A Deep Learning Model of Orientation Perception and Judgement

2019-08-29, ECVP

Pierre Lelièvre [1, 2], Peter Neri [1]

1. Laboratoire des systèmes perceptifs, DEC, ENS, PSL University, Paris, France
2. Laboratoire SACRe (Sciences, Arts, Création, Recherche), ENS, PSL University, Paris, France

Speaker notes

- Hi everyone, my name is Pierre Lelièvre and I am a PhD Student from the ENS in Paris.
- This morning, I will introduce to you "A Deep Learning Model of Orientation Perception and Judgement"

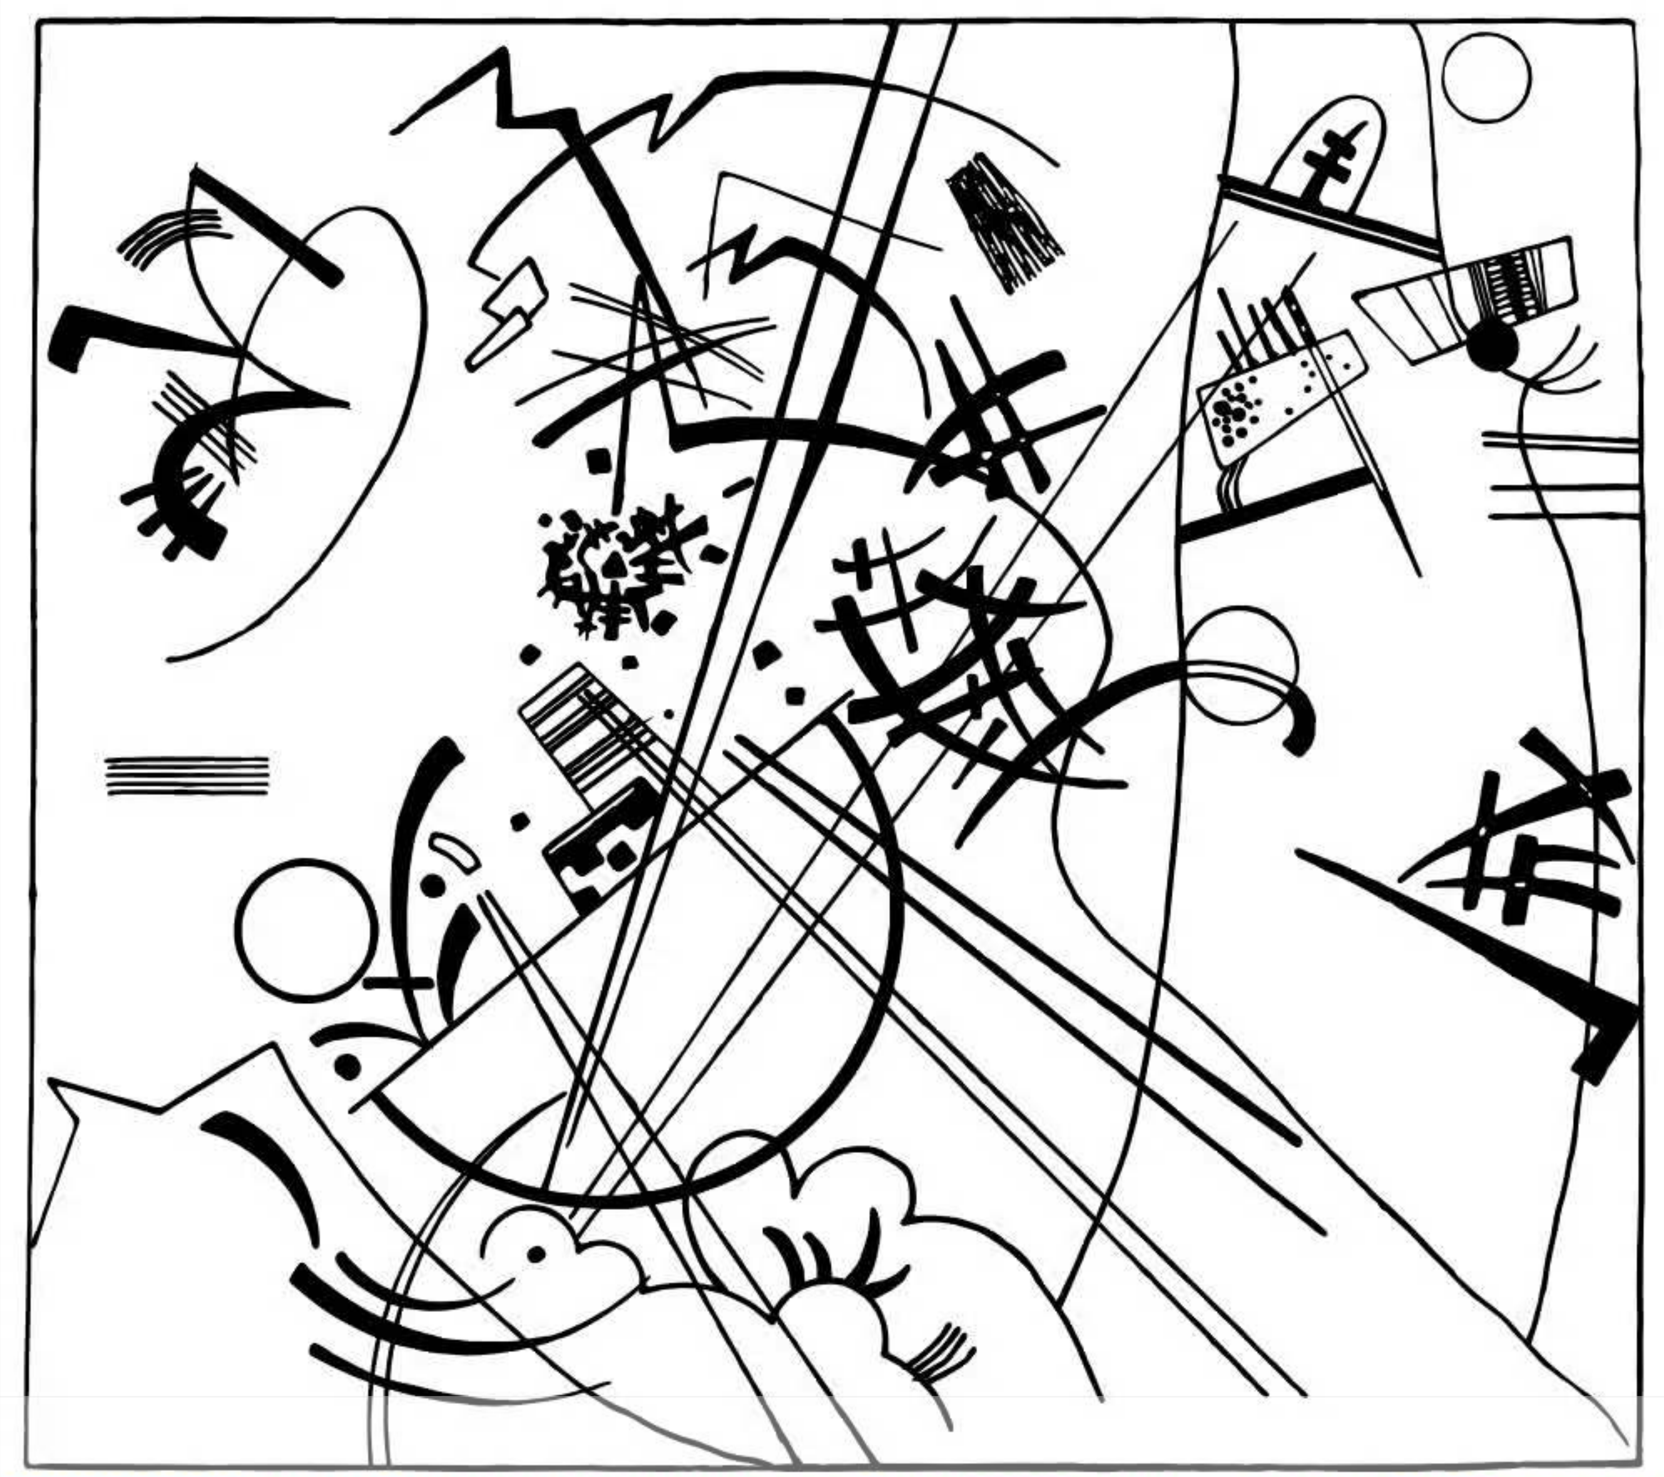

Kandinsky, Small Dream in Red, 1925

Speaker notes

- Artistic composition is often characterized by some basic rules and heuristics, but art history does not offer quantitative tools for segmenting individual elements and measuring their interactions.
- So, to discover whether a metric description of this kind is even possible, our study rely on the painting orientation judgement. Especially because, it is the most simple and objective aspect of the composition. I would also add that it is the less controversial metric from an artistic point of view.
- Moreover, even if the orientation judgement is a global property, we can question the granularity of its cues and begin to define the right scale to look at pictorial elements.

Painting Orientation Judgement

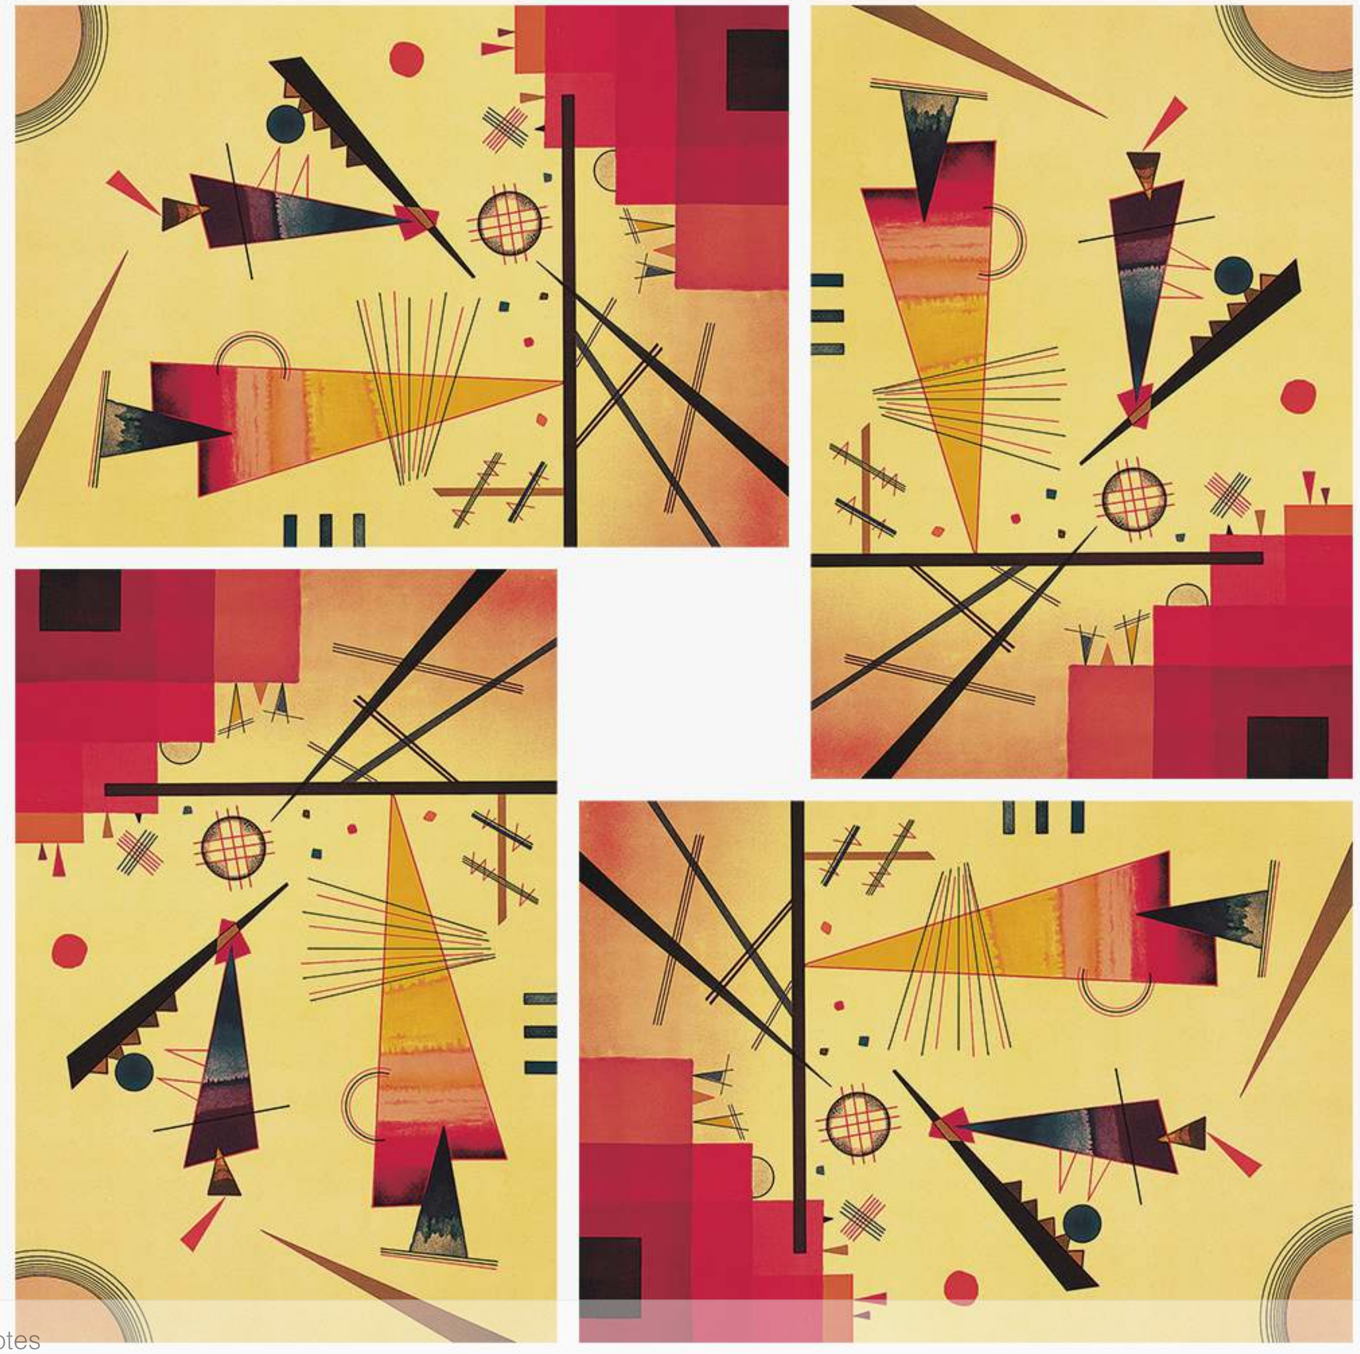

Kandinsky, Struttura Allegra, 1926 & Orange, 1923

Speaker notes

- With the assumption that the true orientation of a painting is the one selected by the artist, are we able to guess when a painting is hung "right-side up" ?
- Particularly for abstract painting, where no "meaningful" content is available.
- I let you a few more seconds to try yourself on these two paintings by Kandinsky...

Painting Orientation Judgement

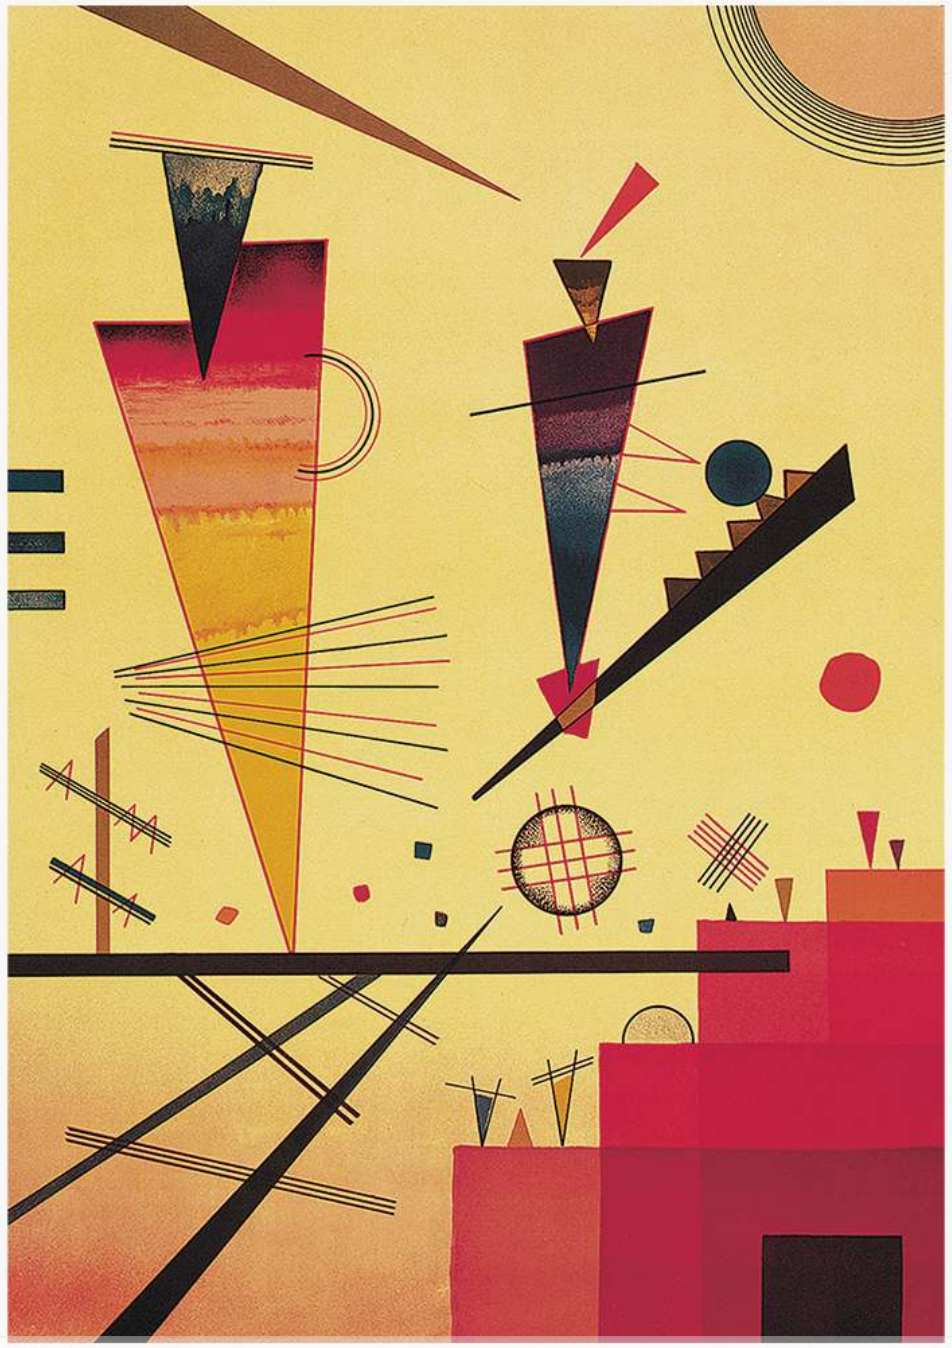

Kandinsky, Struttura Allegra, 1926 & Orange, 1923

Speaker notes

- And here are the results.
- You should roughly have got 1 on 2, because...

Painting Orientation Judgement

- **Lindauer, Martin S. 1969.** “The Orientation of Form in Abstract Art.” Proceedings of the Annual Convention of the American Psychological Association 4 (Pt. 1): 475–76.
- **Lindauer, Martin S. 1987.** “Perceived and Preferred Orientations of Abstract Art.” Empirical Studies of the Arts 5 (1): 47–58.
- **Mather, George. 2012.** “Aesthetic Judgement of Orientation in Modern Art.” I-Perception 3 (1): 18–24.

> *Results showed that the **correct orientation** was selected in **48%** of trials on average, significantly above the 25% chance level.*
>
> *Mather, 2012*

Speaker notes

- From the literature in this domain, the orientation judgement appeared to be a strong behavioral marker.
- Quoting Mather's 2012 paper :

Results showed that the correct orientation was selected in 48% of trials on average, significantly above the 25% chance level.

Web Experiments & Deep Learning Model

Speaker notes

- But is this finding still true with paintings of a wider range of styles and abstraction levels ?
- Can a deep learning model do the task and explain some human's judgement mechanisms ?

