{\delta'}^{attr}\right) - \left(E_{\delta}^{img} + E_{\delta}^{attr}\right)\right]} \} \end{split}$$

Probability of choosing a nonterminal h in tree τ

Image and attribute terms



ETH Eidgenössische Technische Hochschule Zürich Swiss Federal Institute of Technology Zurich

Parsing results

- ECP dataset
- Example: fold 1 •



monge_1



monge_3





monge_9



monge_71



monge_101



monge_29bis



monge_74





monge_61



monge_79bis









monge_86



Parsing results

Г

• Quantitative analysis



Class	RF [22]	RL [21]	Ours	SOA [8]
Window	29	62	66	75
Wall	58	82	80	88
Balcony	35	58	49	70
Door	79	47	50	67
Roof	51	66	71	74
Sky	73	95	91	97
Shop	20	88	81	93
Overall	48.55	74.71	74.82	84.17


Generating novel designs









Generating novel designs

• The effect of prior weight w on grammar generalization





w = 0.3

Future work

KU LEUVEN

- State of the art in facade parsing
 - Images → Labelings

Eidgenössische Technische Hochschule Zürich Swiss Federal Institute of Technology Zurich

- This paper
 - Labelings → Grammar
- Combine the two approaches
 - "Facade bootstrapping"

A. Martinović, M. Mathias, J. Weissenberg, and L. Van Gool. A three-layered approach to facade parsing. In ECCV, 2012







Thank you

Eidgenössische Technische Hochschule Zürich Swiss Federal Institute of Technology Zurich