Web Experiments

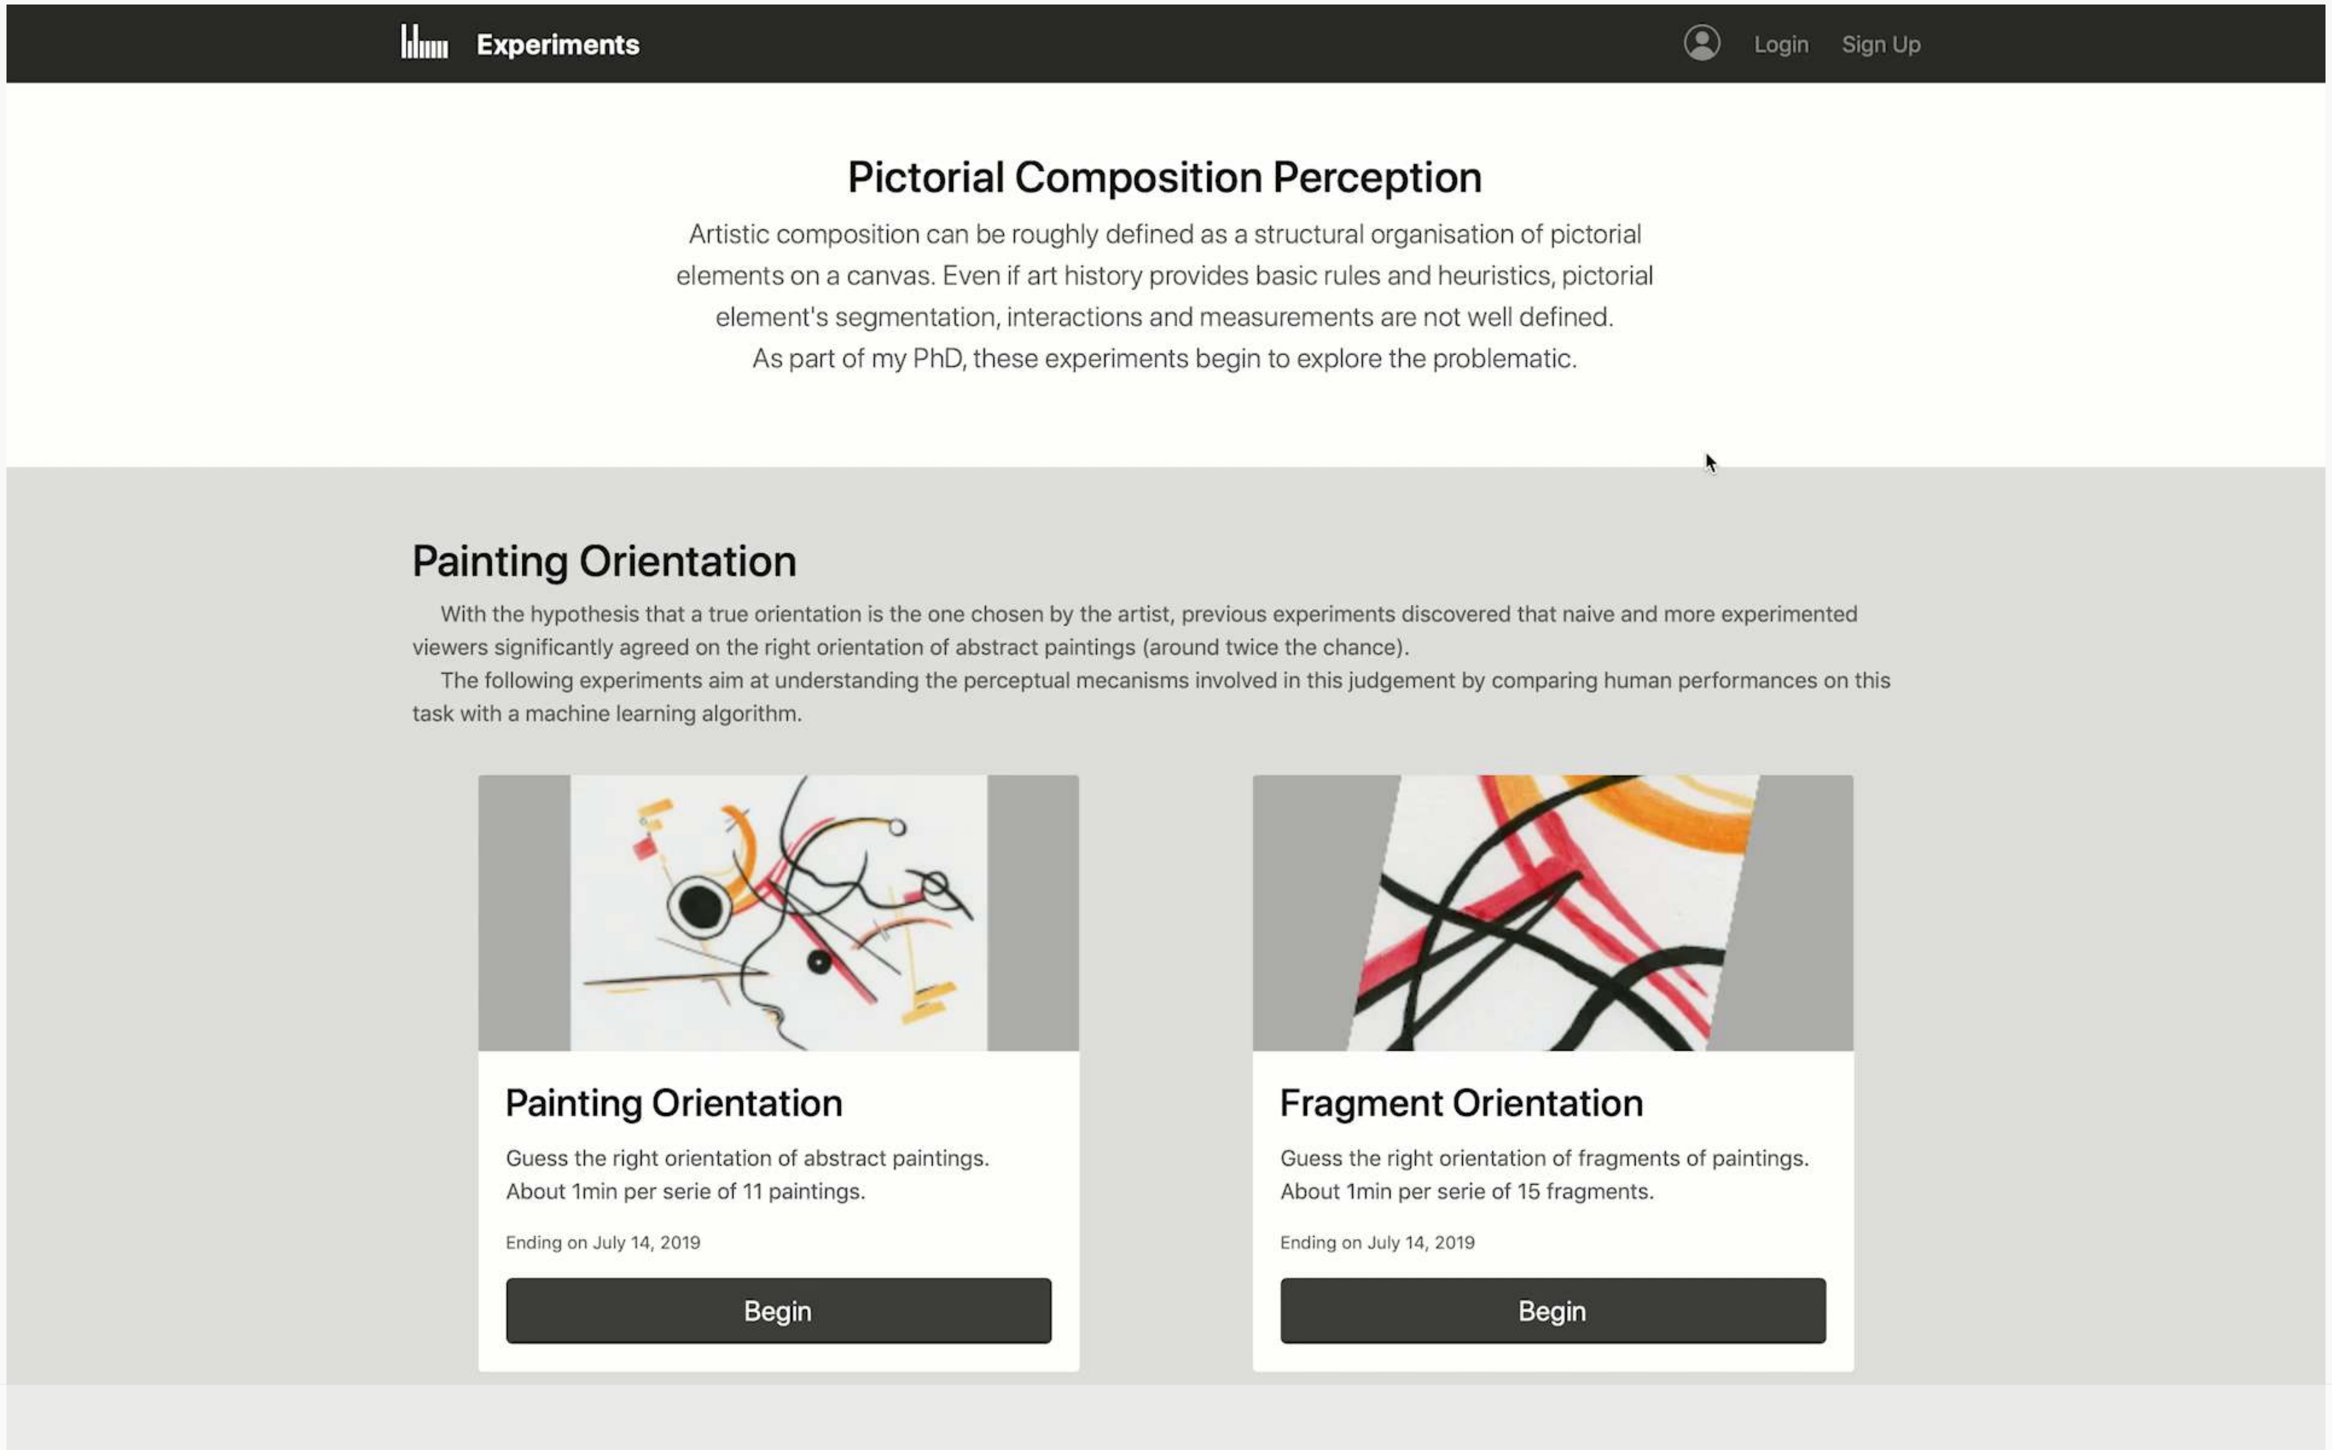

Paintings images and metadata from **wikiart.org**

Speaker notes

- To collect human data, a dedicated website have been developed.
- After having signed up and given their age and art knowledge participants have access to 2 experiments.

Web Experiments

Paintings images and metadata from **wikiart.org**

Speaker notes

- The first one ask people to guess the original orientation of randomly picked abstract paintings among more than 10000.
- If something in the painting gives an obvious hint of the orientation (like a word or a signature), people are invited to report it.

Web Experiments

Paintings images and metadata from **wikiart.org**

Speaker notes

- Some realistic paintings are also introduced to check whether participants are really doing the task.
- Finally, after each series, they get their score and some details about the paintings.

Web Experiments

- **71** participants (15 → 67 yo)
- **47%** accuracy on abstract painting (whole painting)
- contribute at experiments.plelievre.com

Speaker notes

- Before continuing, I can already tell you that from our 71 participants ranging from 15 to 67 yo, that performance reproduces the literature with an average accuracy of 47%.
- The experiment is still running, so feel free to contribute on this website.

Web Experiments

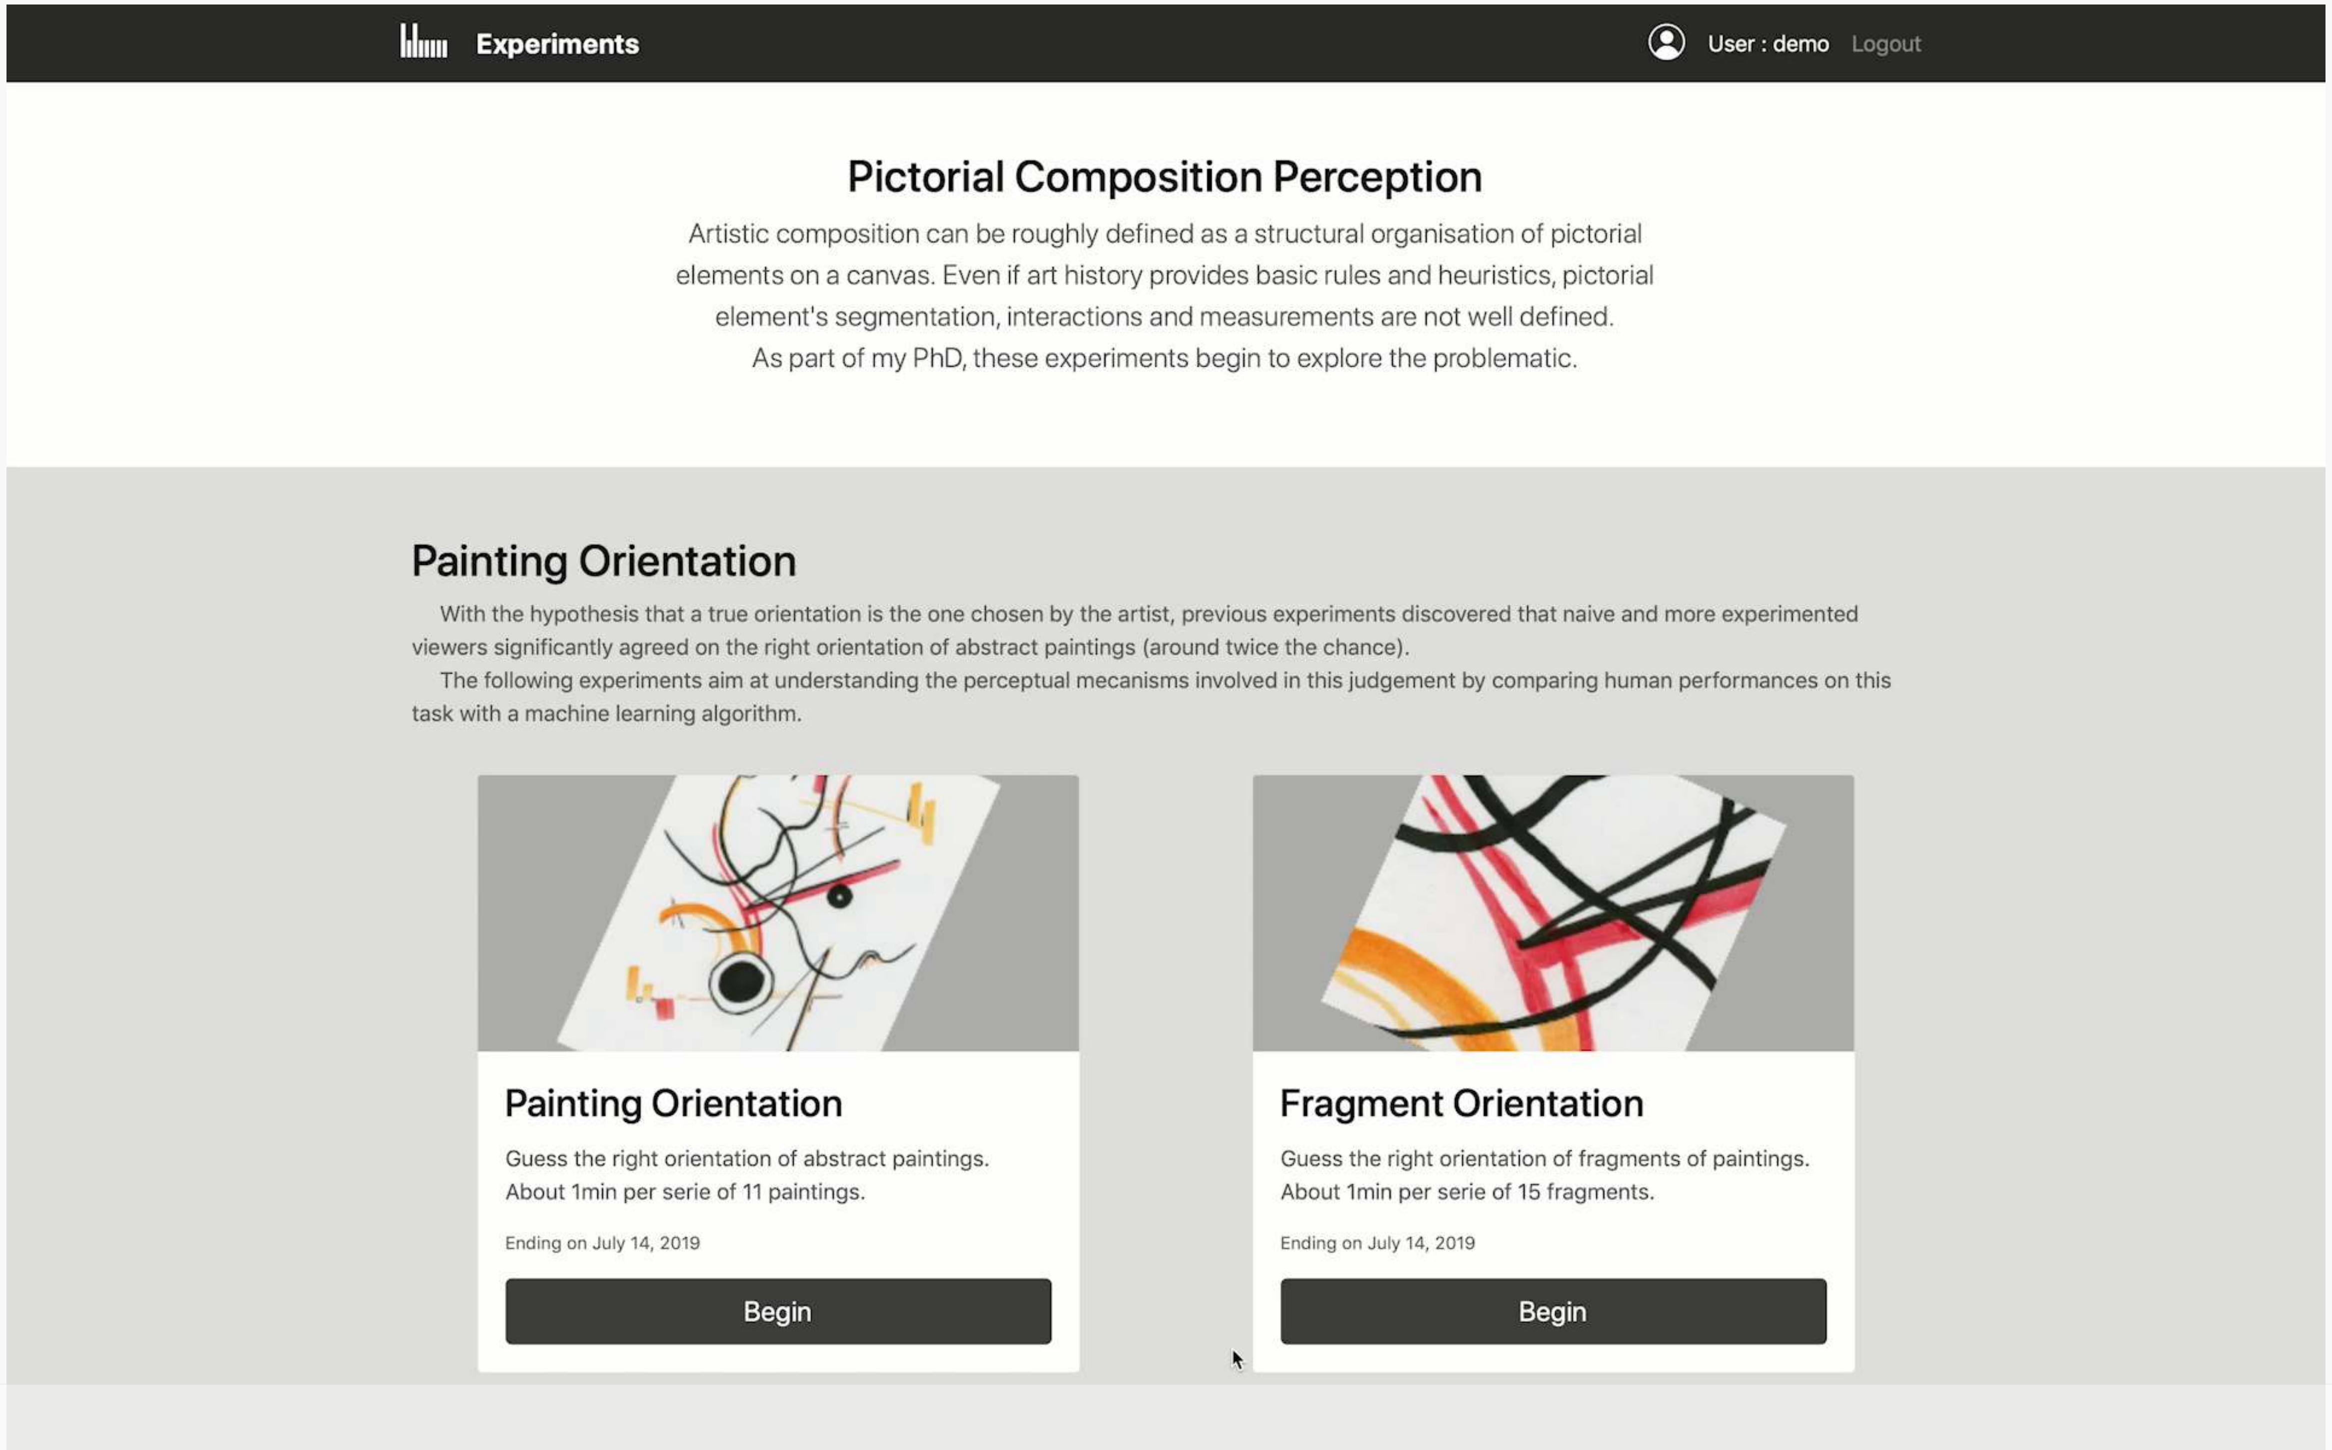

Speaker notes

- The second experiment is far more difficult because it only displays some fragments of abstract and realistic paintings which can be very small and blurry.
- This task aims at understanding the composition perception and judgement at different granularities. You will see in few seconds that it actually has a direct link with the deep learning model we have used.

Deep Learning Model

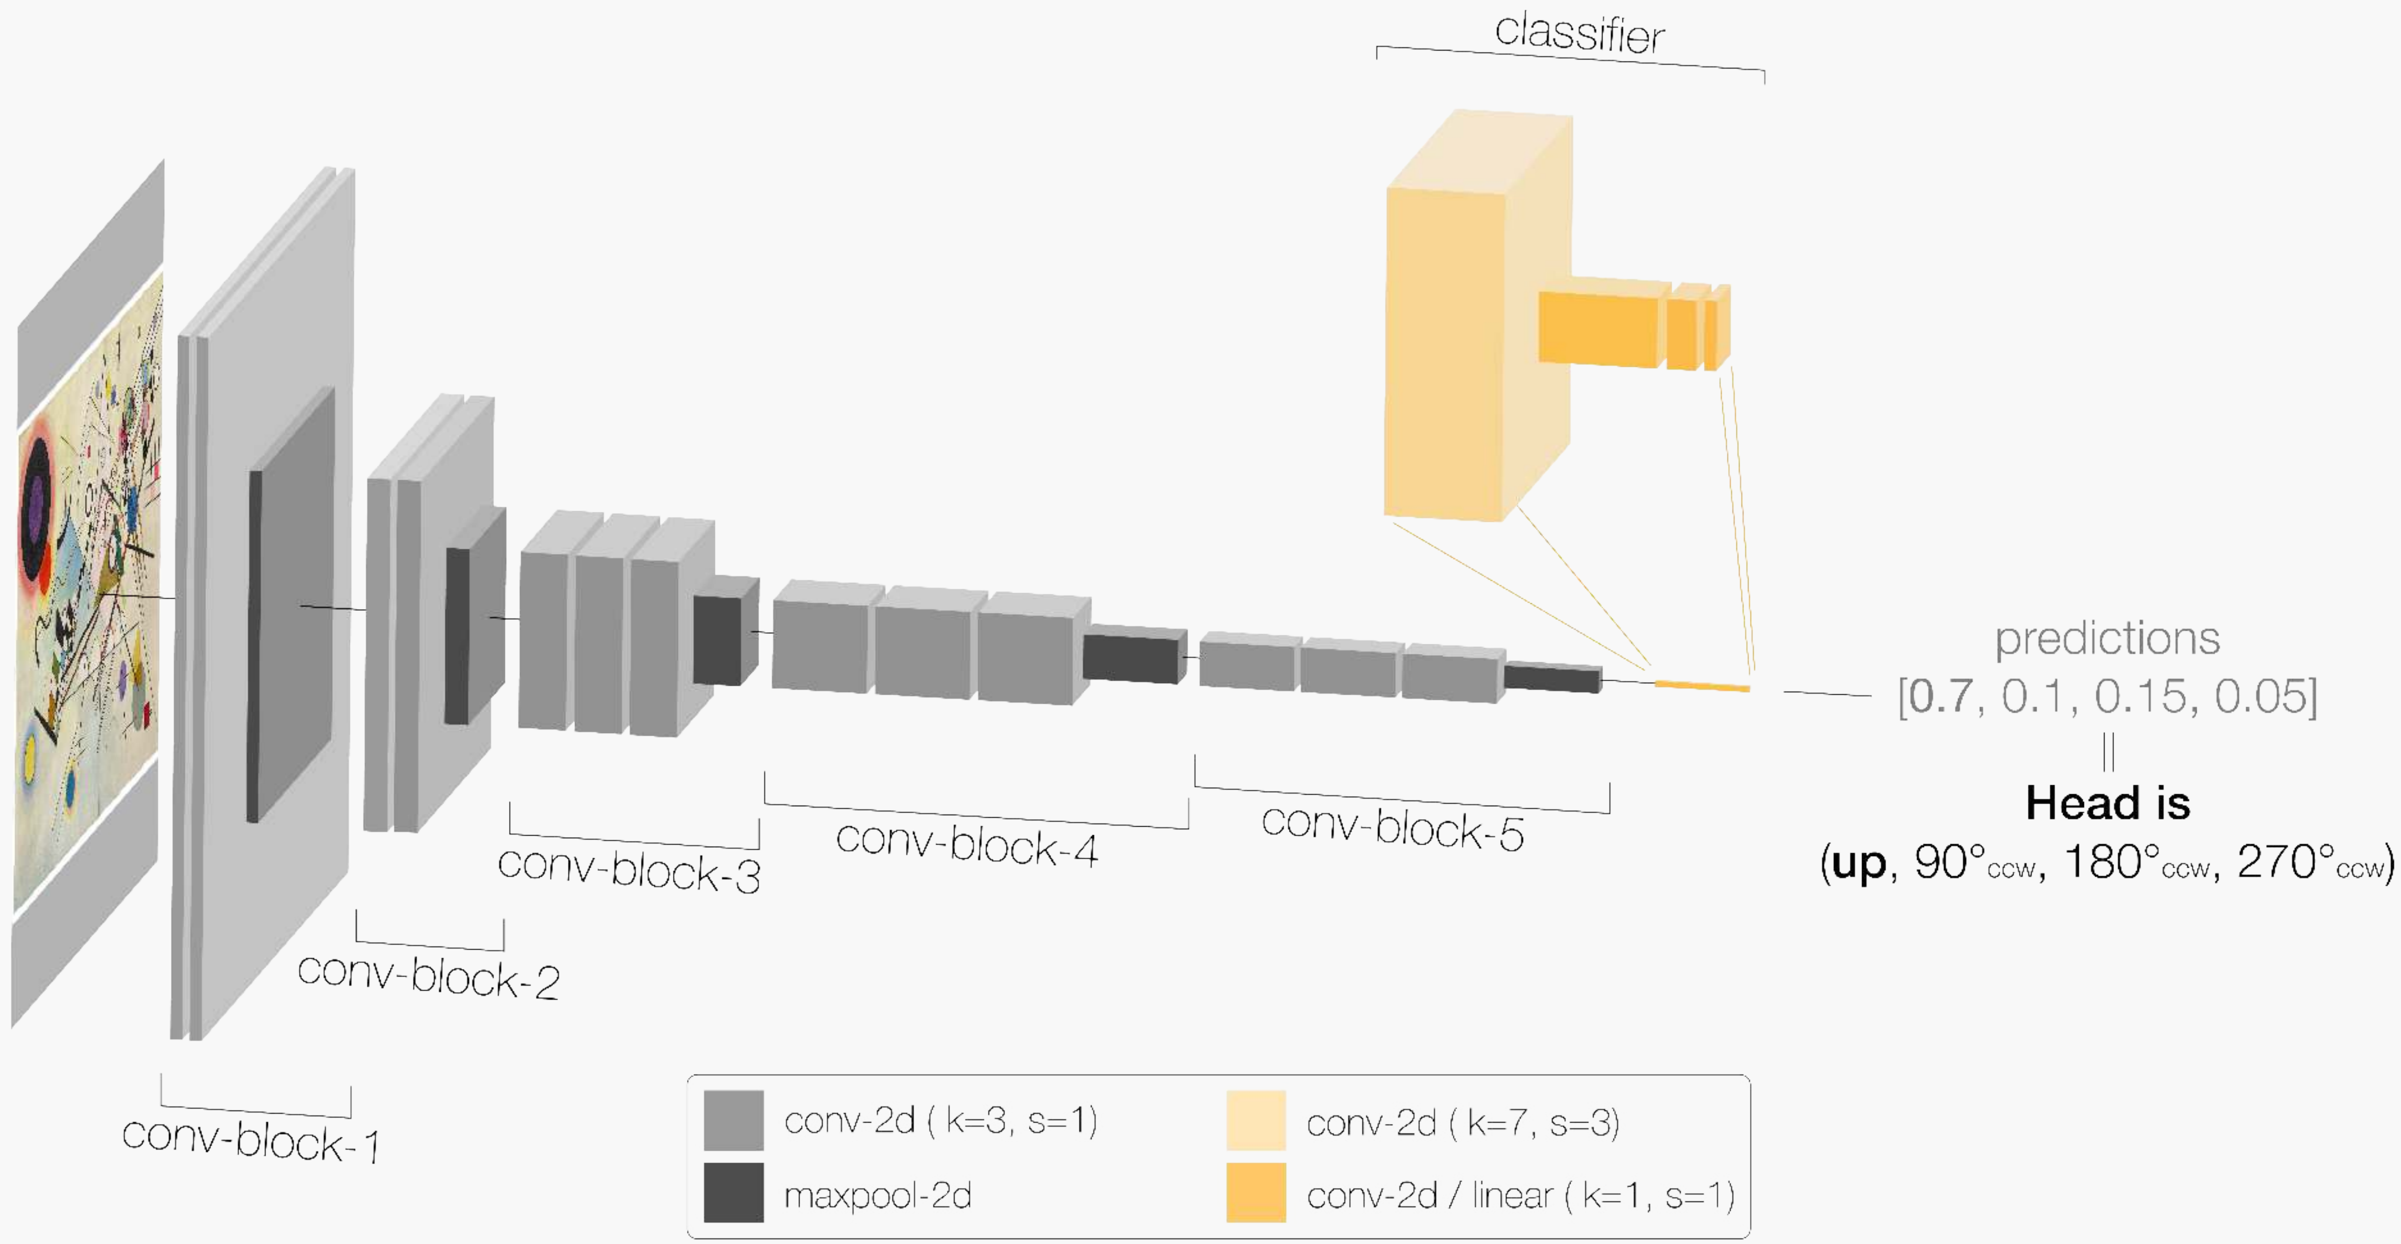

Simonyan, Karen, and Andrew Zisserman. 2014. "Very Deep Convolutional Networks for Large-Scale Image Recognition." ArXiv:1409.1556 [Cs]

Speaker notes

- Our model is based on the very popular VGG (ViJayJay) architecture, which was initially used to do image "content" classification and that have been reported to have similarities with human visual processing.
- In our case, we have actually only 4 classes corresponding to the 4 painting orientations, so it has been quite easy to transform a pre-trained VGG to do our task.

Deep Learning Model

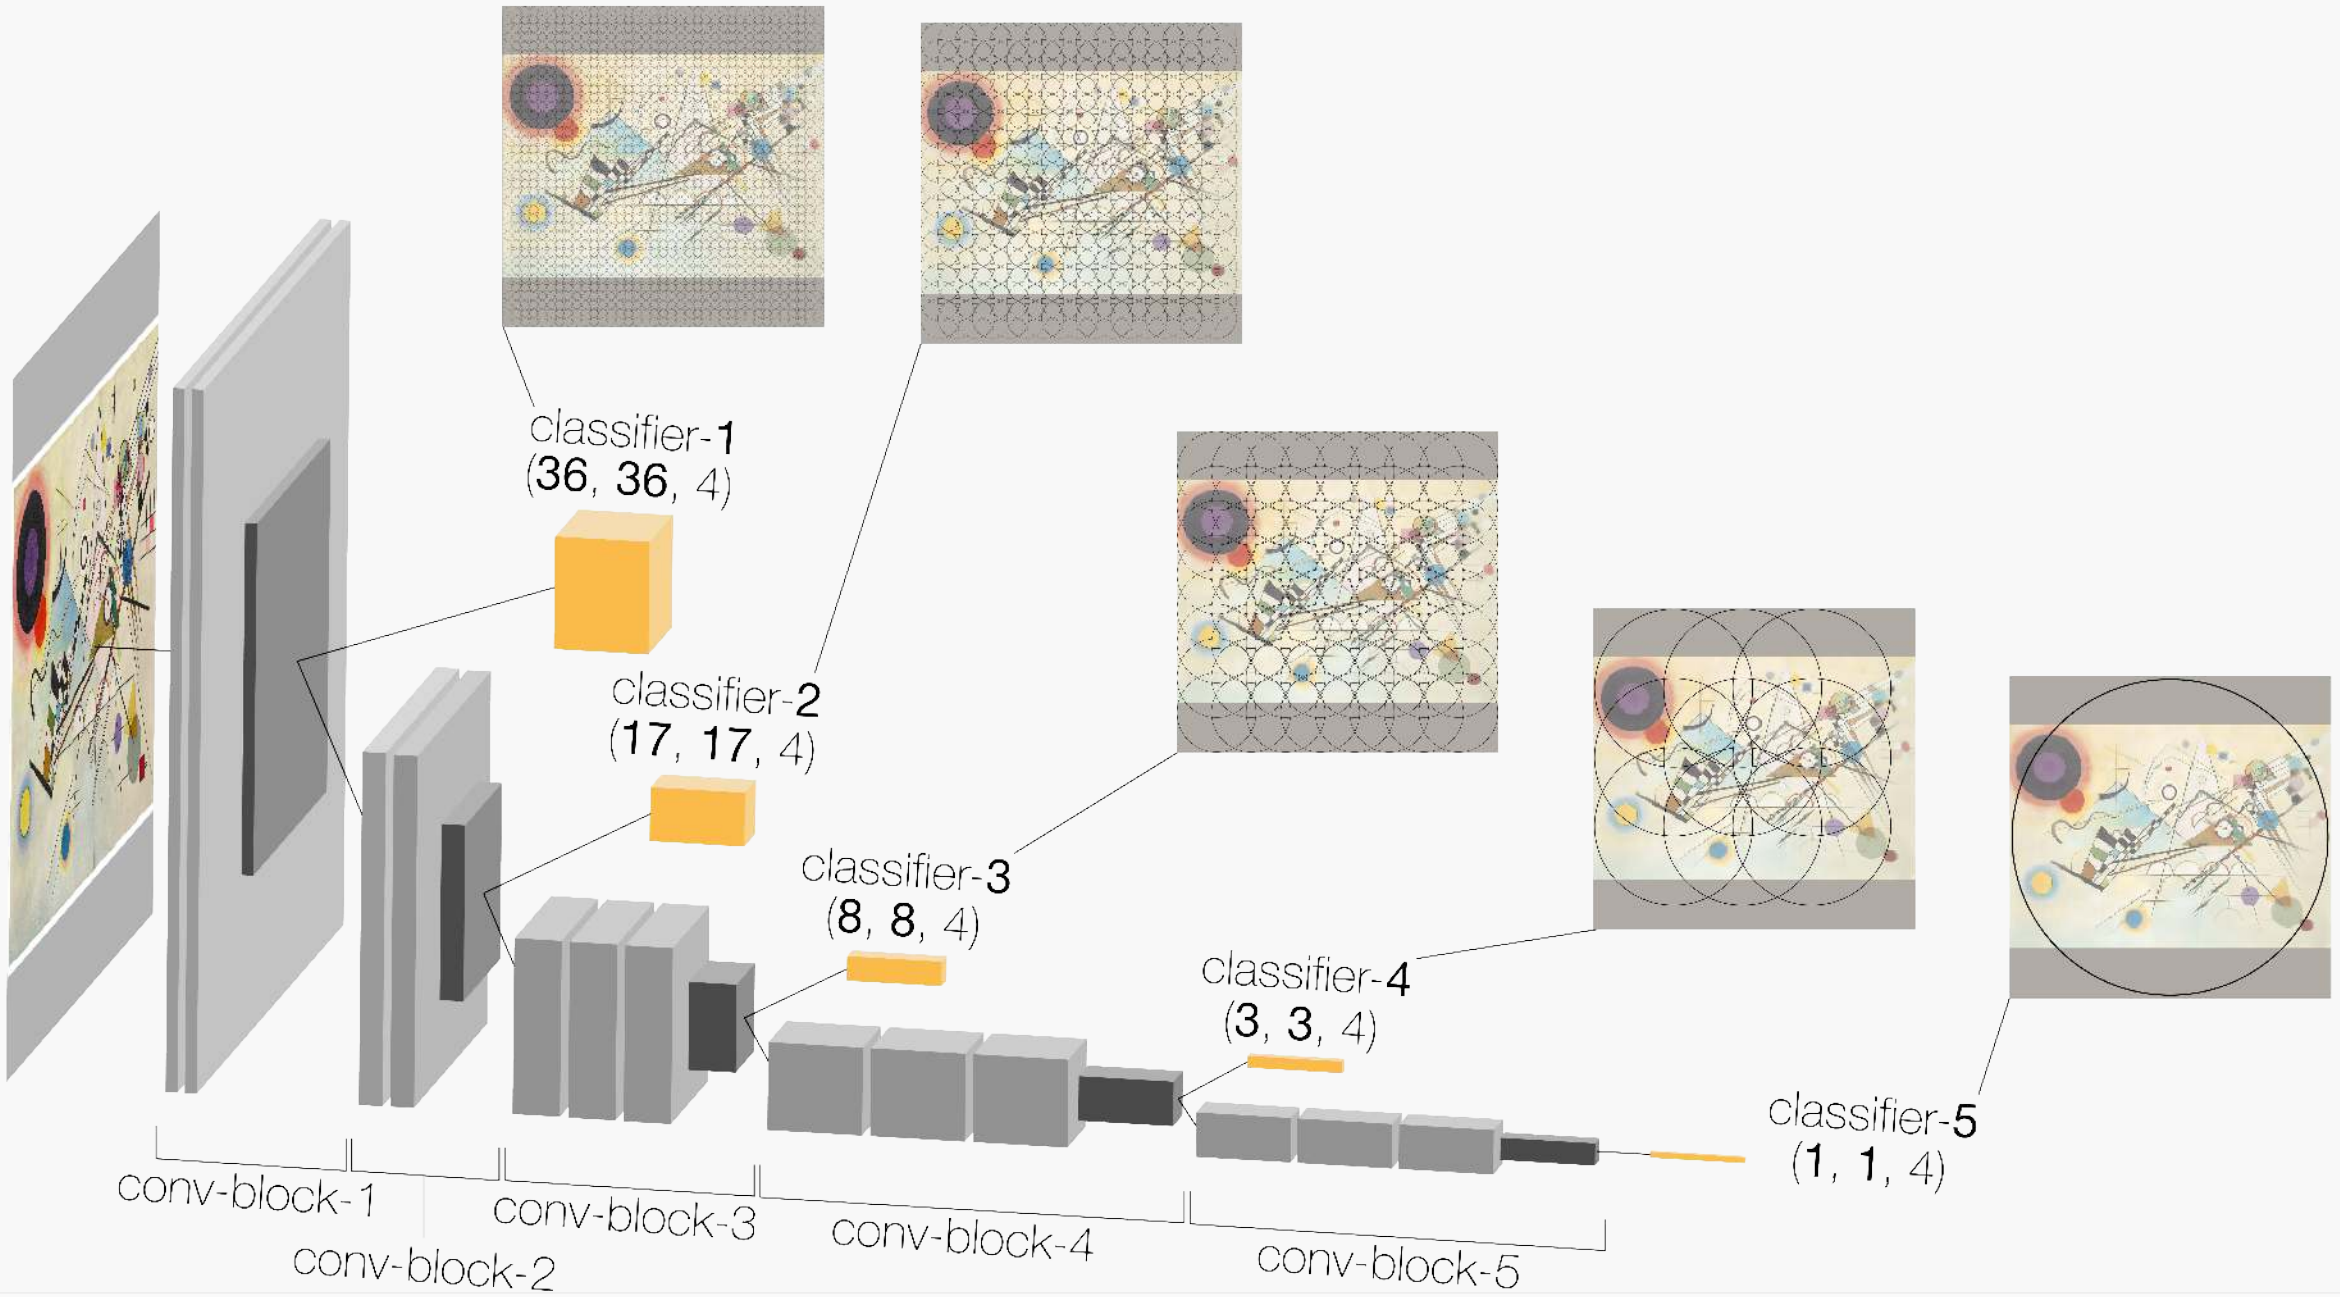

Speaker notes

But to better understand how the network behaves, we have added 4 other classifiers after each convolutional blocks. As a consequence, earlier levels classifiers corresponding to earlier visual perceptive areas have access to smaller receptive fields. It is like if the network was simultaneously judging the orientation of multiple fragments in the picture.

Model Performance

Speaker notes

- So, first, lets have a look at the model performance.

Model Performance

Accuracy per Genre

Abstract

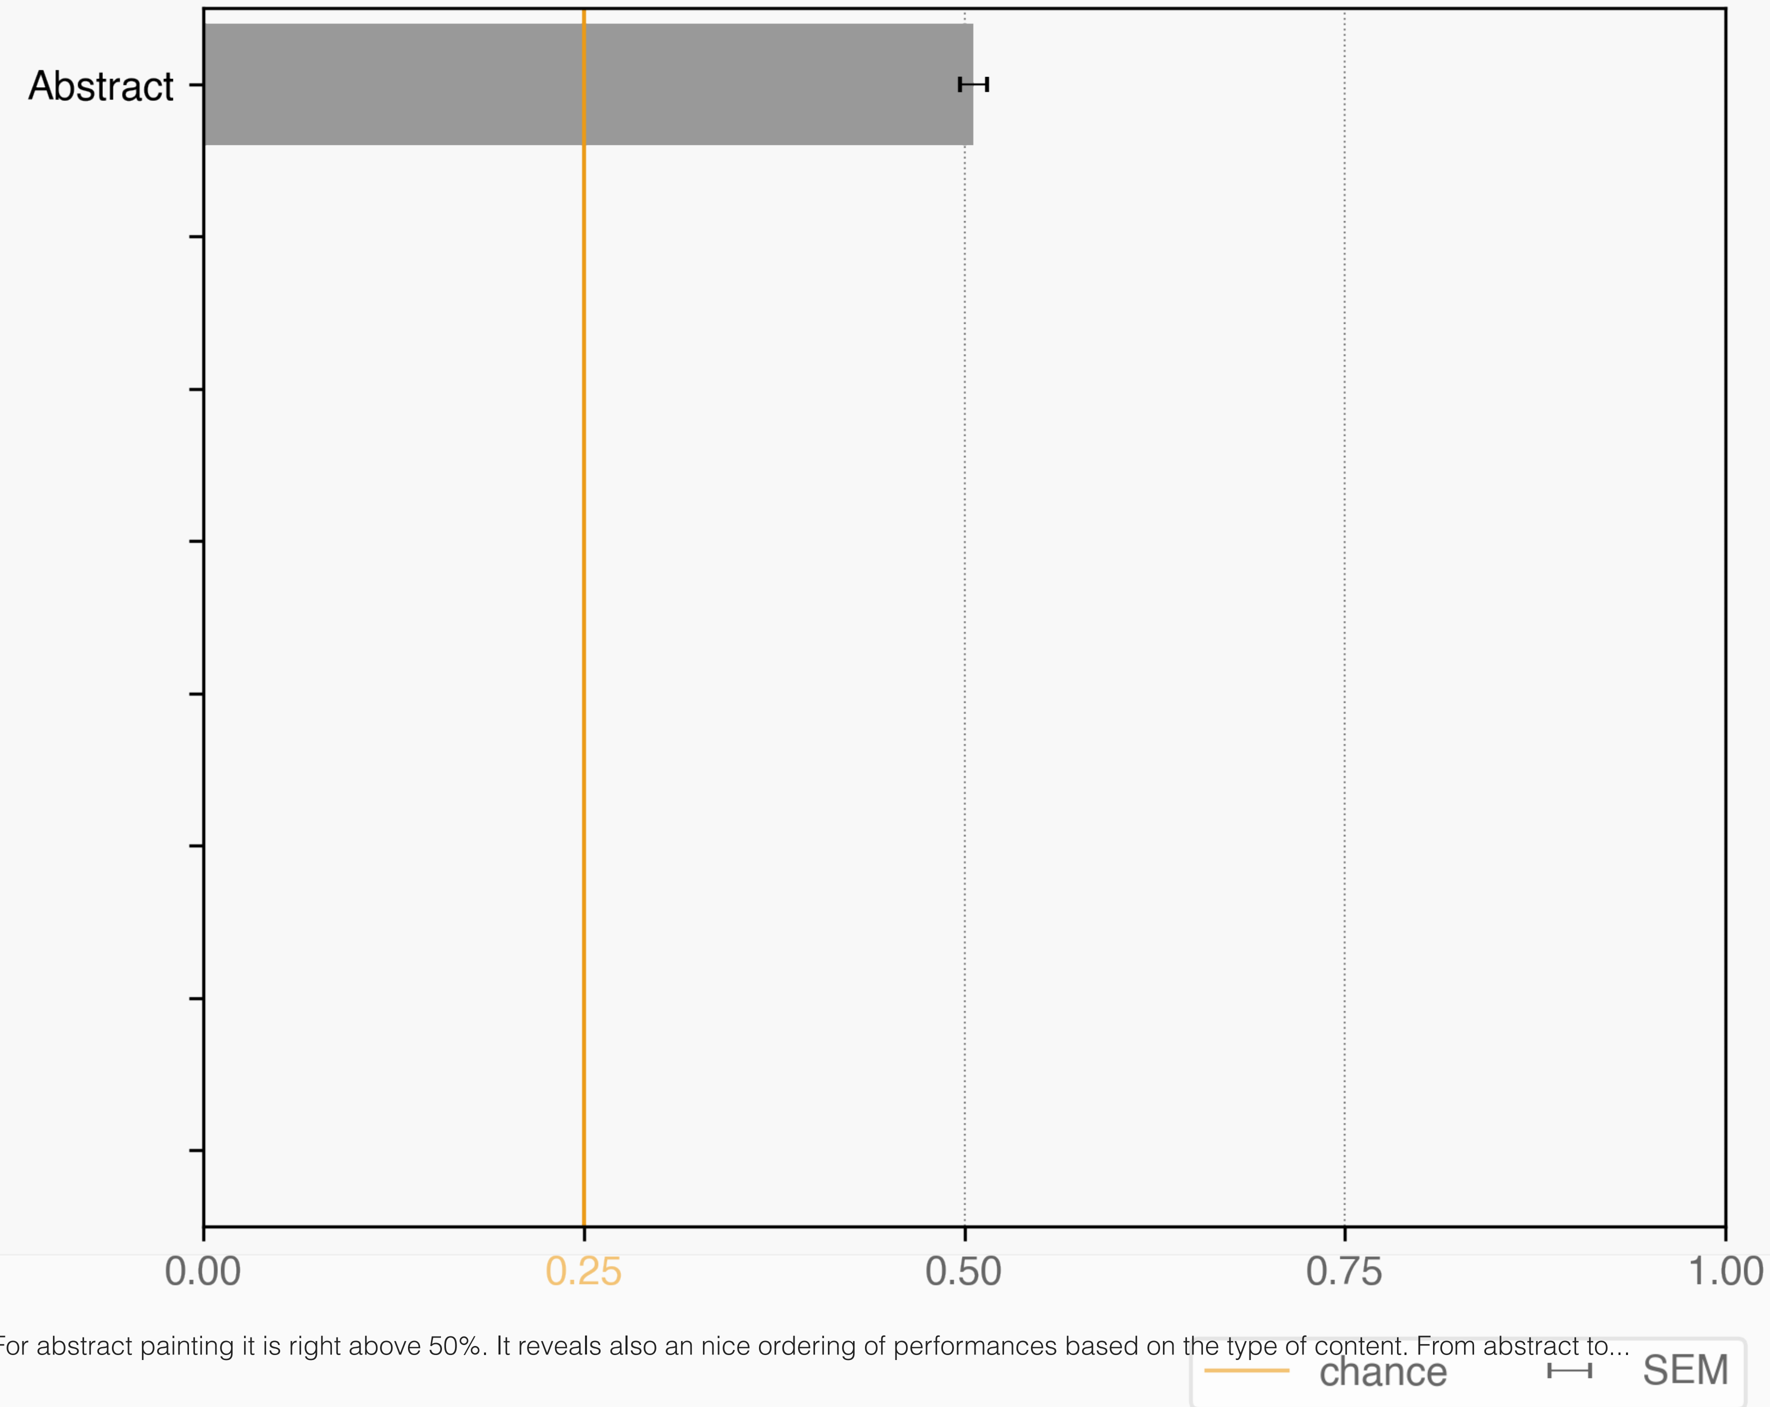

Speaker notes

- After training, performance on the whole painting are in accordance with the literature. For abstract painting it is right above 50%. It reveals also an nice ordering of performances based on the type of content. From abstract to...
- Objects.
- Landscapes.
- and Portraits.

Model Performance

Accuracy per Genre

Still Life

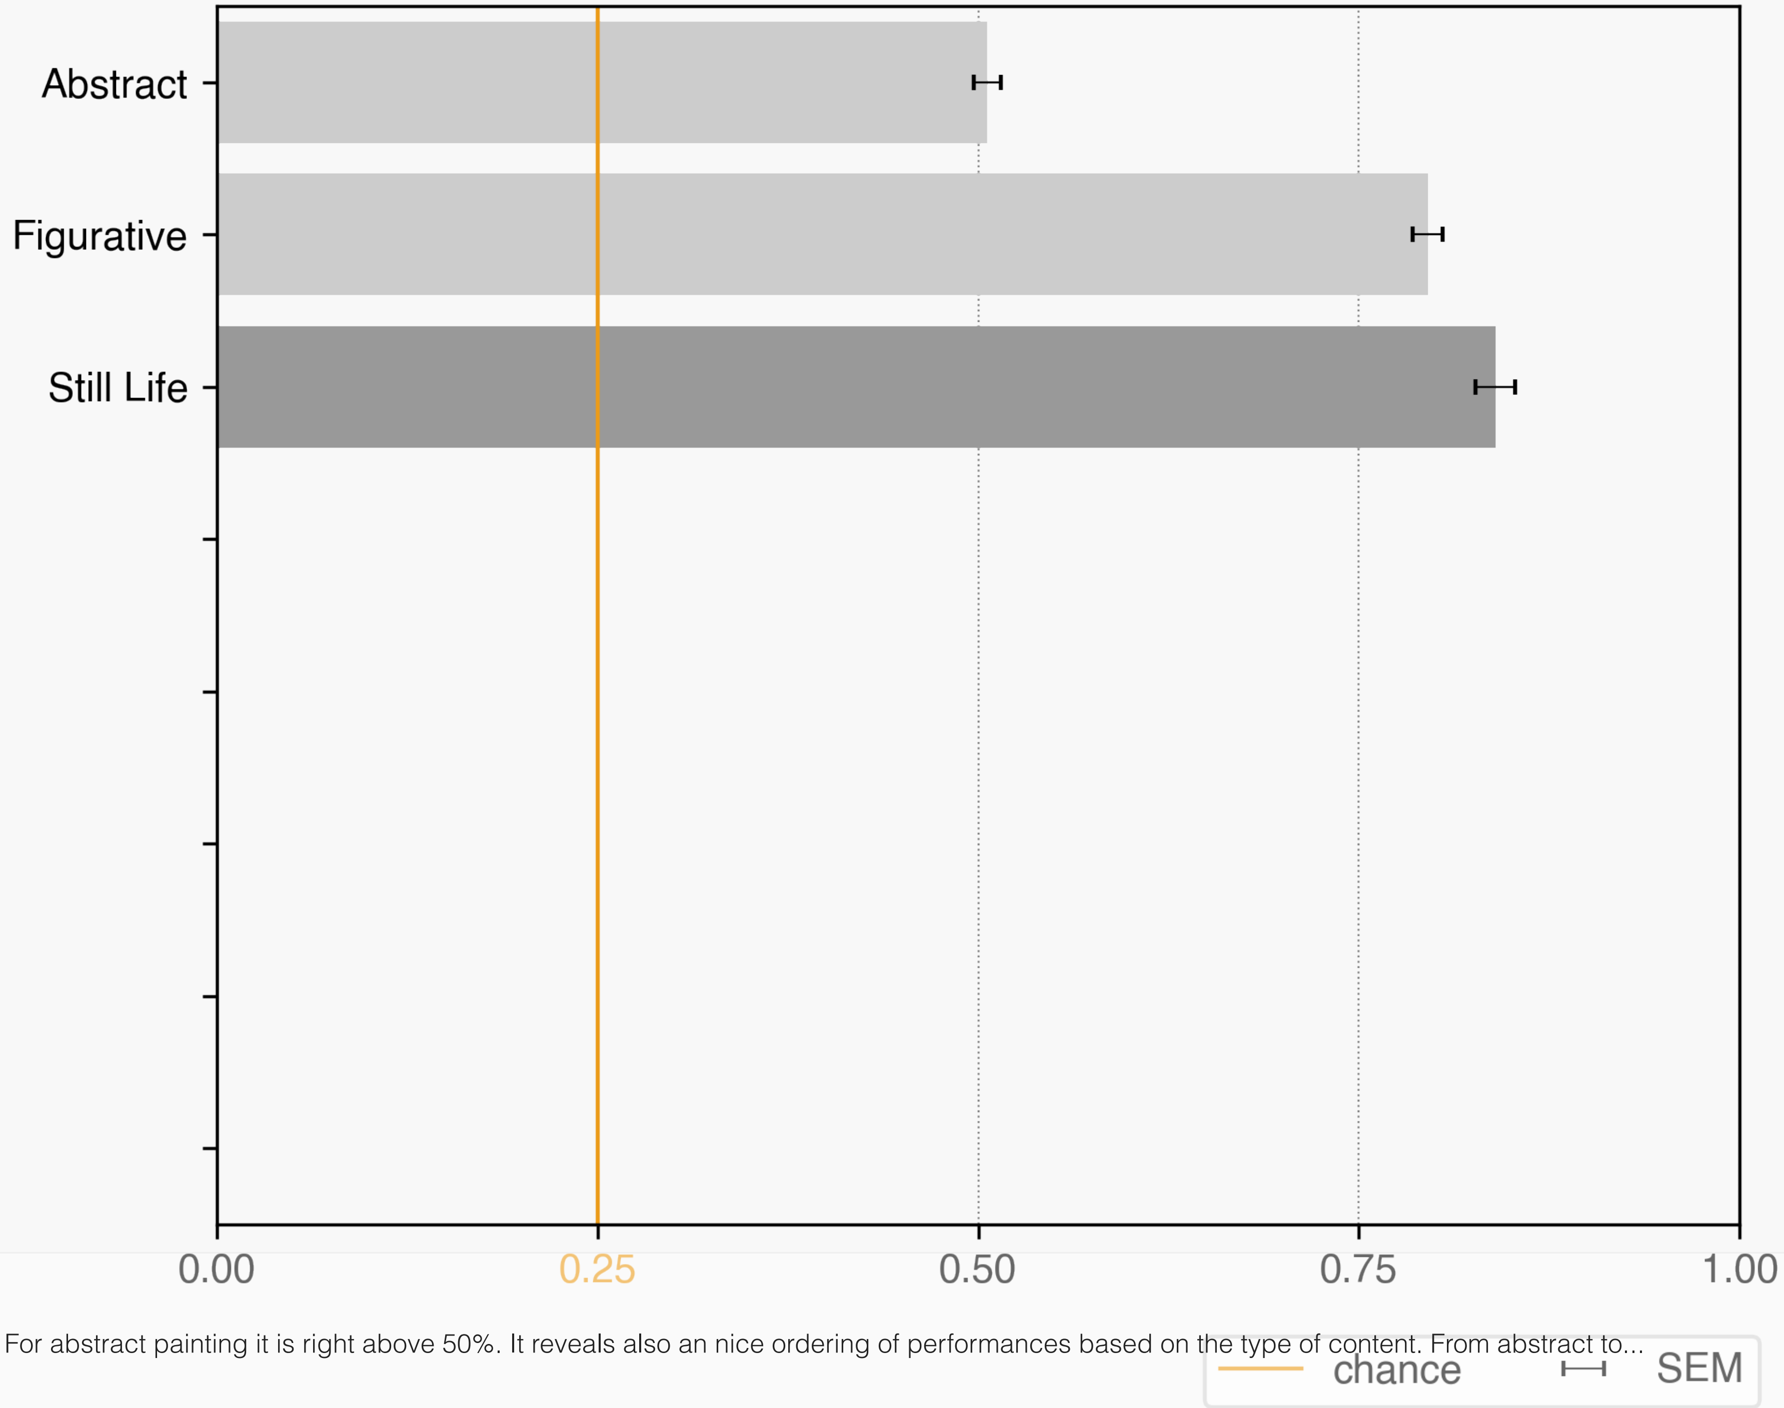

Speaker notes

- After training, performance on the whole painting are in accordance with the literature. For abstract painting it is right above 50%. It reveals also an nice ordering of performances based on the type of content. From abstract to...
- Objects.
- Landscapes.
- and Portraits.

Model Performance

Accuracy per Genre

Landscape

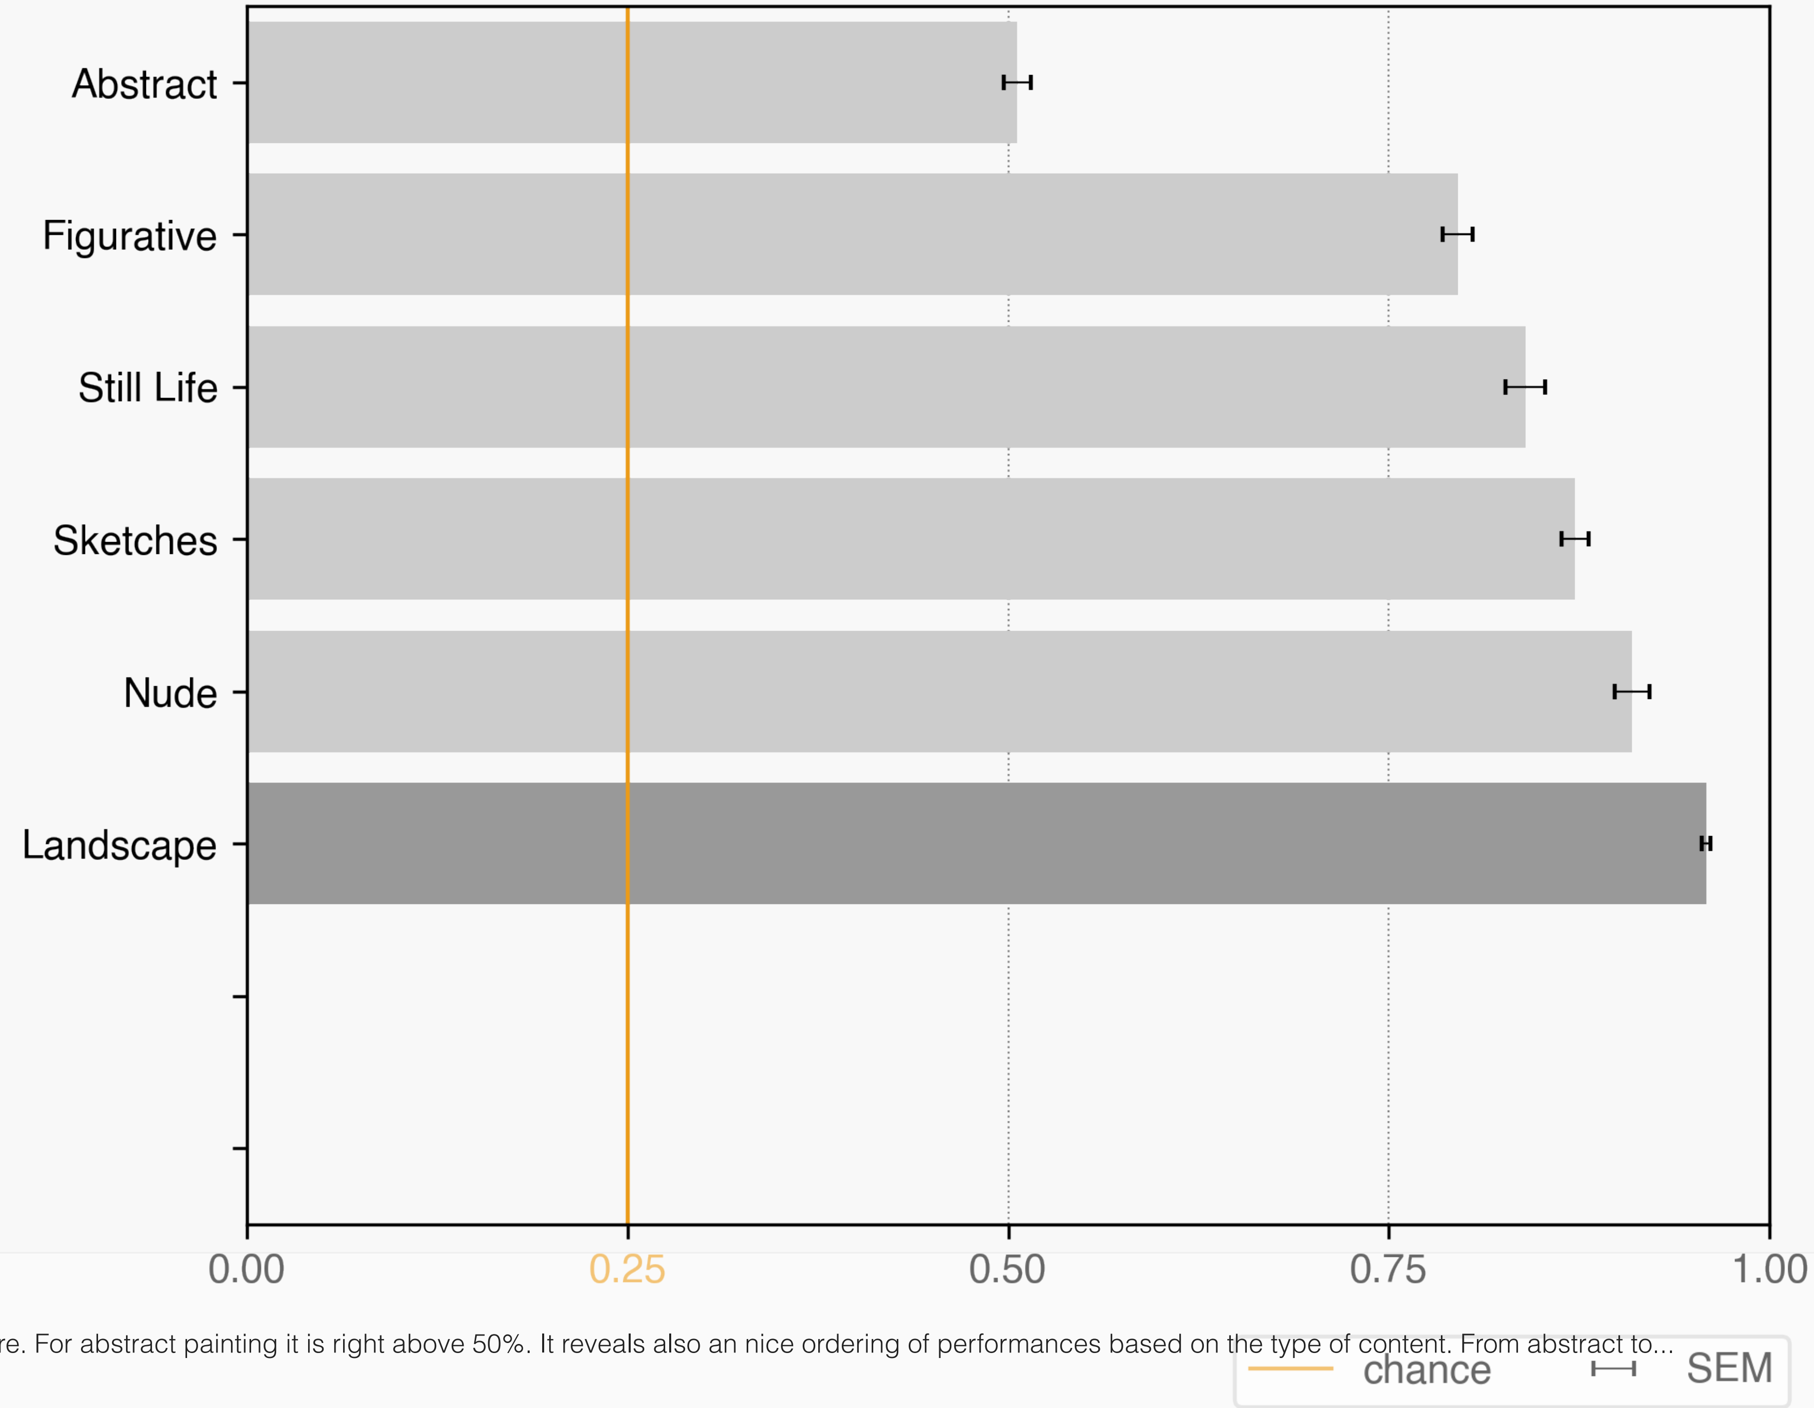

Speaker notes

- After training, performance on the whole painting are in accordance with the literature. For abstract painting it is right above 50%. It reveals also an nice ordering of performances based on the type of content. From abstract to...
- Objects.
- Landscapes.
- and Portraits.

Model Performance

Accuracy per Genre

Portrait

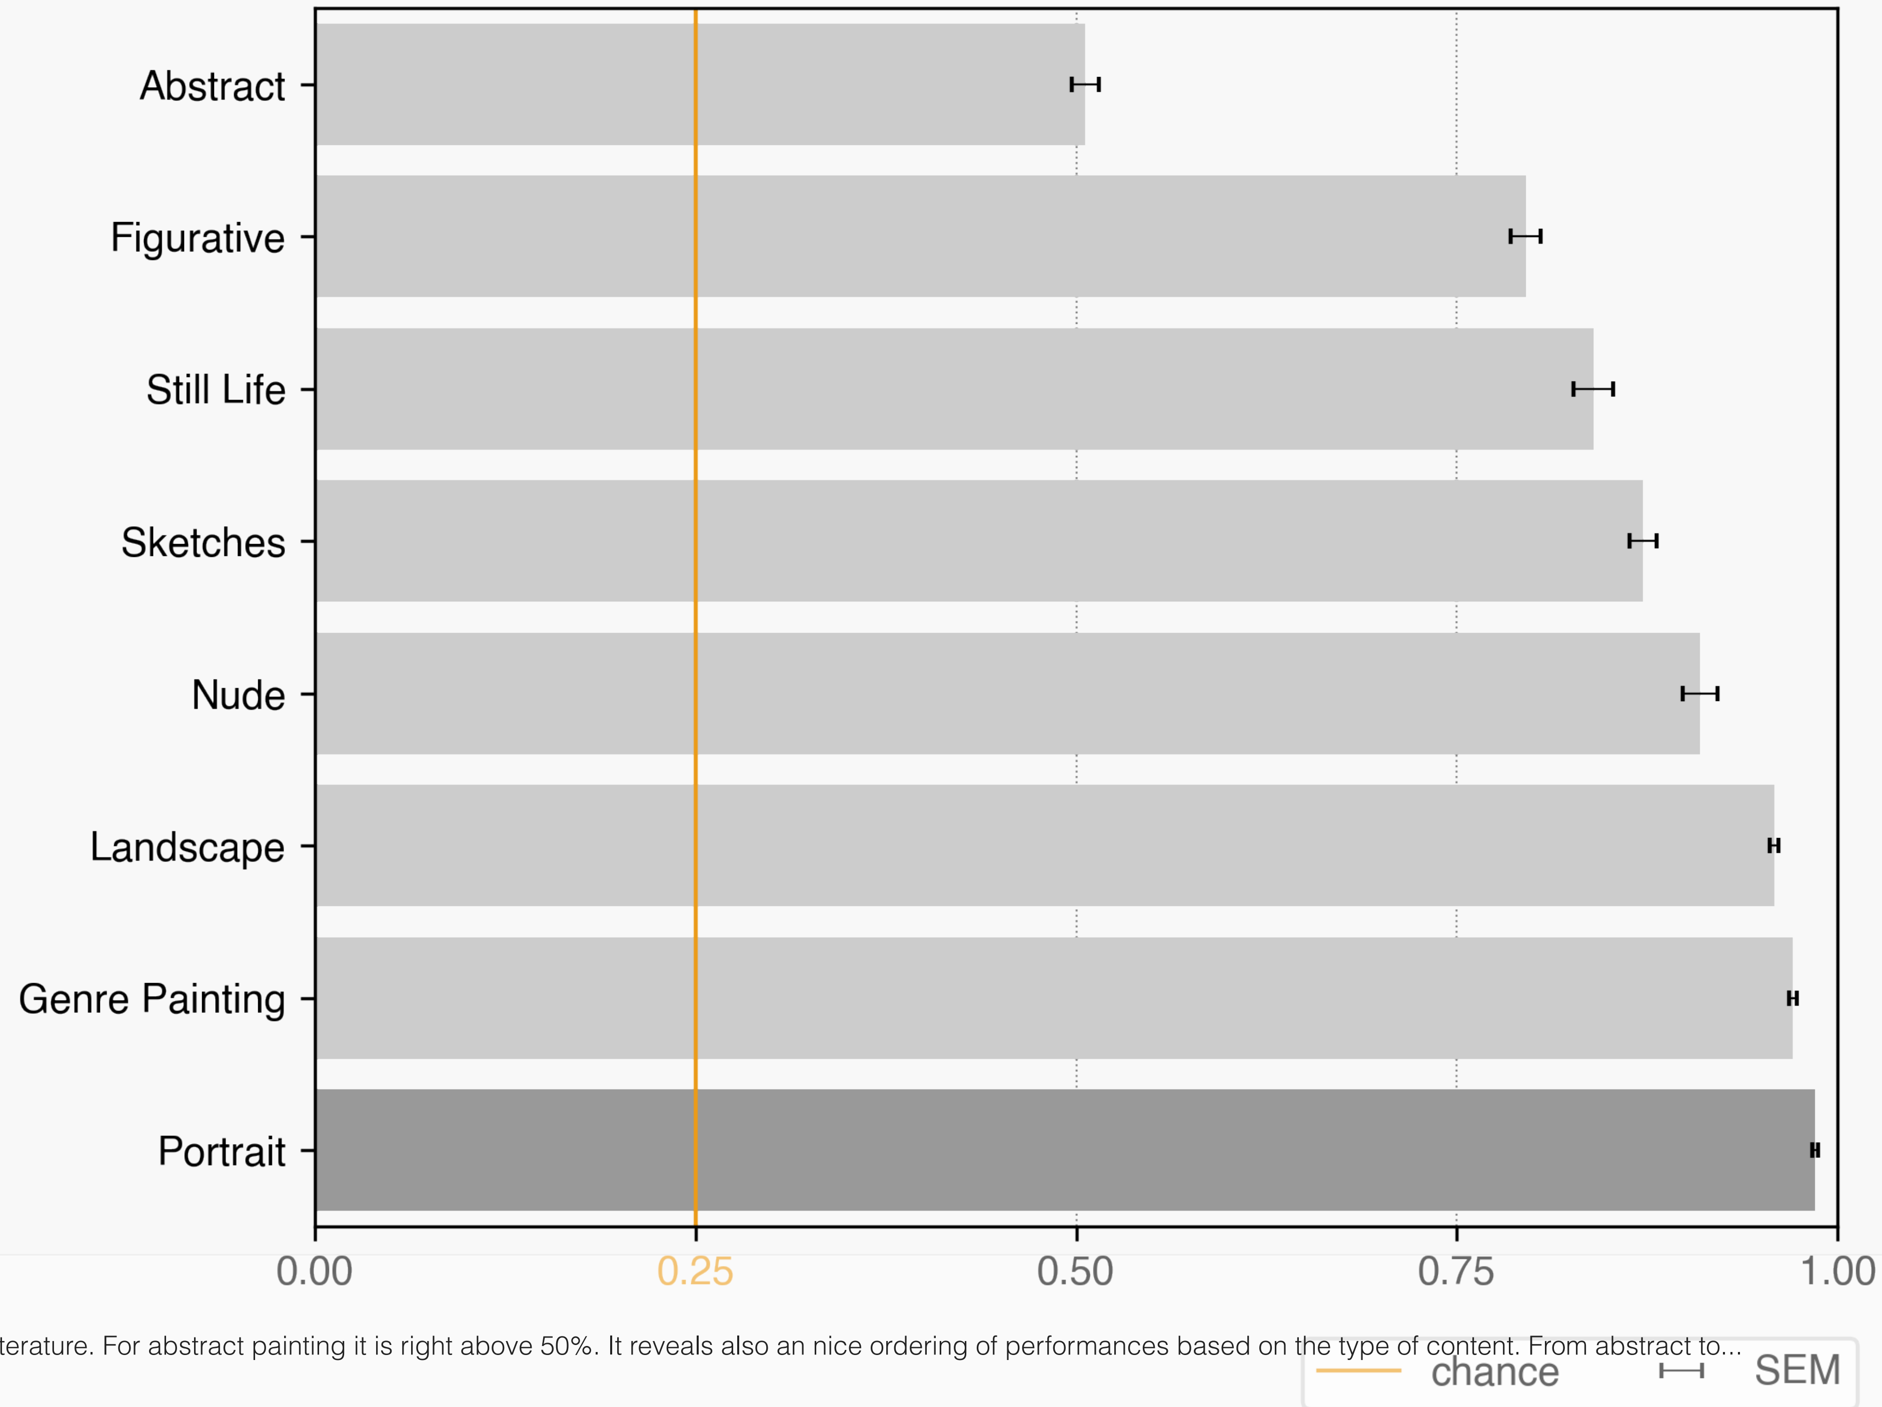

Speaker notes

- After training, performance on the whole painting are in accordance with the literature. For abstract painting it is right above 50%. It reveals also an nice ordering of performances based on the type of content. From abstract to...
- Objects.
- Landscapes.
- and Portraits.

Model Performance

Accuracy per Style

Abstract

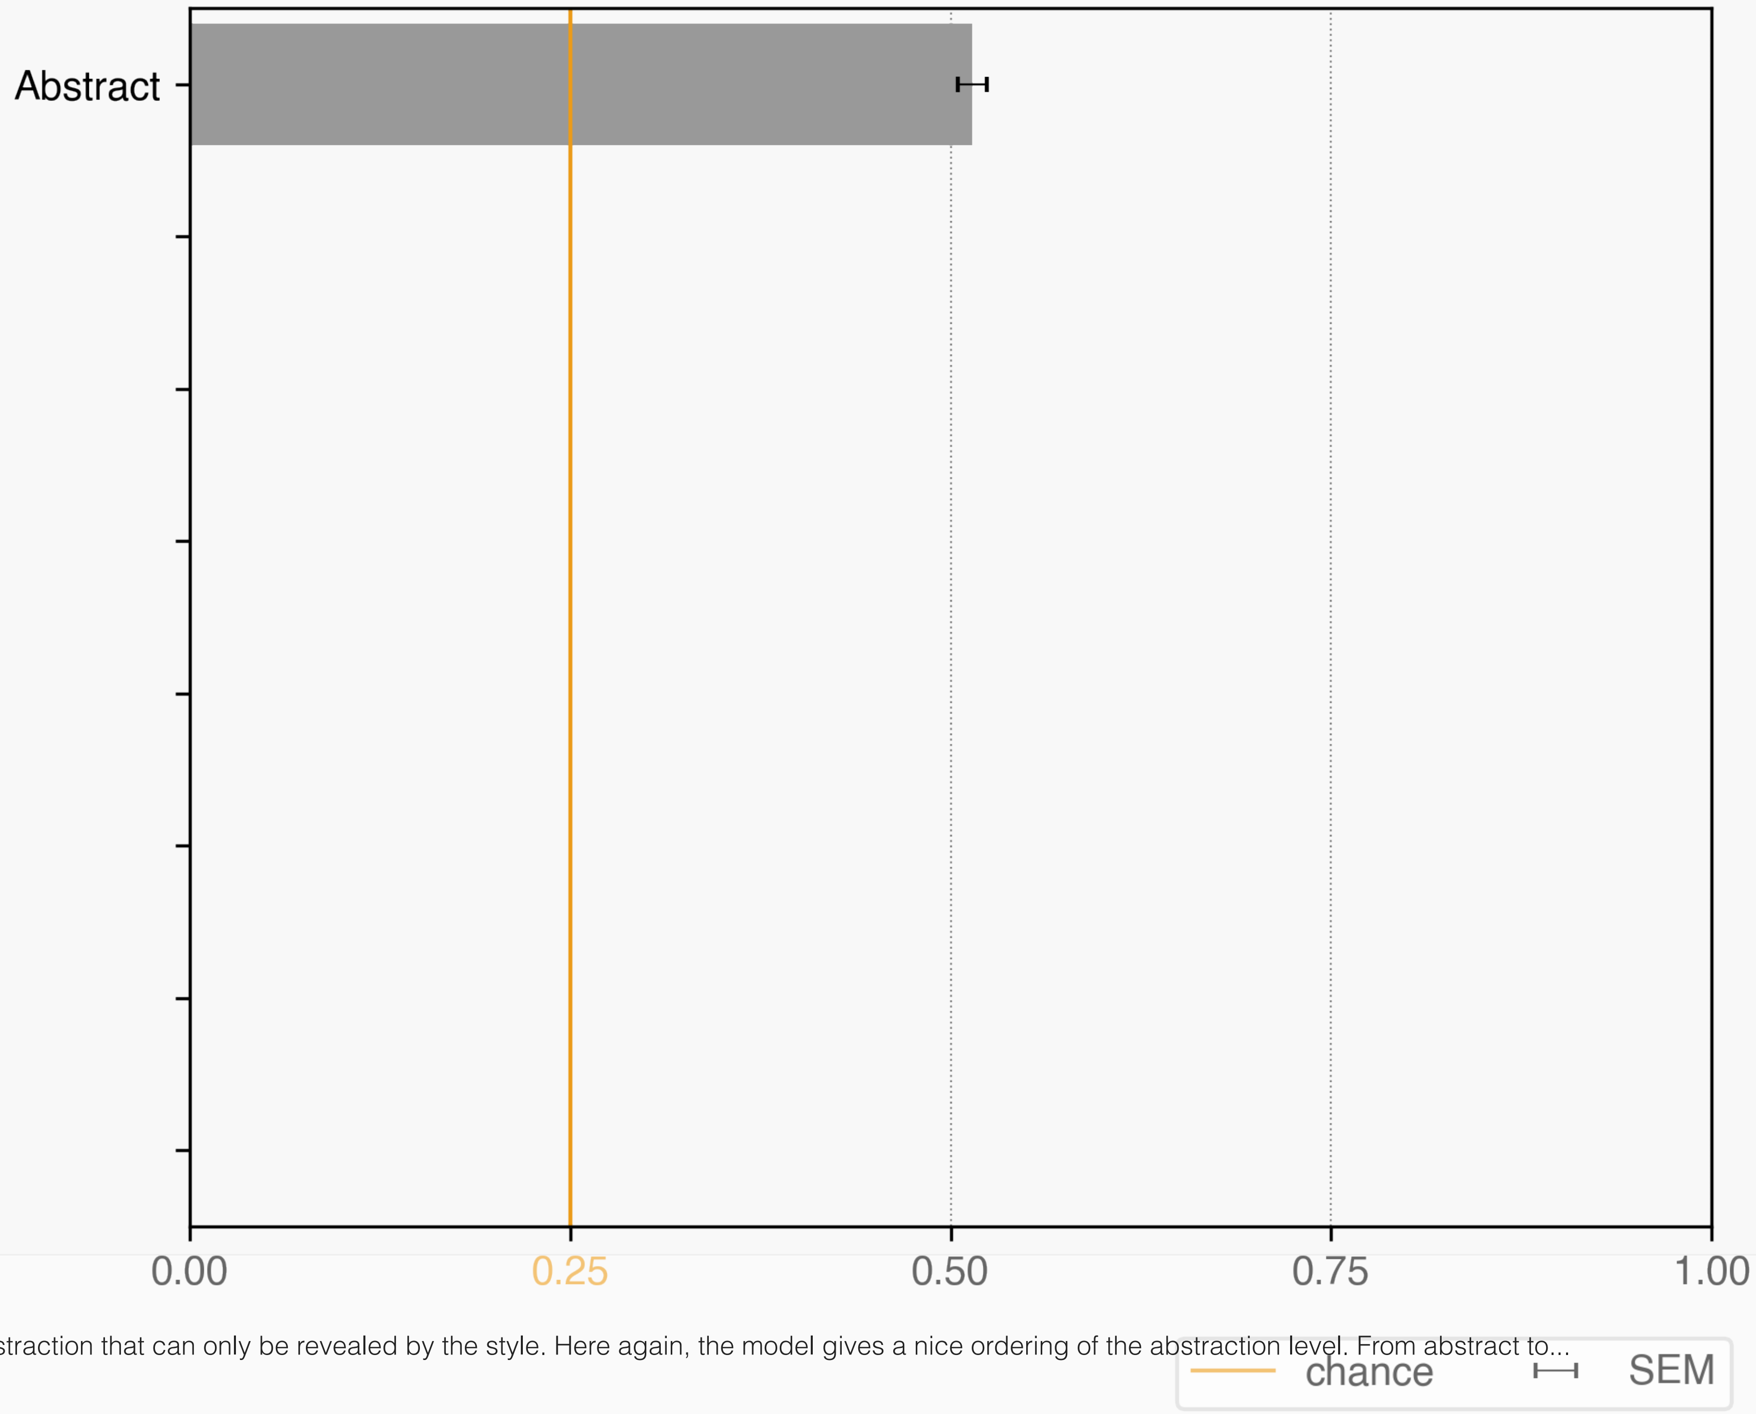

Speaker notes

- But for example, portraits from Vermeer and Picasso encompass different degrees of abstraction that can only be revealed by the style. Here again, the model gives a nice ordering of the abstraction level. From abstract to...
- Cubism.
- Symbolism.
- and Post-Renaissance's Romanticism.

Model Performance

Accuracy per Style

Cubism

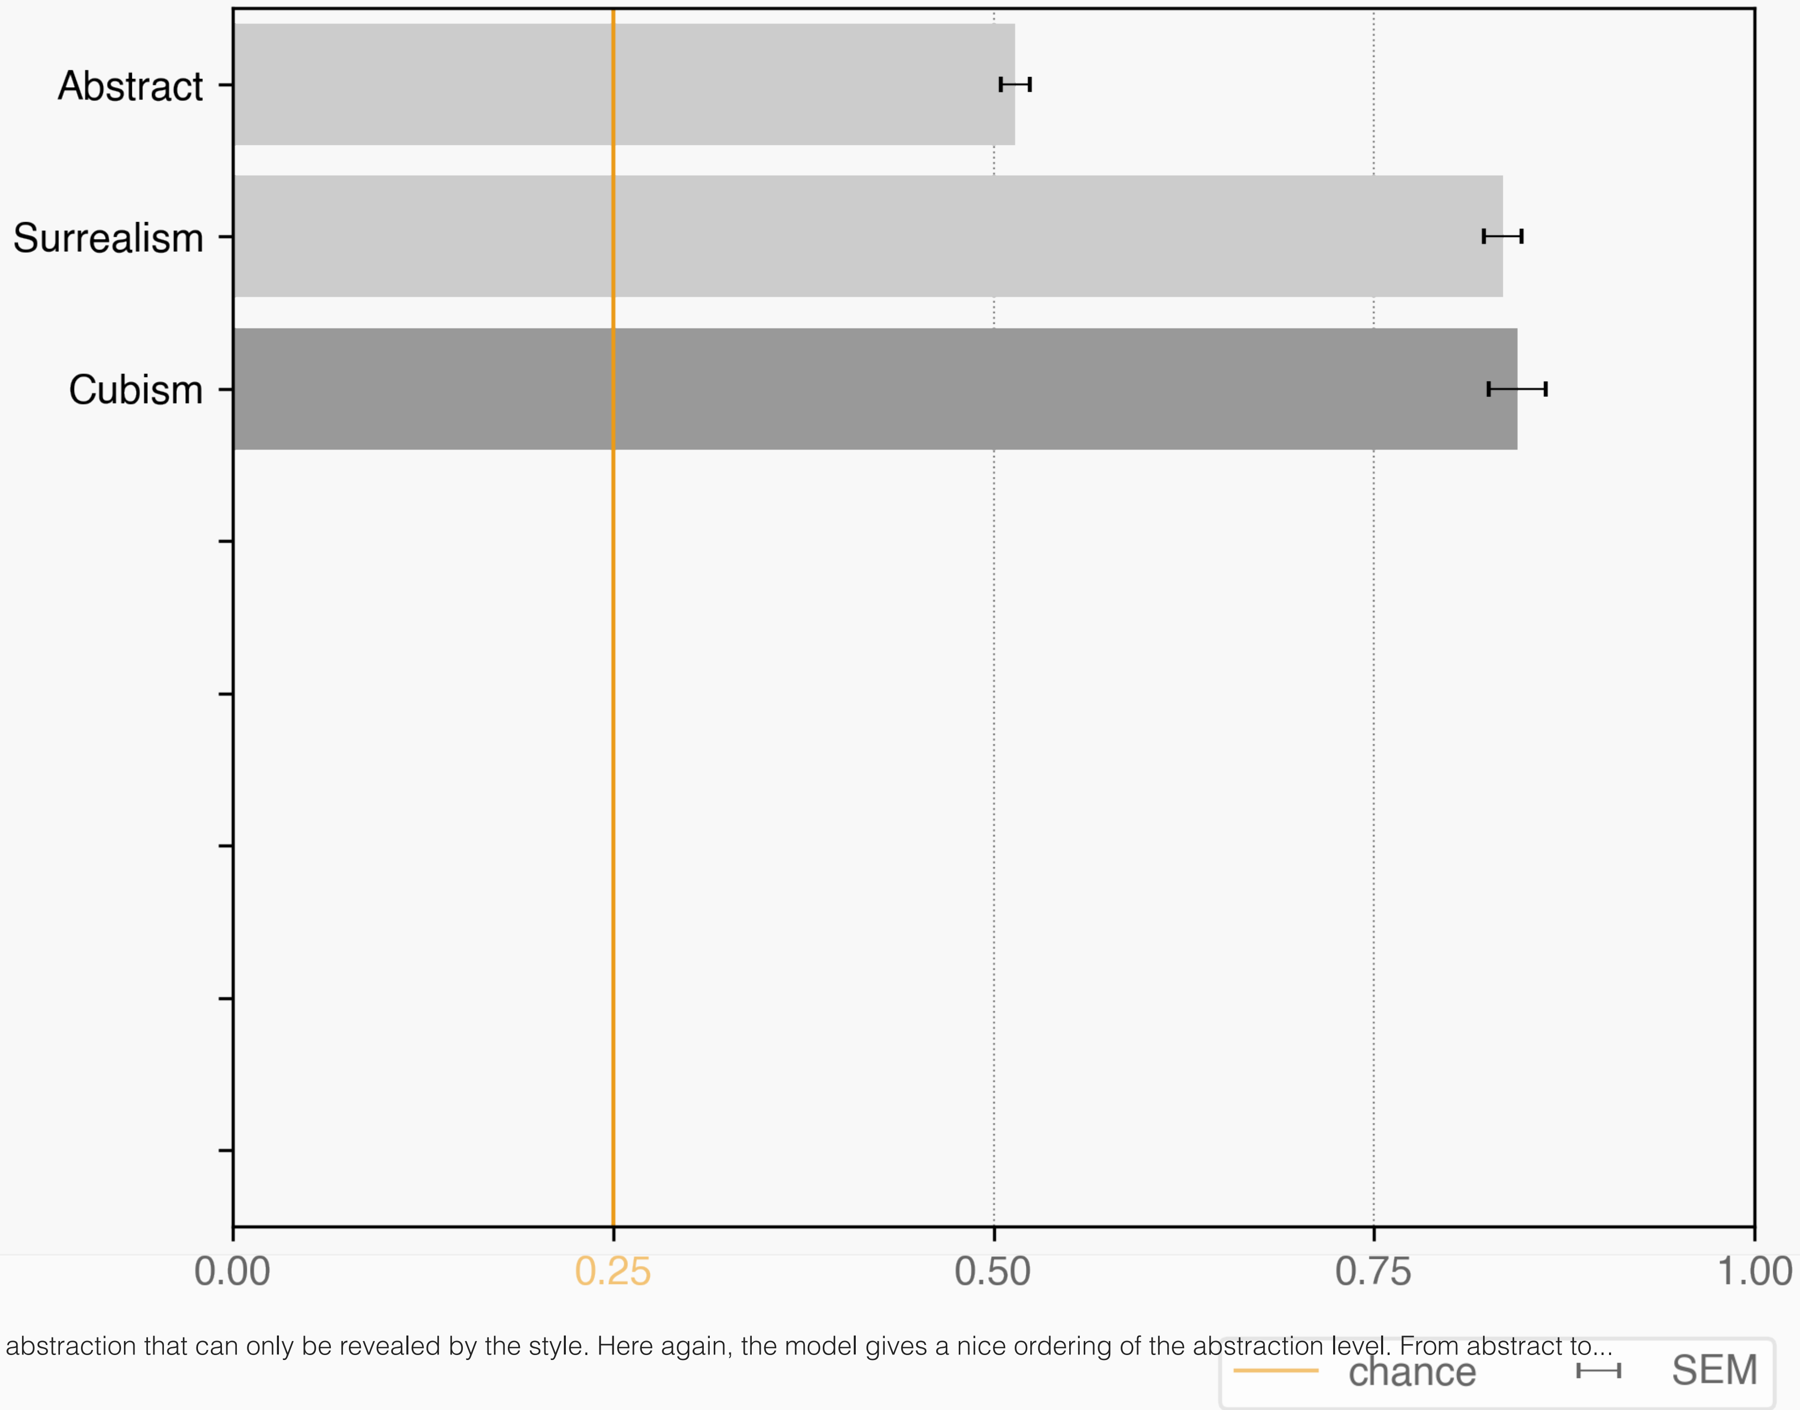

Speaker notes

- But for example, portraits from Vermeer and Picasso encompass different degrees of abstraction that can only be revealed by the style. Here again, the model gives a nice ordering of the abstraction level. From abstract to...
- Cubism.
- Symbolism.
- and Post-Renaissance's Romanticism.

Model Performance

Accuracy per Style

Symbolism

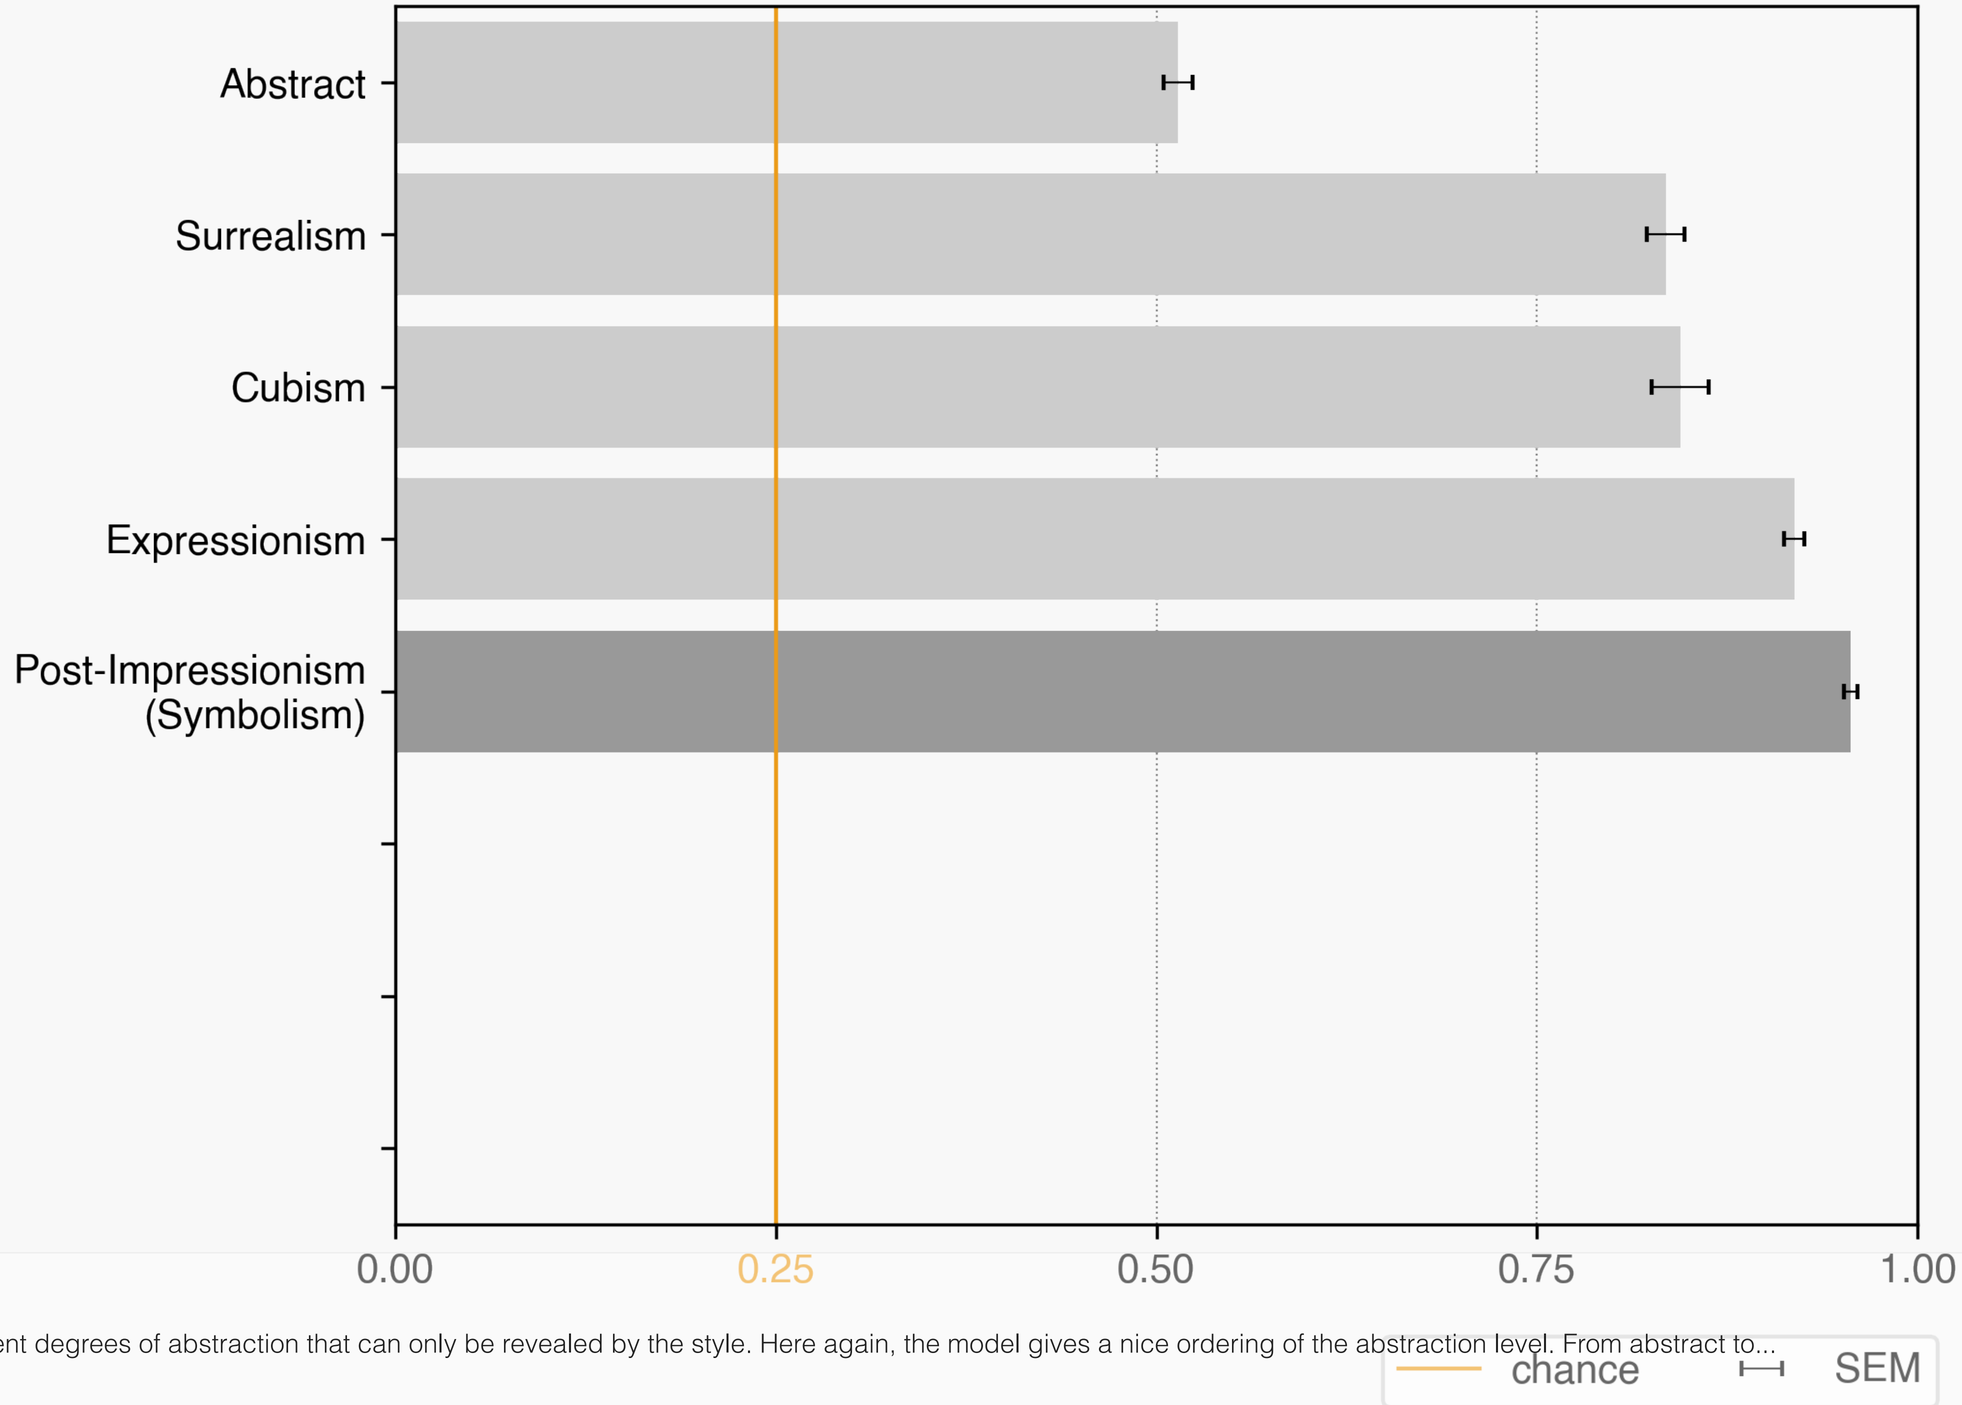

Speaker notes

- But for example, portraits from Vermeer and Picasso encompass different degrees of abstraction that can only be revealed by the style. Here again, the model gives a nice ordering of the abstraction level. From abstract to...
- Cubism.
- Symbolism.
- and Post-Renaissance's Romanticism.

Model Performance

Accuracy per Style

Romanticism

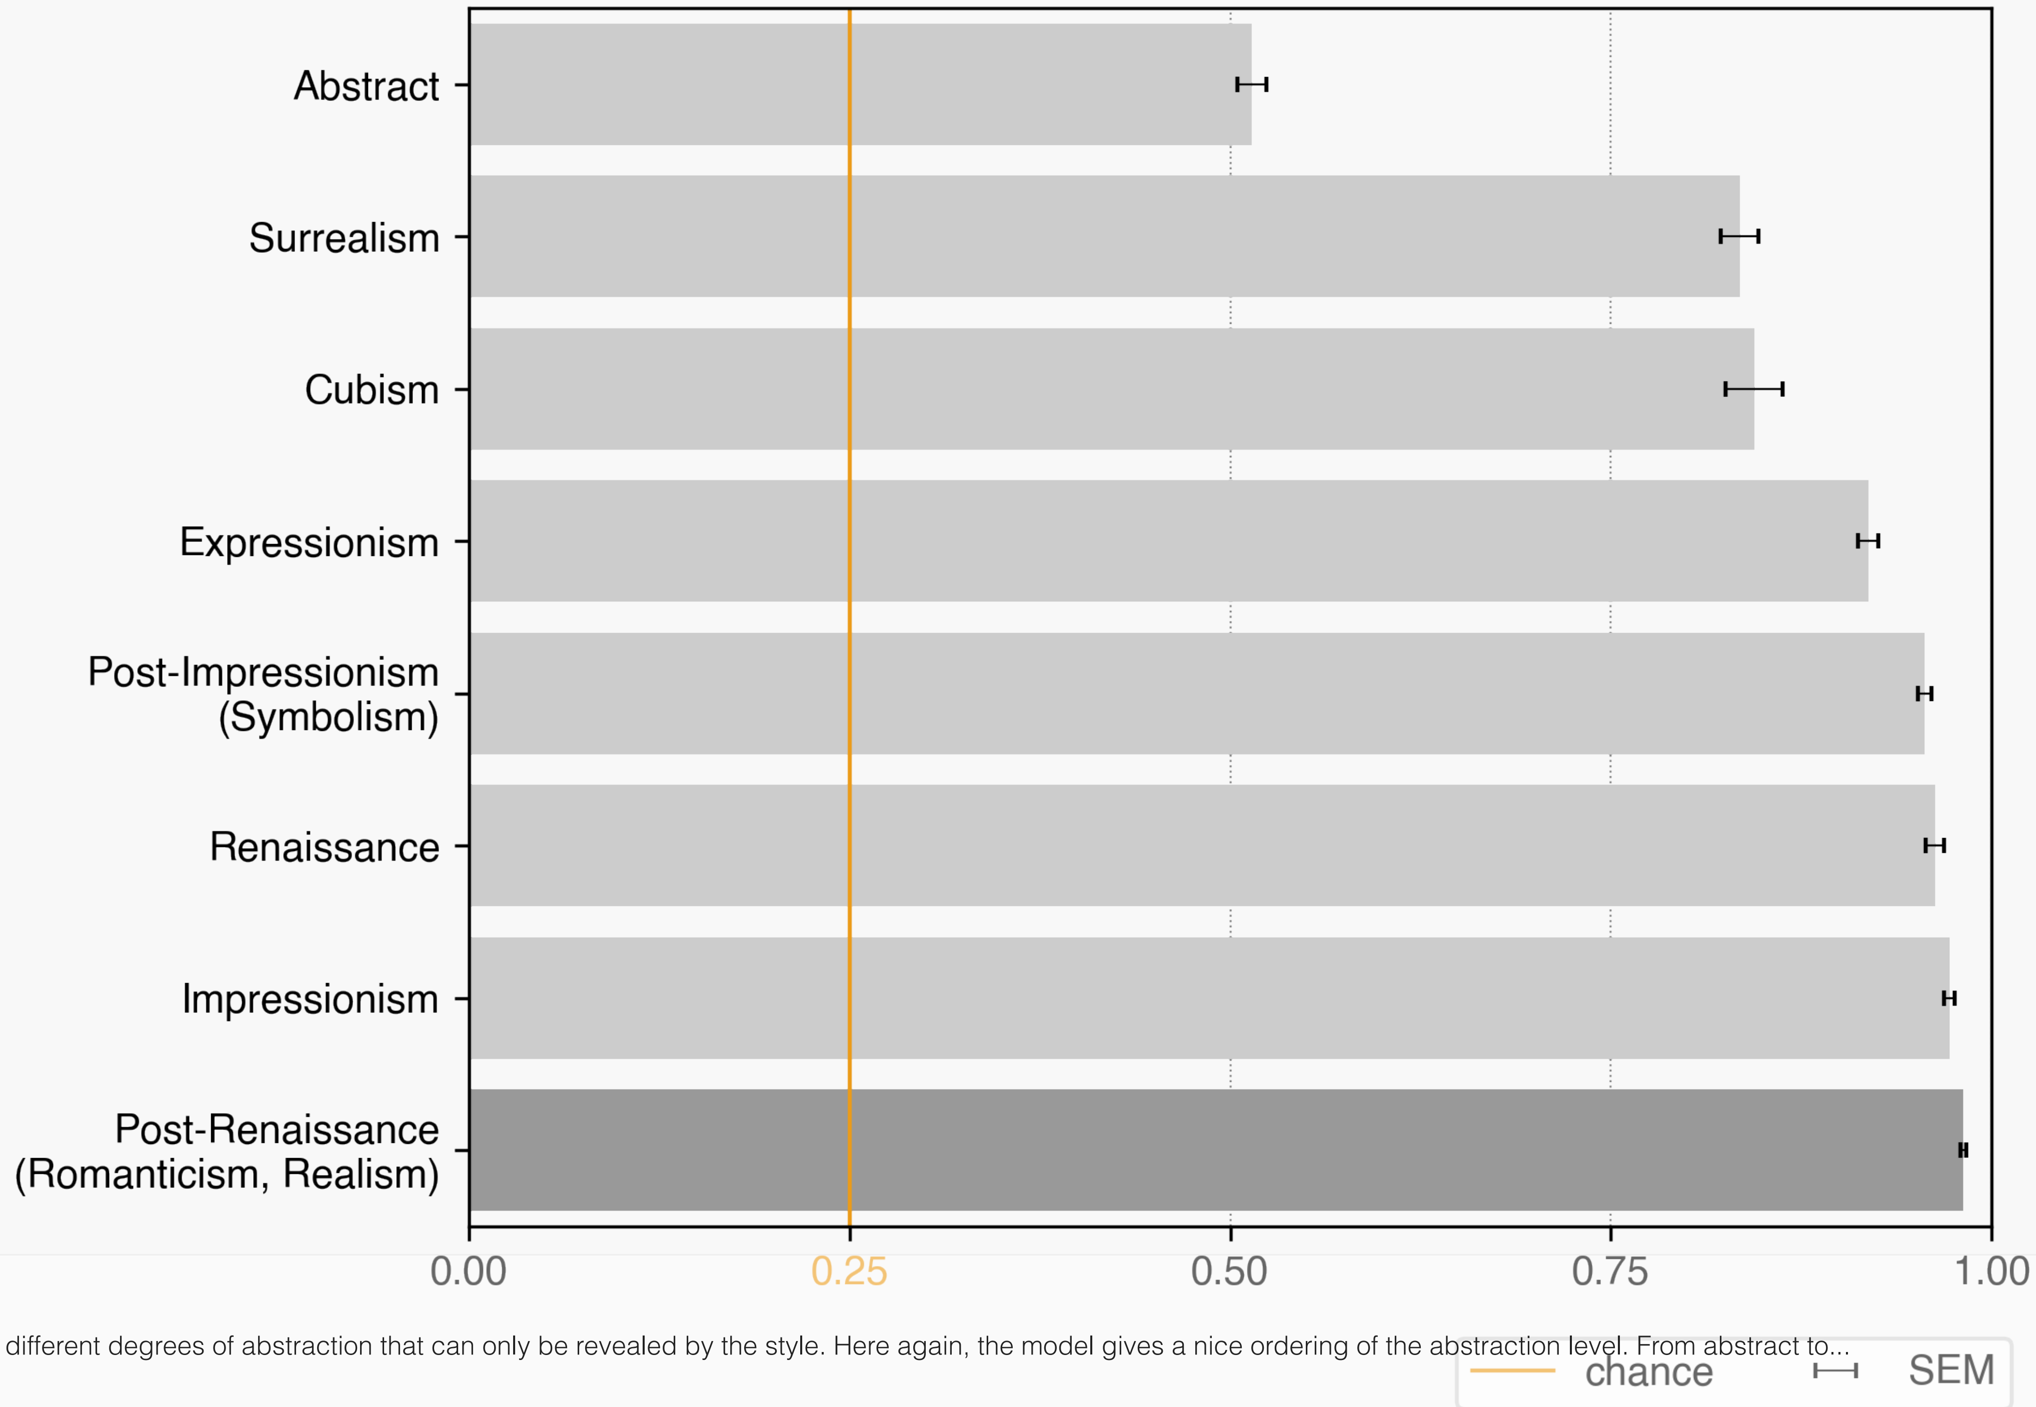

Speaker notes

- But for example, portraits from Vermeer and Picasso encompass different degrees of abstraction that can only be revealed by the style. Here again, the model gives a nice ordering of the abstraction level. From abstract to...
- Cubism.
- Symbolism.
- and Post-Renaissance's Romanticism.

Model Performance

Accuracy per Layer

Layer 5

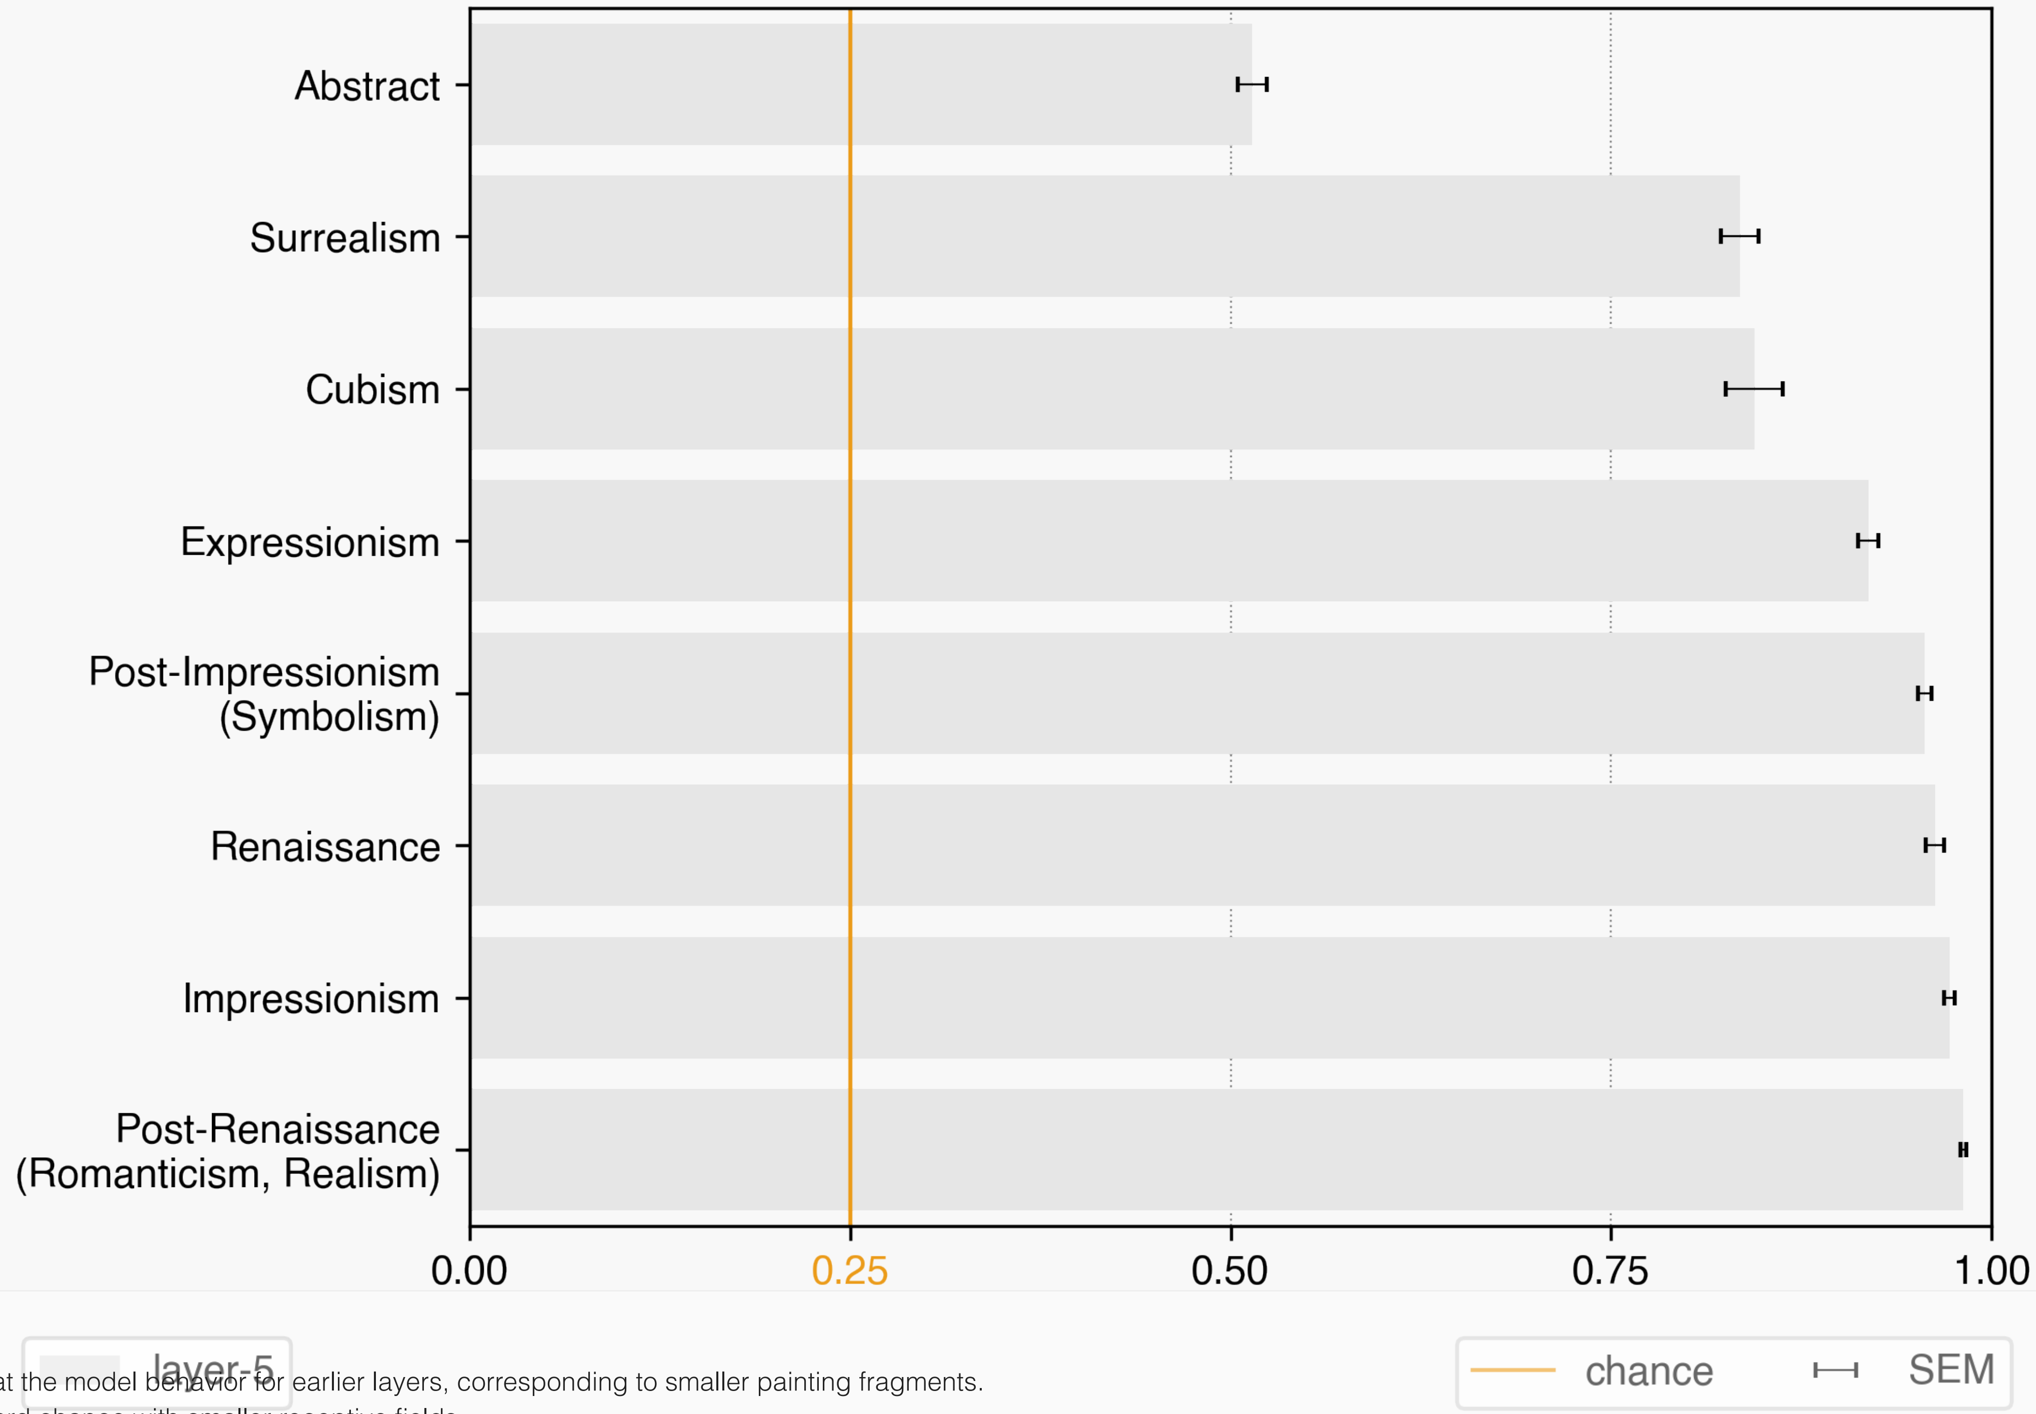

Speaker notes

- These results were for the whole paintings. But, lets have a look at the model behavior for earlier layers, corresponding to smaller painting fragments.
- As we can expect, the model gives performances that shifts toward chance with smaller receptive fields.
- And if we scale the results between chance and performance at layer-5, it seems that abstract styles depend more on deeper layers than more realistic styles.

Model Performance

Accuracy per Layer

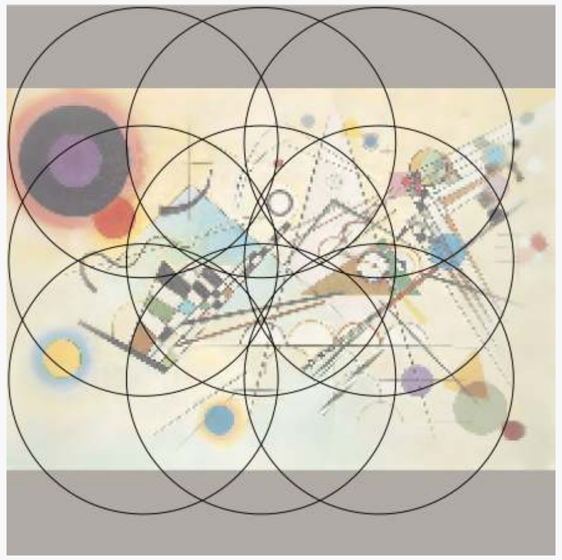

Layer 4

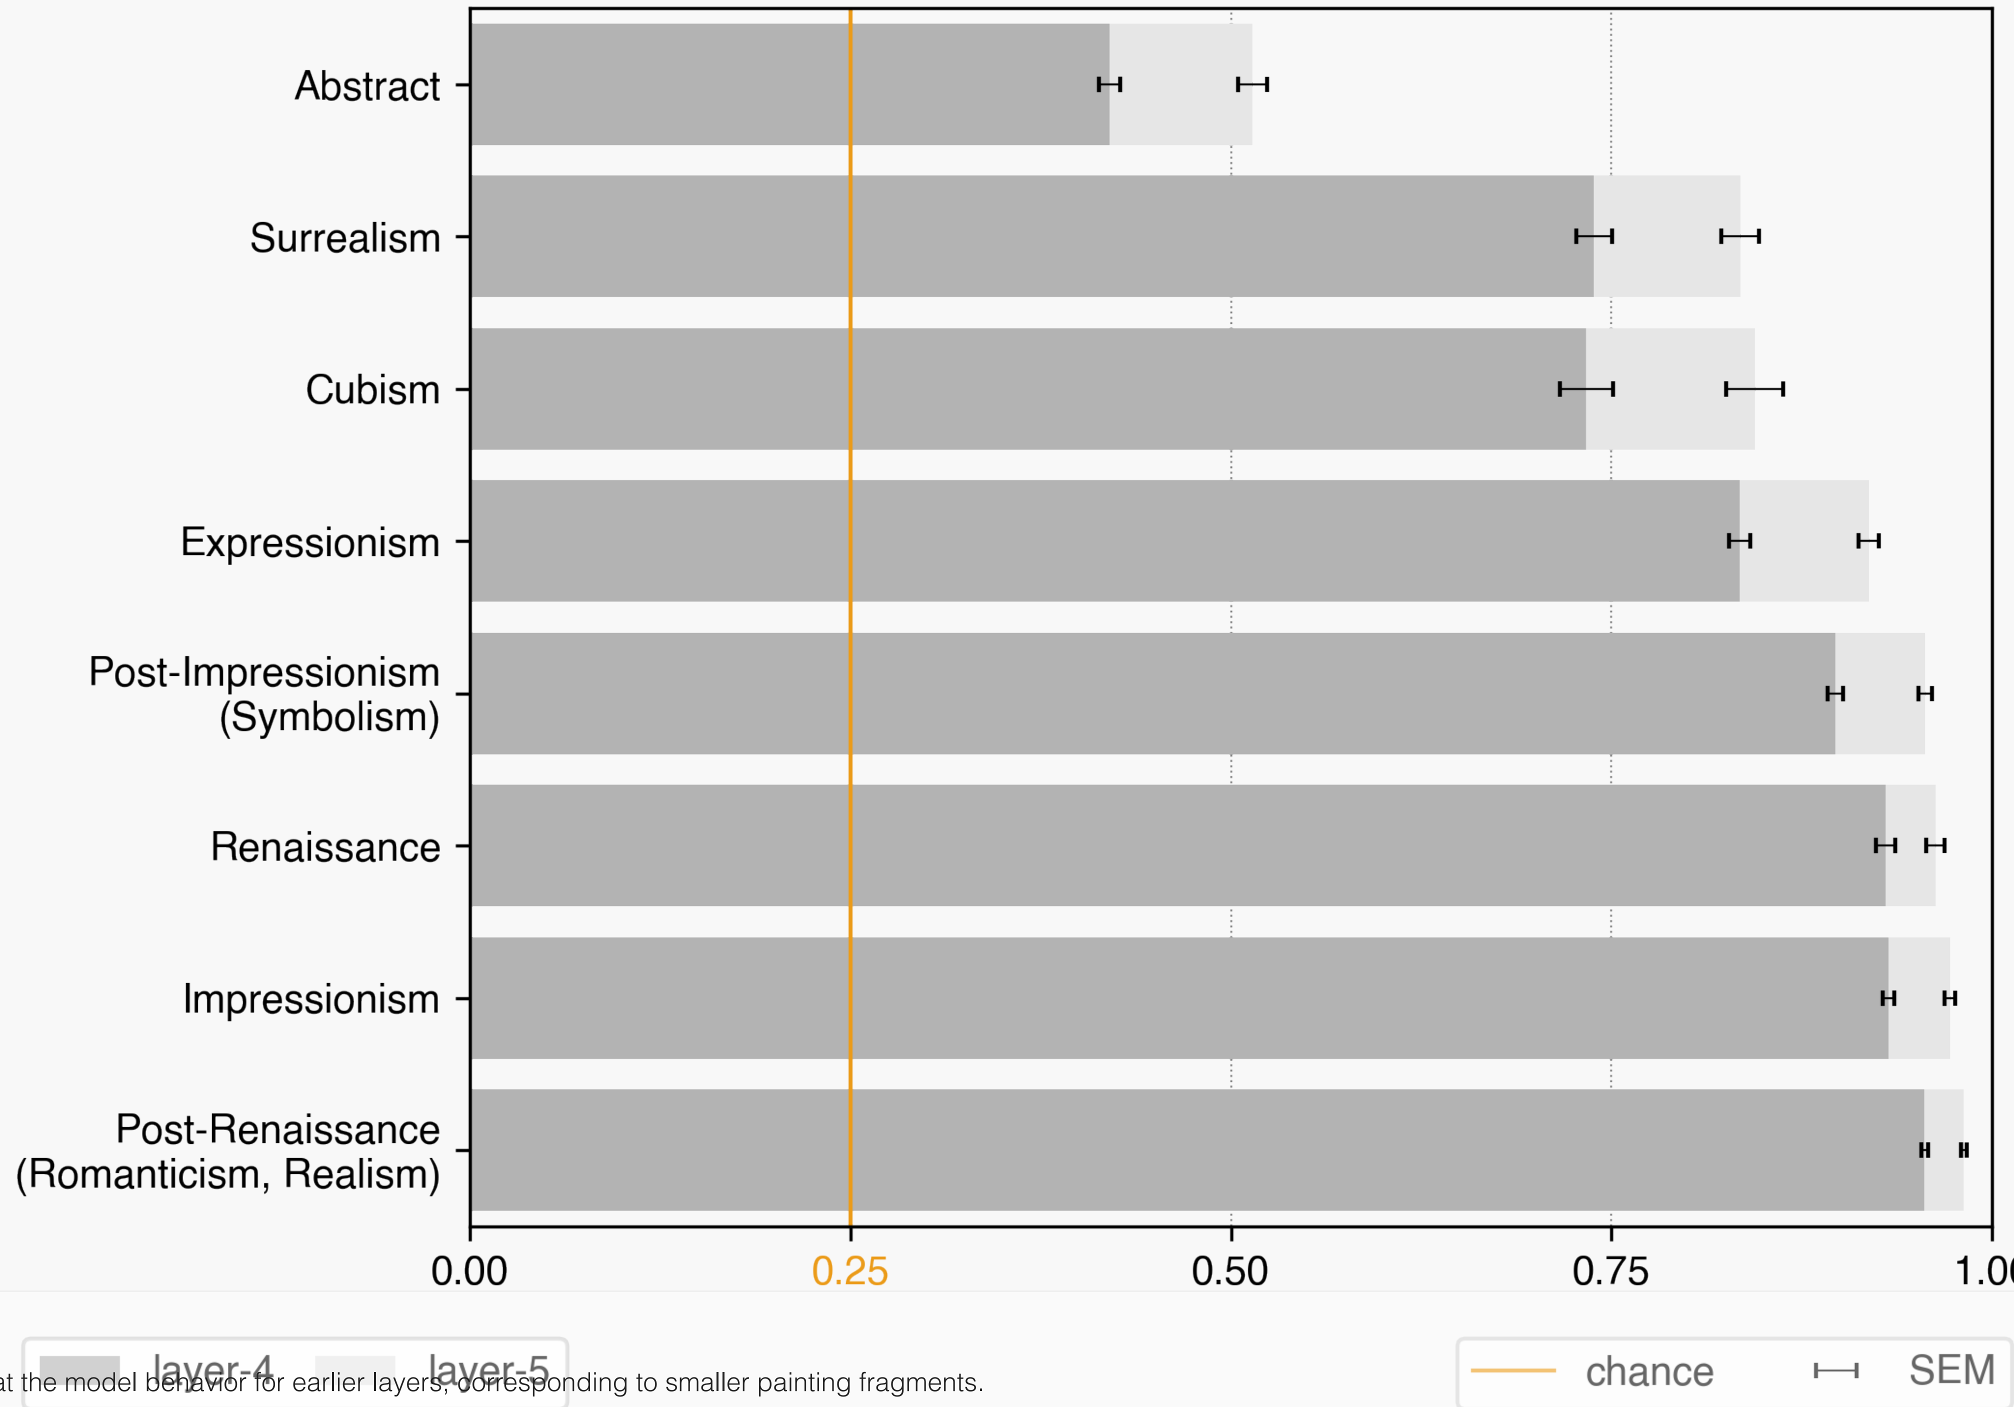

Speaker notes

- These results were for the whole paintings. But, lets have a look at the model behavior for earlier layers, corresponding to smaller painting fragments.
- As we can expect, the model gives performances that shifts toward chance with smaller receptive fields.
- And if we scale the results between chance and performance at layer-5, it seems that abstract styles depend more on deeper layers than more realistic styles.

Model Performance

Accuracy per Layer

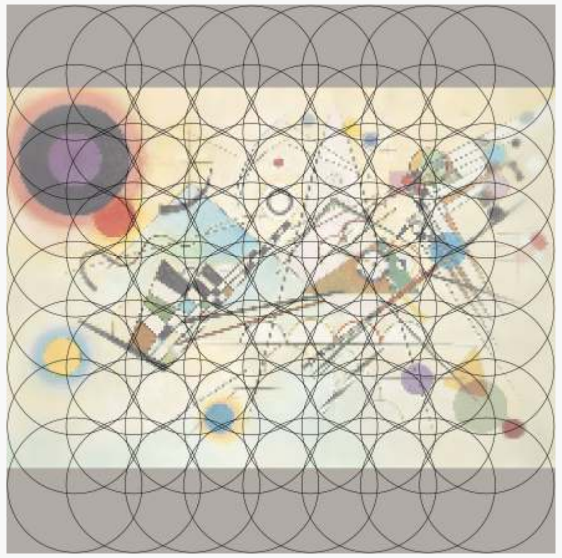

Layer 3

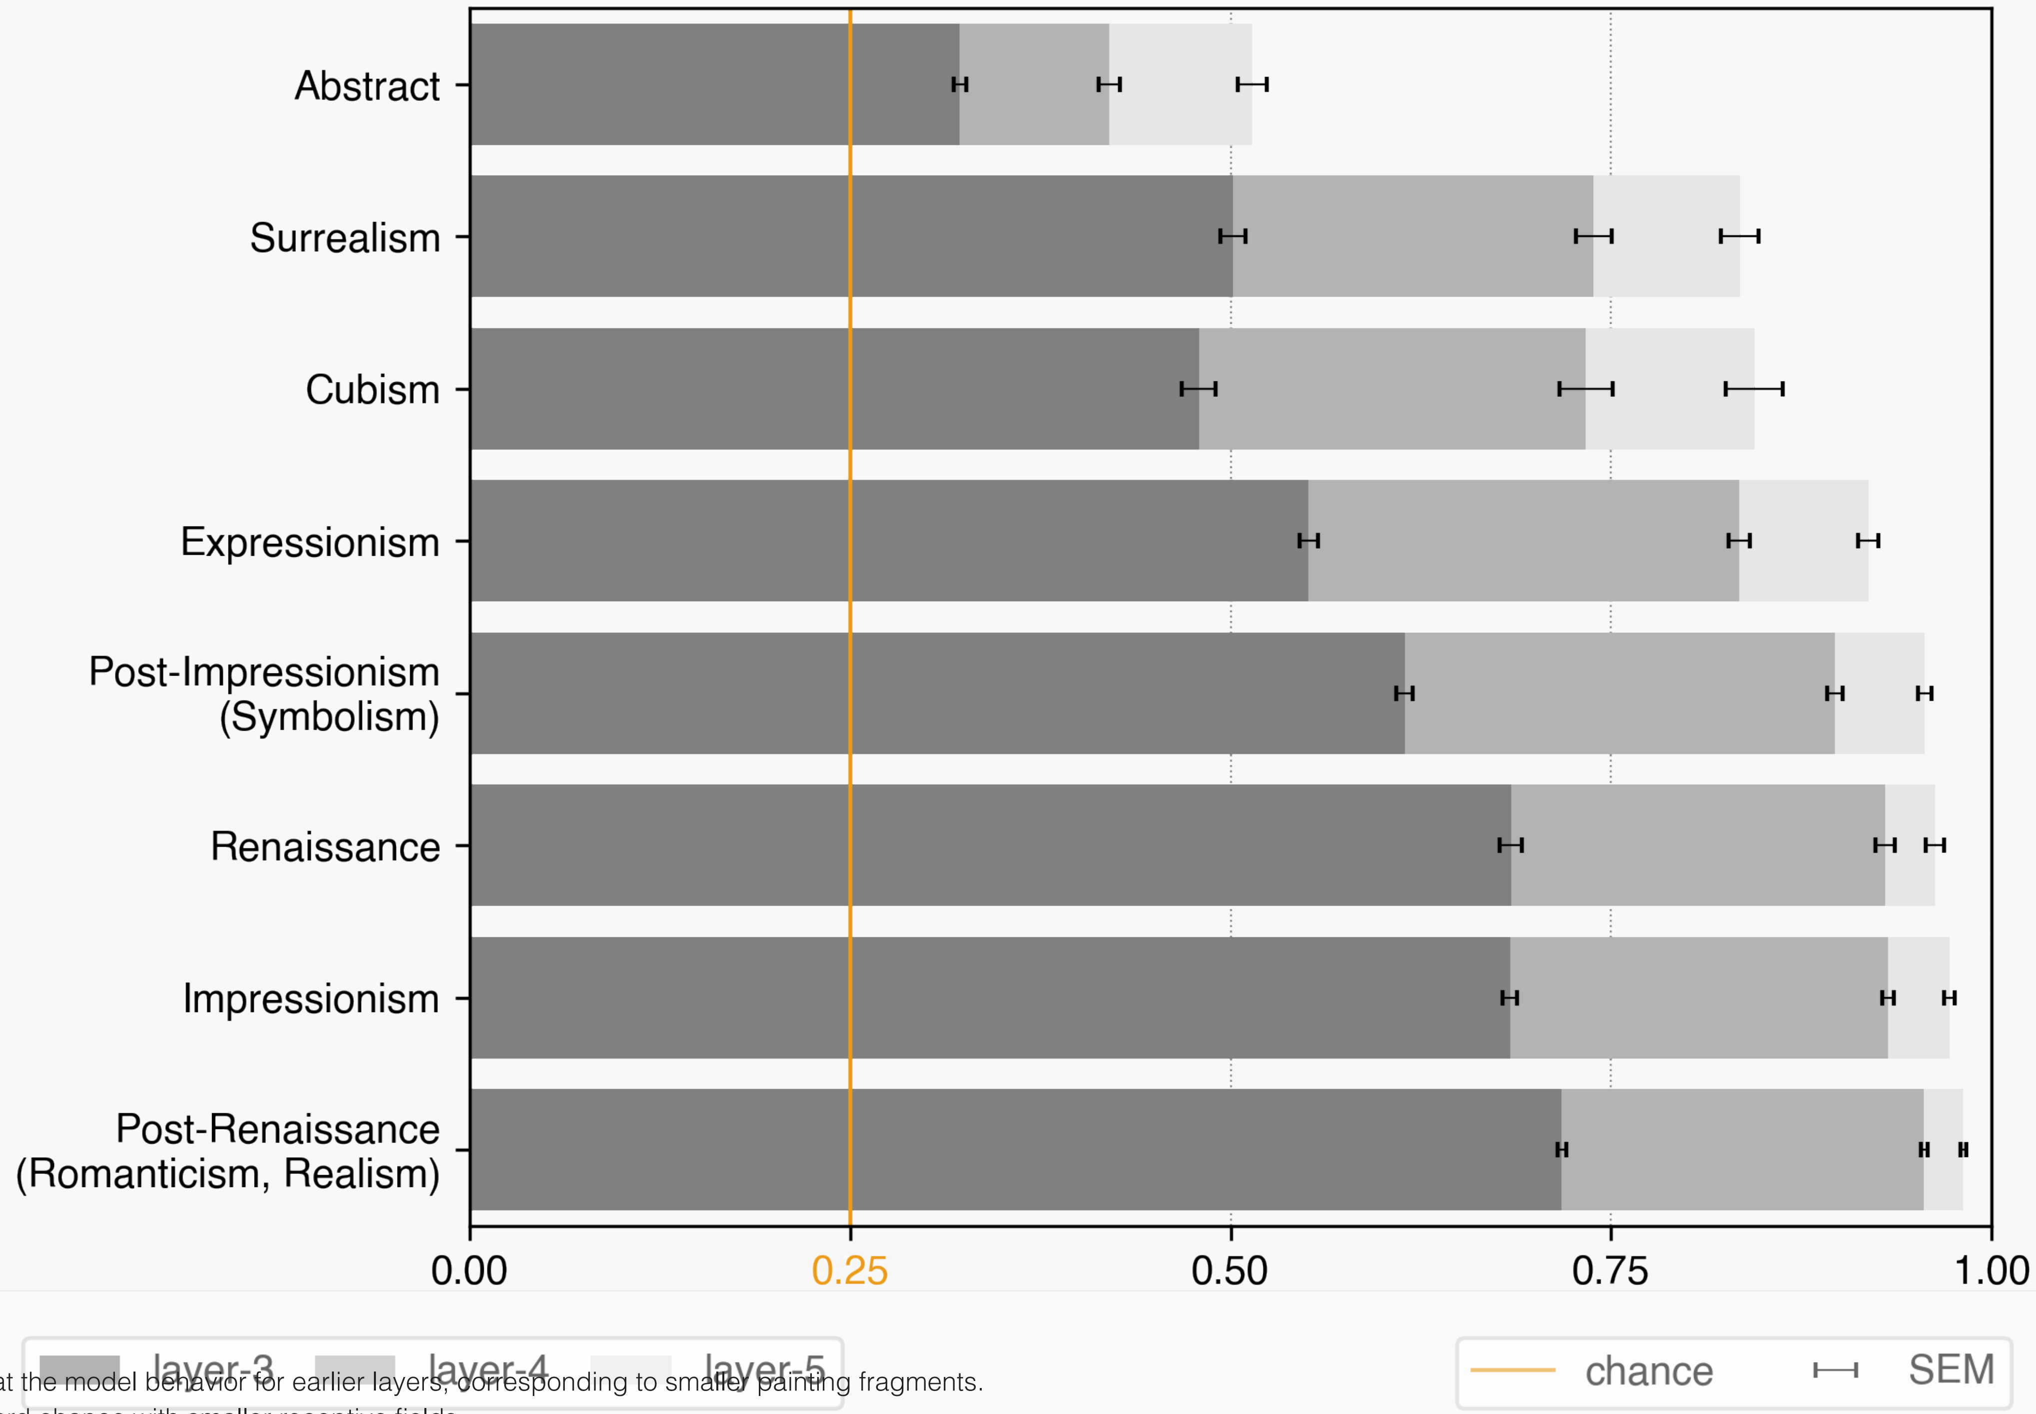

Speaker notes

- These results were for the whole paintings. But, lets have a look at the model behavior for earlier layers, corresponding to smaller painting fragments.
- As we can expect, the model gives performances that shifts toward chance with smaller receptive fields.
- And if we scale the results between chance and performance at layer-5, it seems that abstract styles depend more on deeper layers than more realistic styles.

Model Performance

Accuracy per Layer

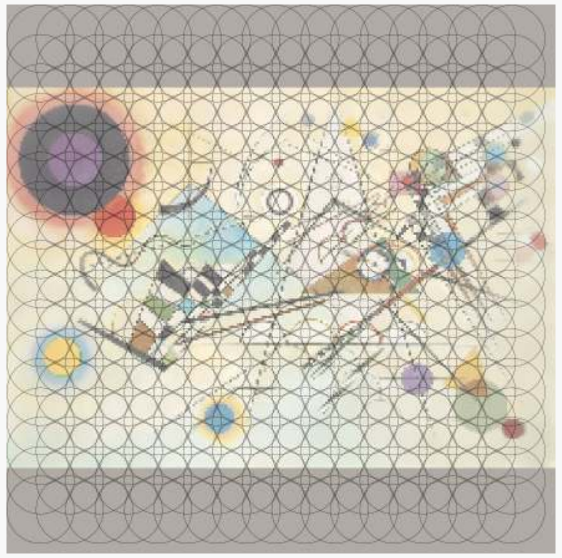

Layer 2

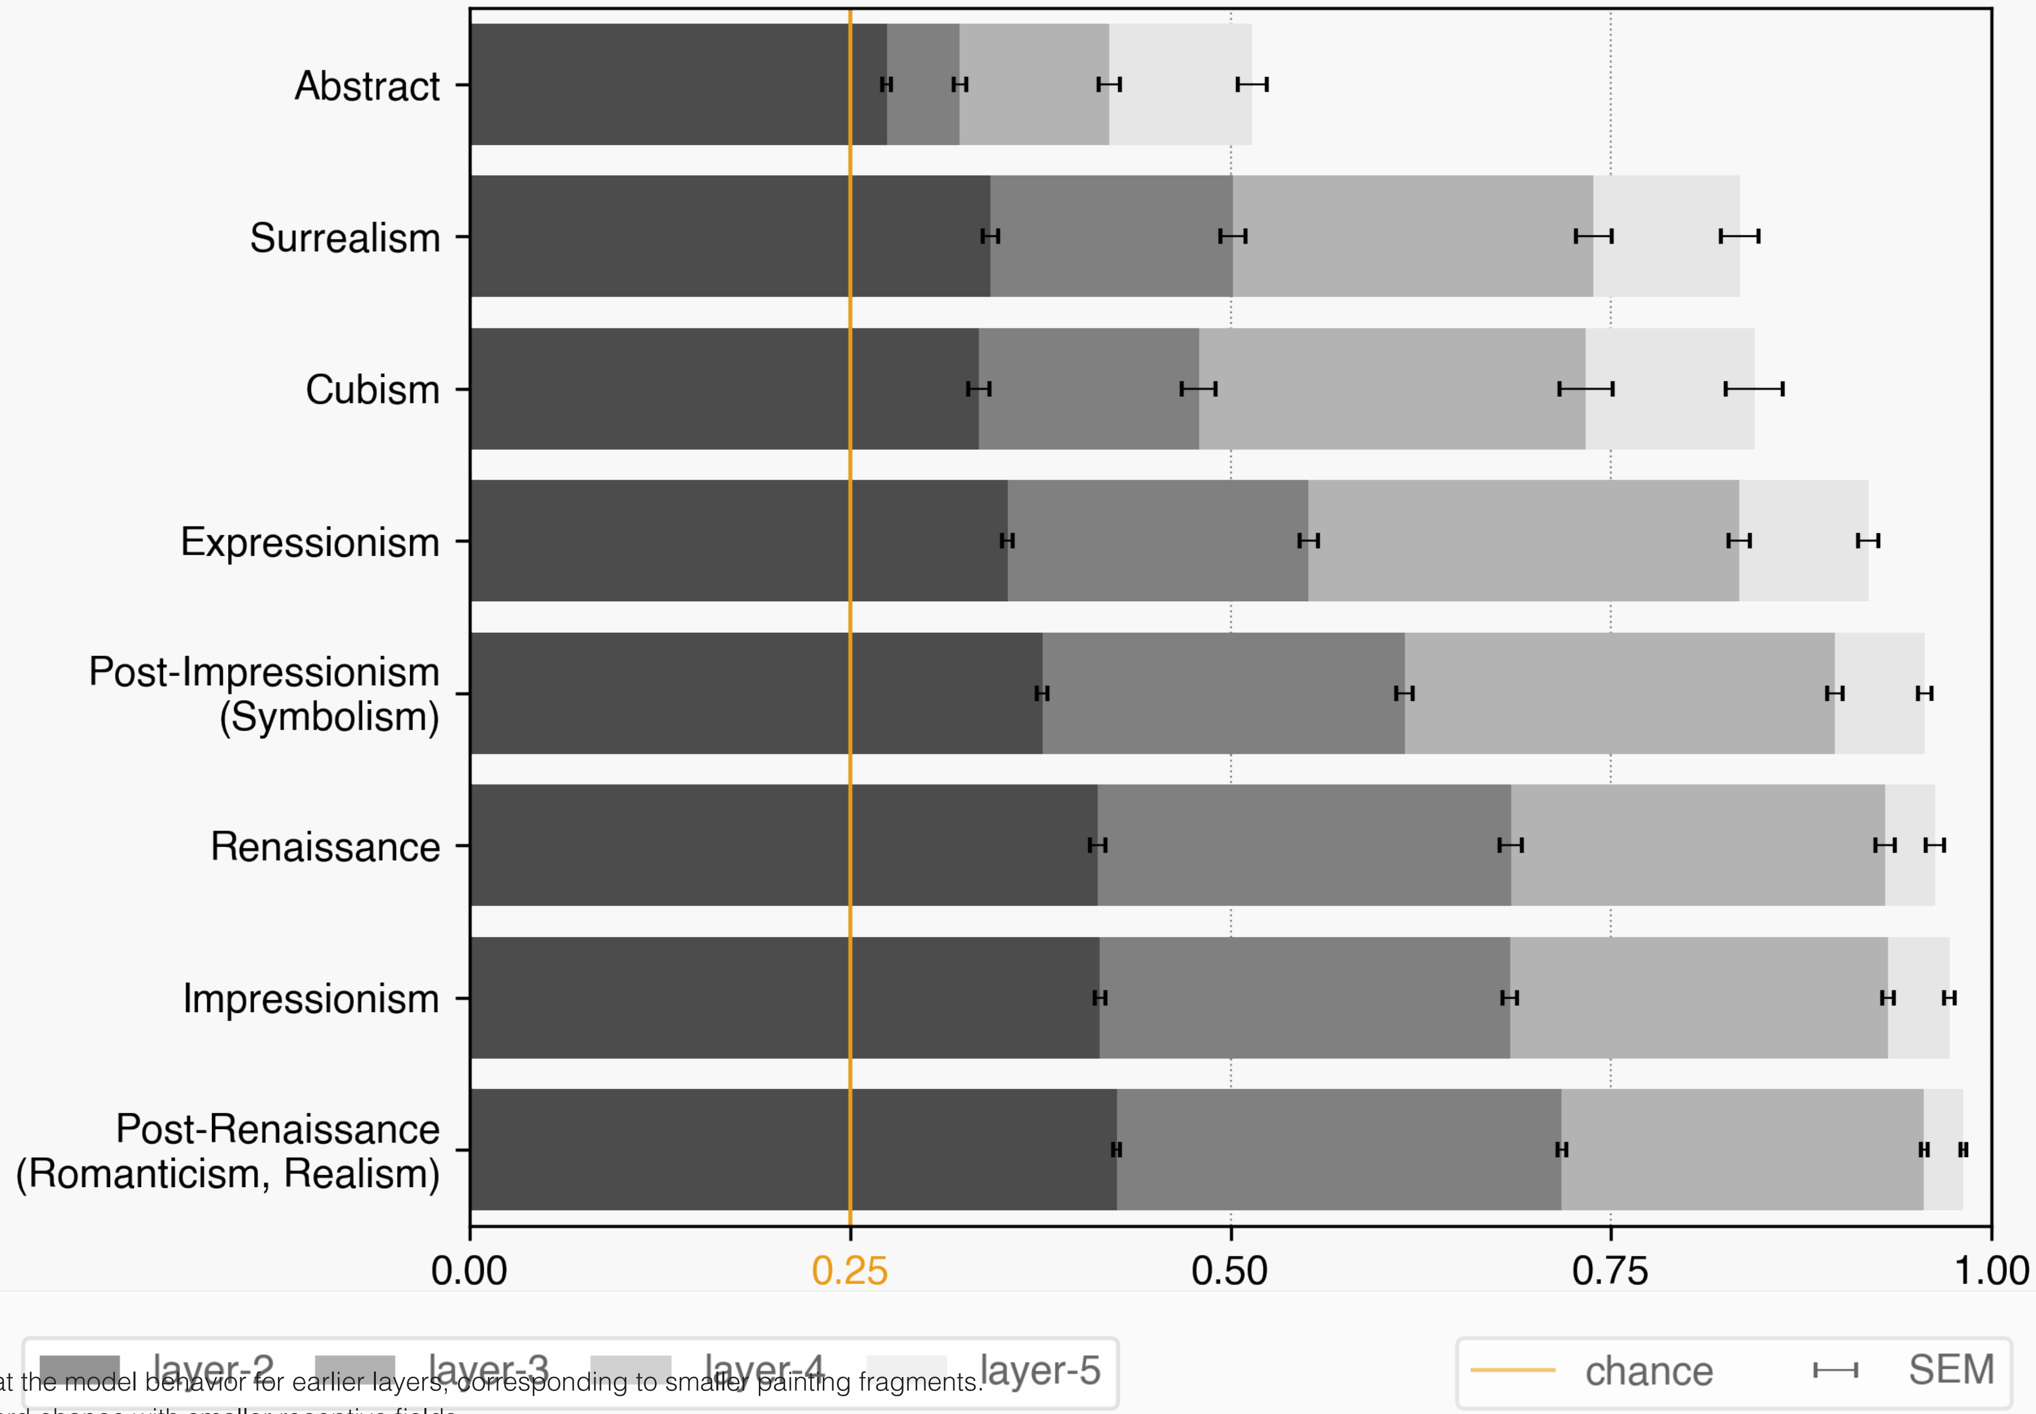

Speaker notes

- These results were for the whole paintings. But, lets have a look at the model behavior for earlier layers, corresponding to smaller painting fragments.
- As we can expect, the model gives performances that shifts toward chance with smaller receptive fields.
- And if we scale the results between chance and performance at layer-5, it seems that abstract styles depend more on deeper layers than more realistic styles.

Model Performance

Accuracy per Layer

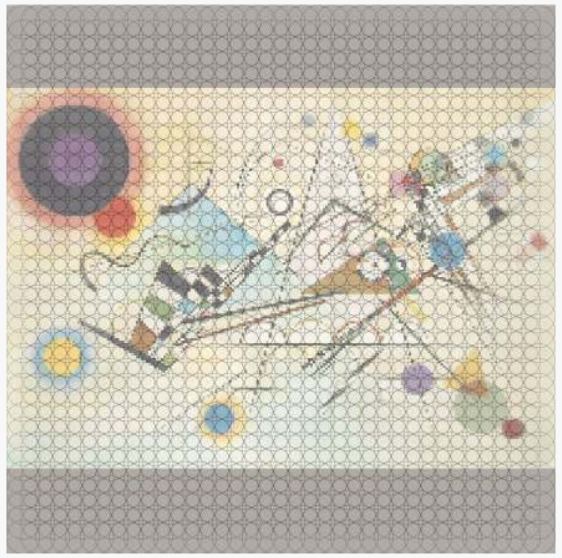

Layer 1

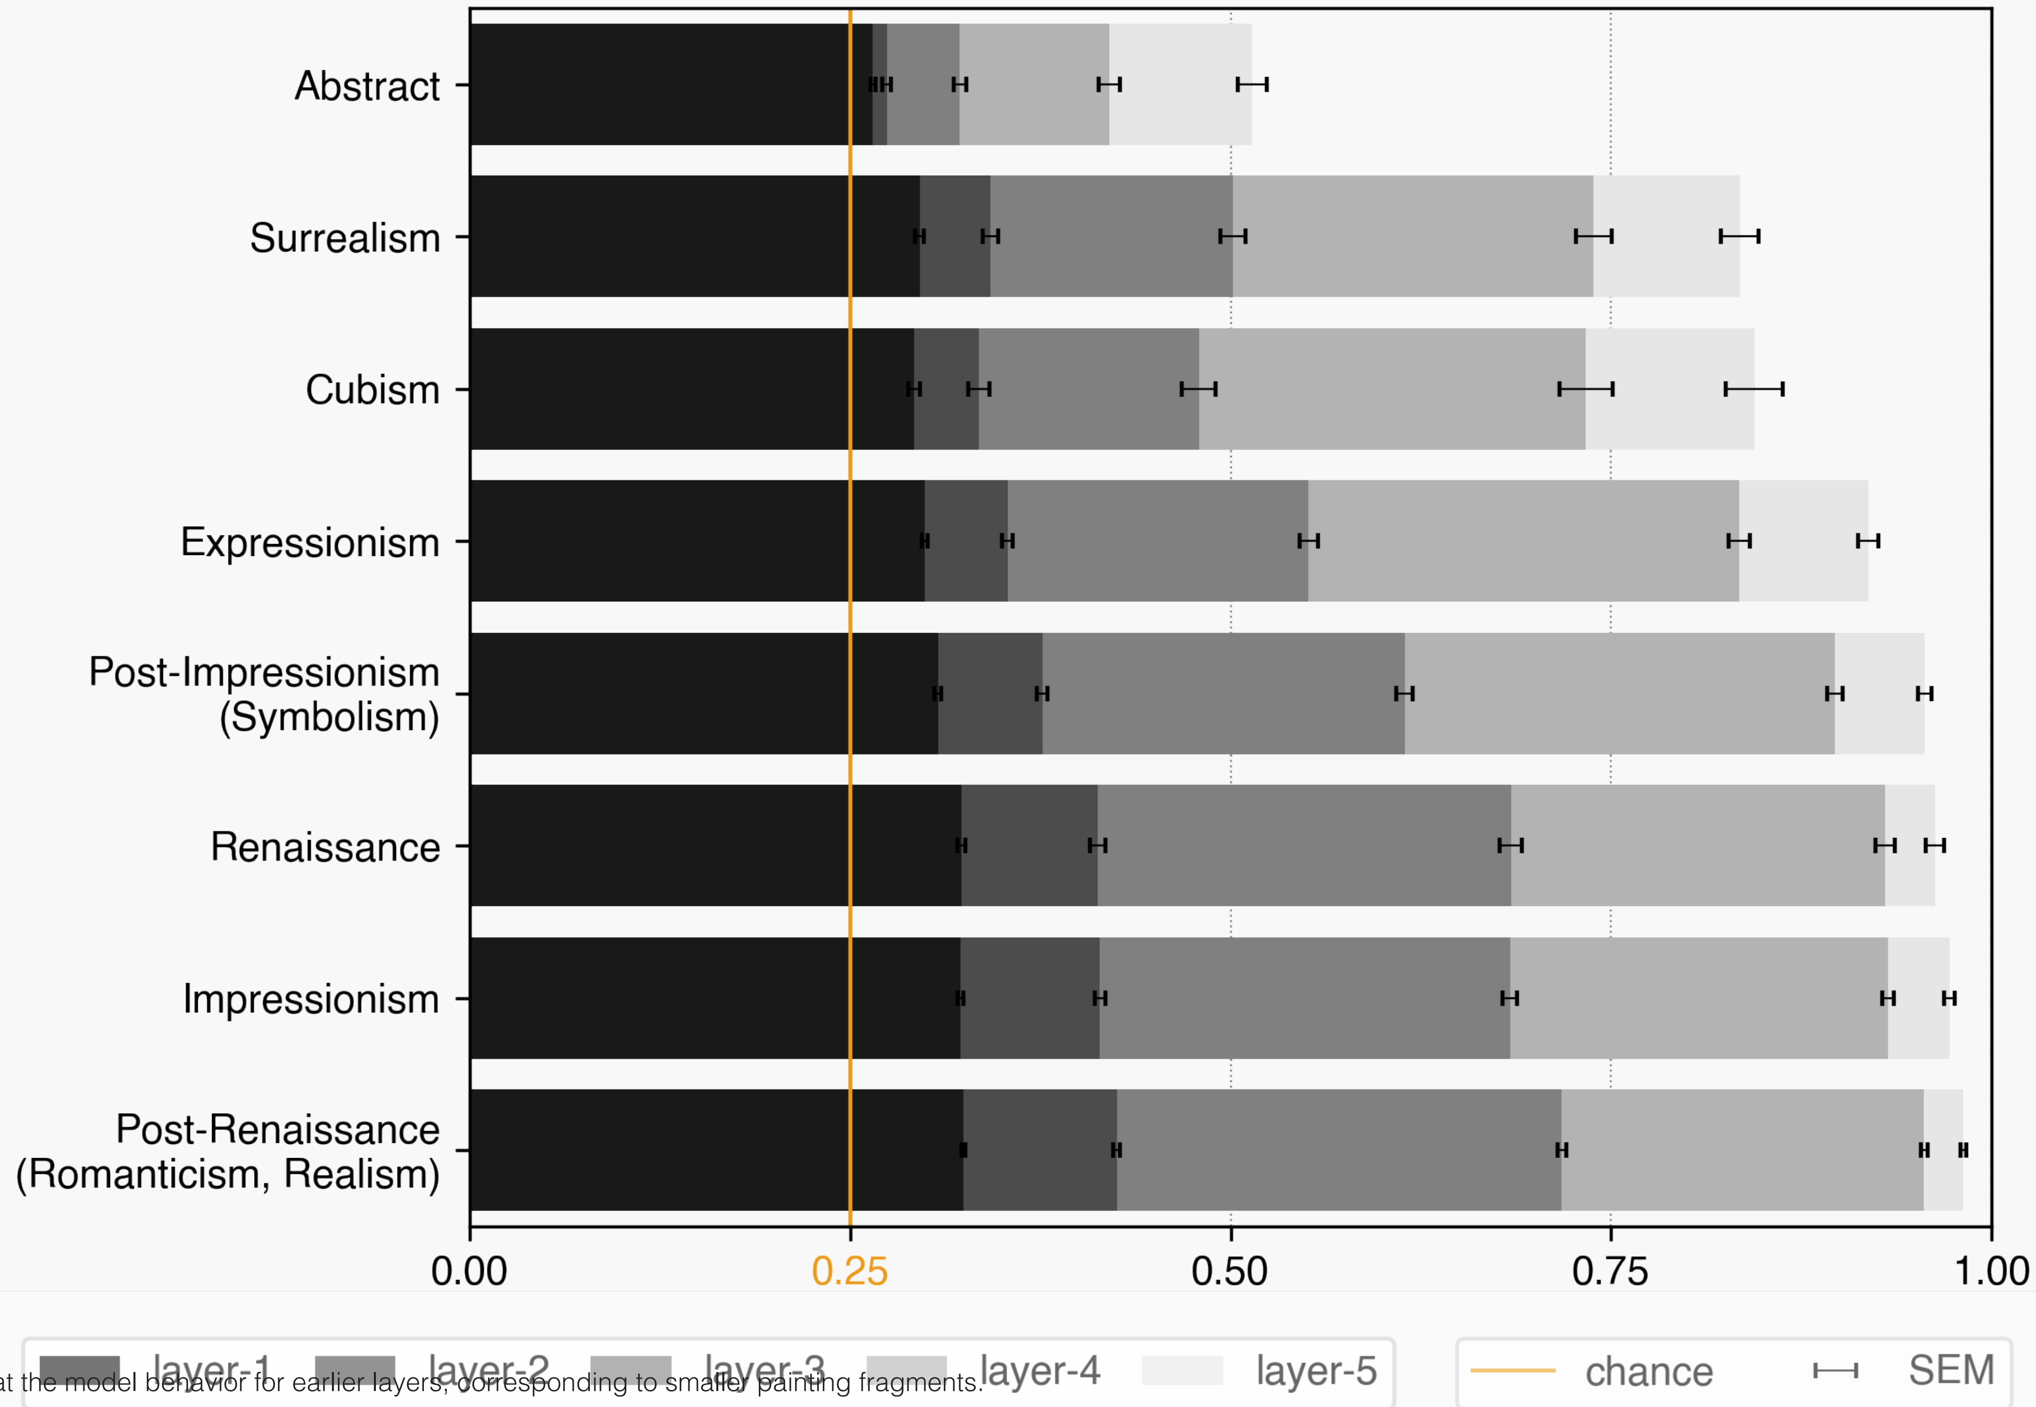

Speaker notes

- These results were for the whole paintings. But, lets have a look at the model behavior for earlier layers, corresponding to smaller painting fragments.
- As we can expect, the model gives performances that shifts toward chance with smaller receptive fields.
- And if we scale the results between chance and performance at layer-5, it seems that abstract styles depend more on deeper layers than more realistic styles.

Model Performance

Accuracy per Layer

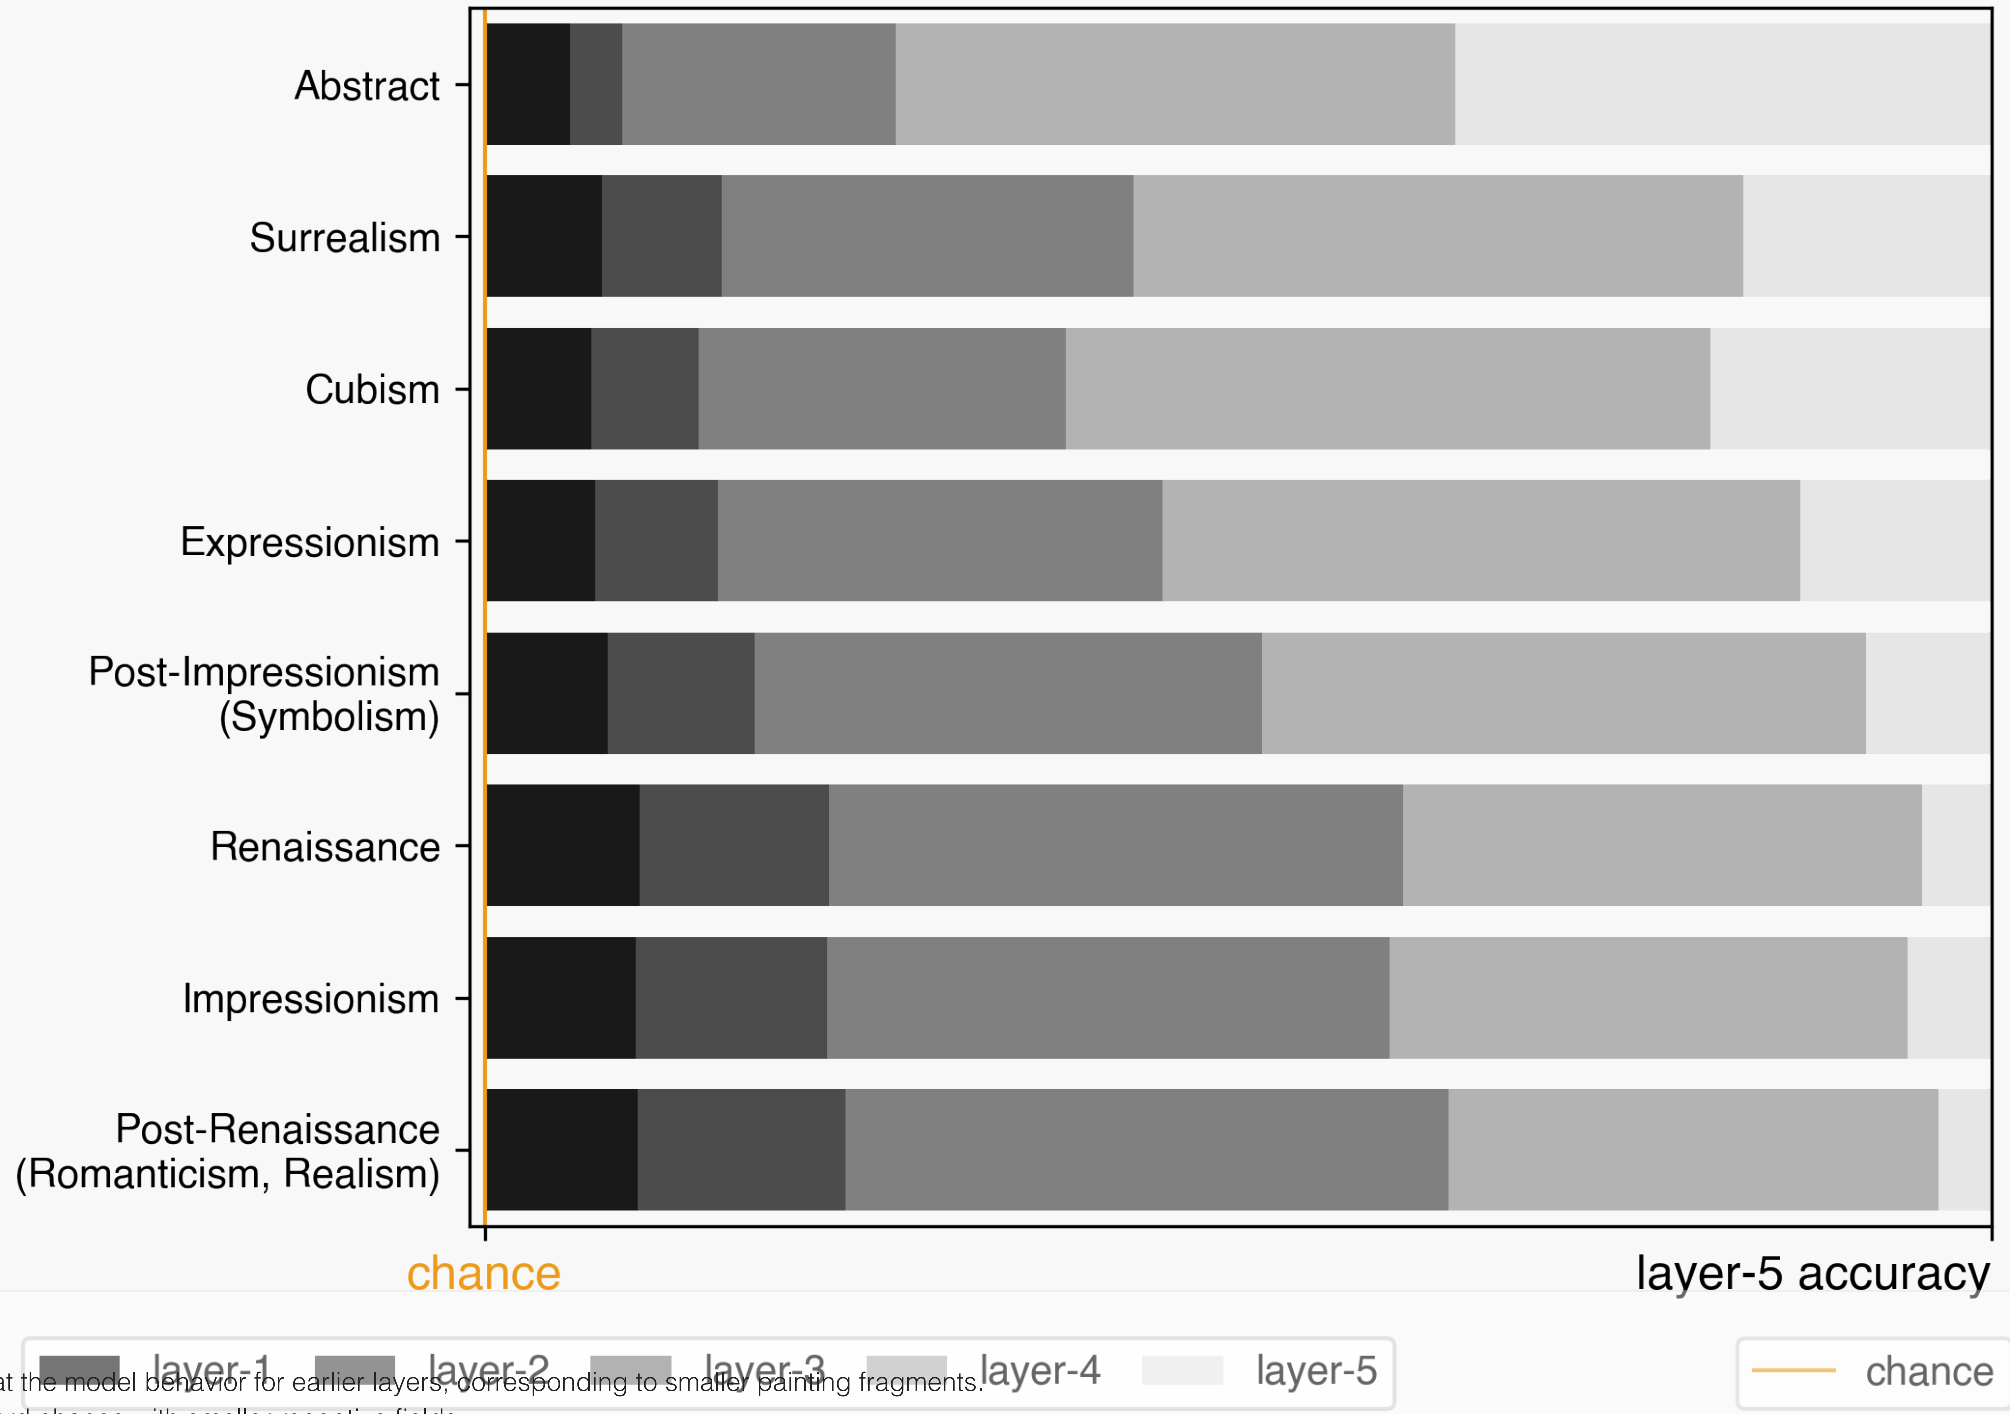

Speaker notes

- These results were for the whole paintings. But, lets have a look at the model behavior for earlier layers, corresponding to smaller painting fragments.
- As we can expect, the model gives performances that shifts toward chance with smaller receptive fields.
- And if we scale the results between chance and performance at layer-5, it seems that abstract styles depend more on deeper layers than more realistic styles.

Model Performance

Back Propagation

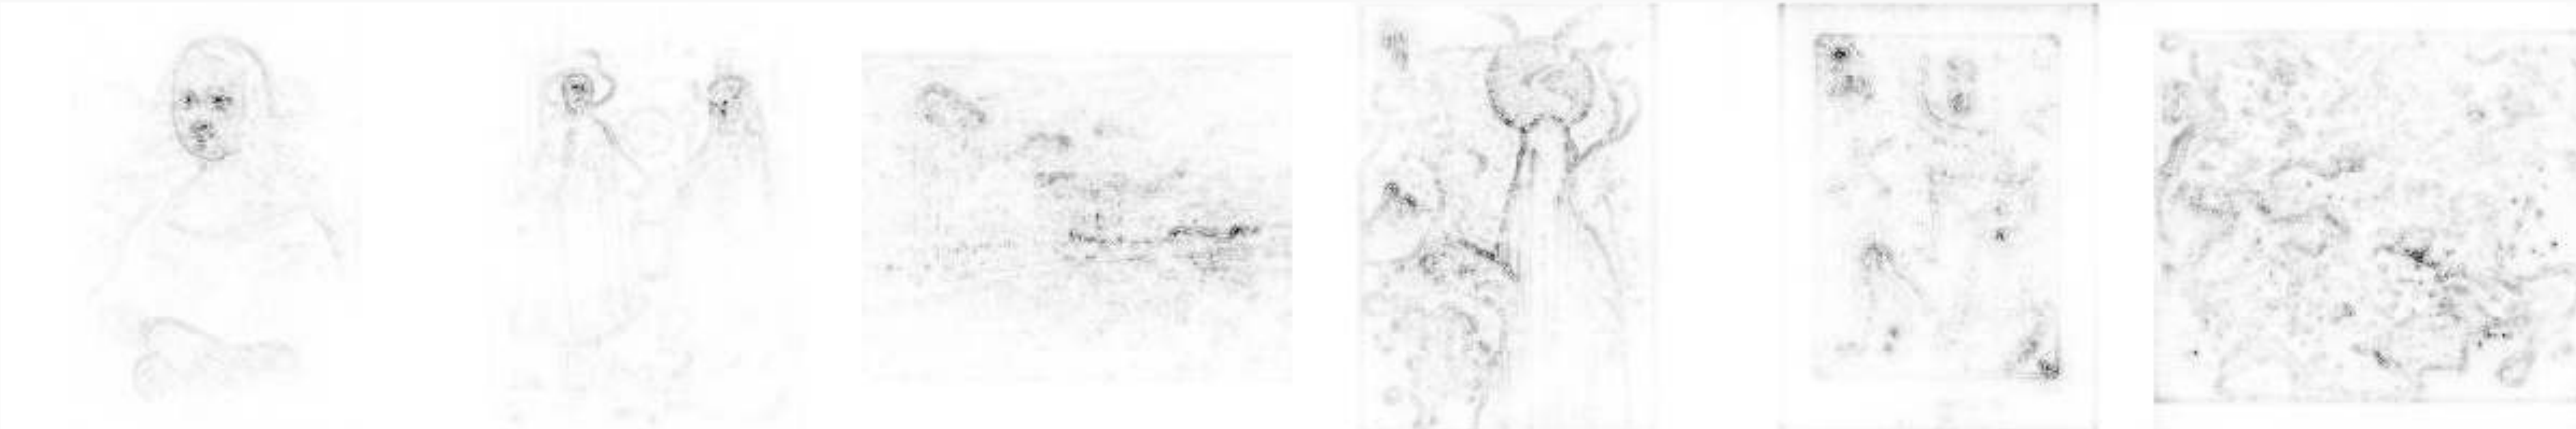

Springenberg, Jost Tobias, Alexey Dosovitskiy, Thomas Brox, and Martin Riedmiller. 2014. “Striving for Simplicity: The All Convolutional Net.” ArXiv:1412.6806 [Cs]

Speaker notes

- Through the visualization of the network's error back propagation, we can observe where the network addresses its attention to answer the question.
- It appears that for realistic paintings, the network can focus on very small regions
- compared to abstract one that need to look at details over the whole area.

Model Performance

Back Propagation

Speaker notes

- Through the visualization of the network's error back propagation, we can observe where the network addresses its attention to answer the question.
- It appears that for realistic paintings, the network can focus on very small regions
- compared to abstract one that need to look at details over the whole area.

Model Performance

Back Propagation

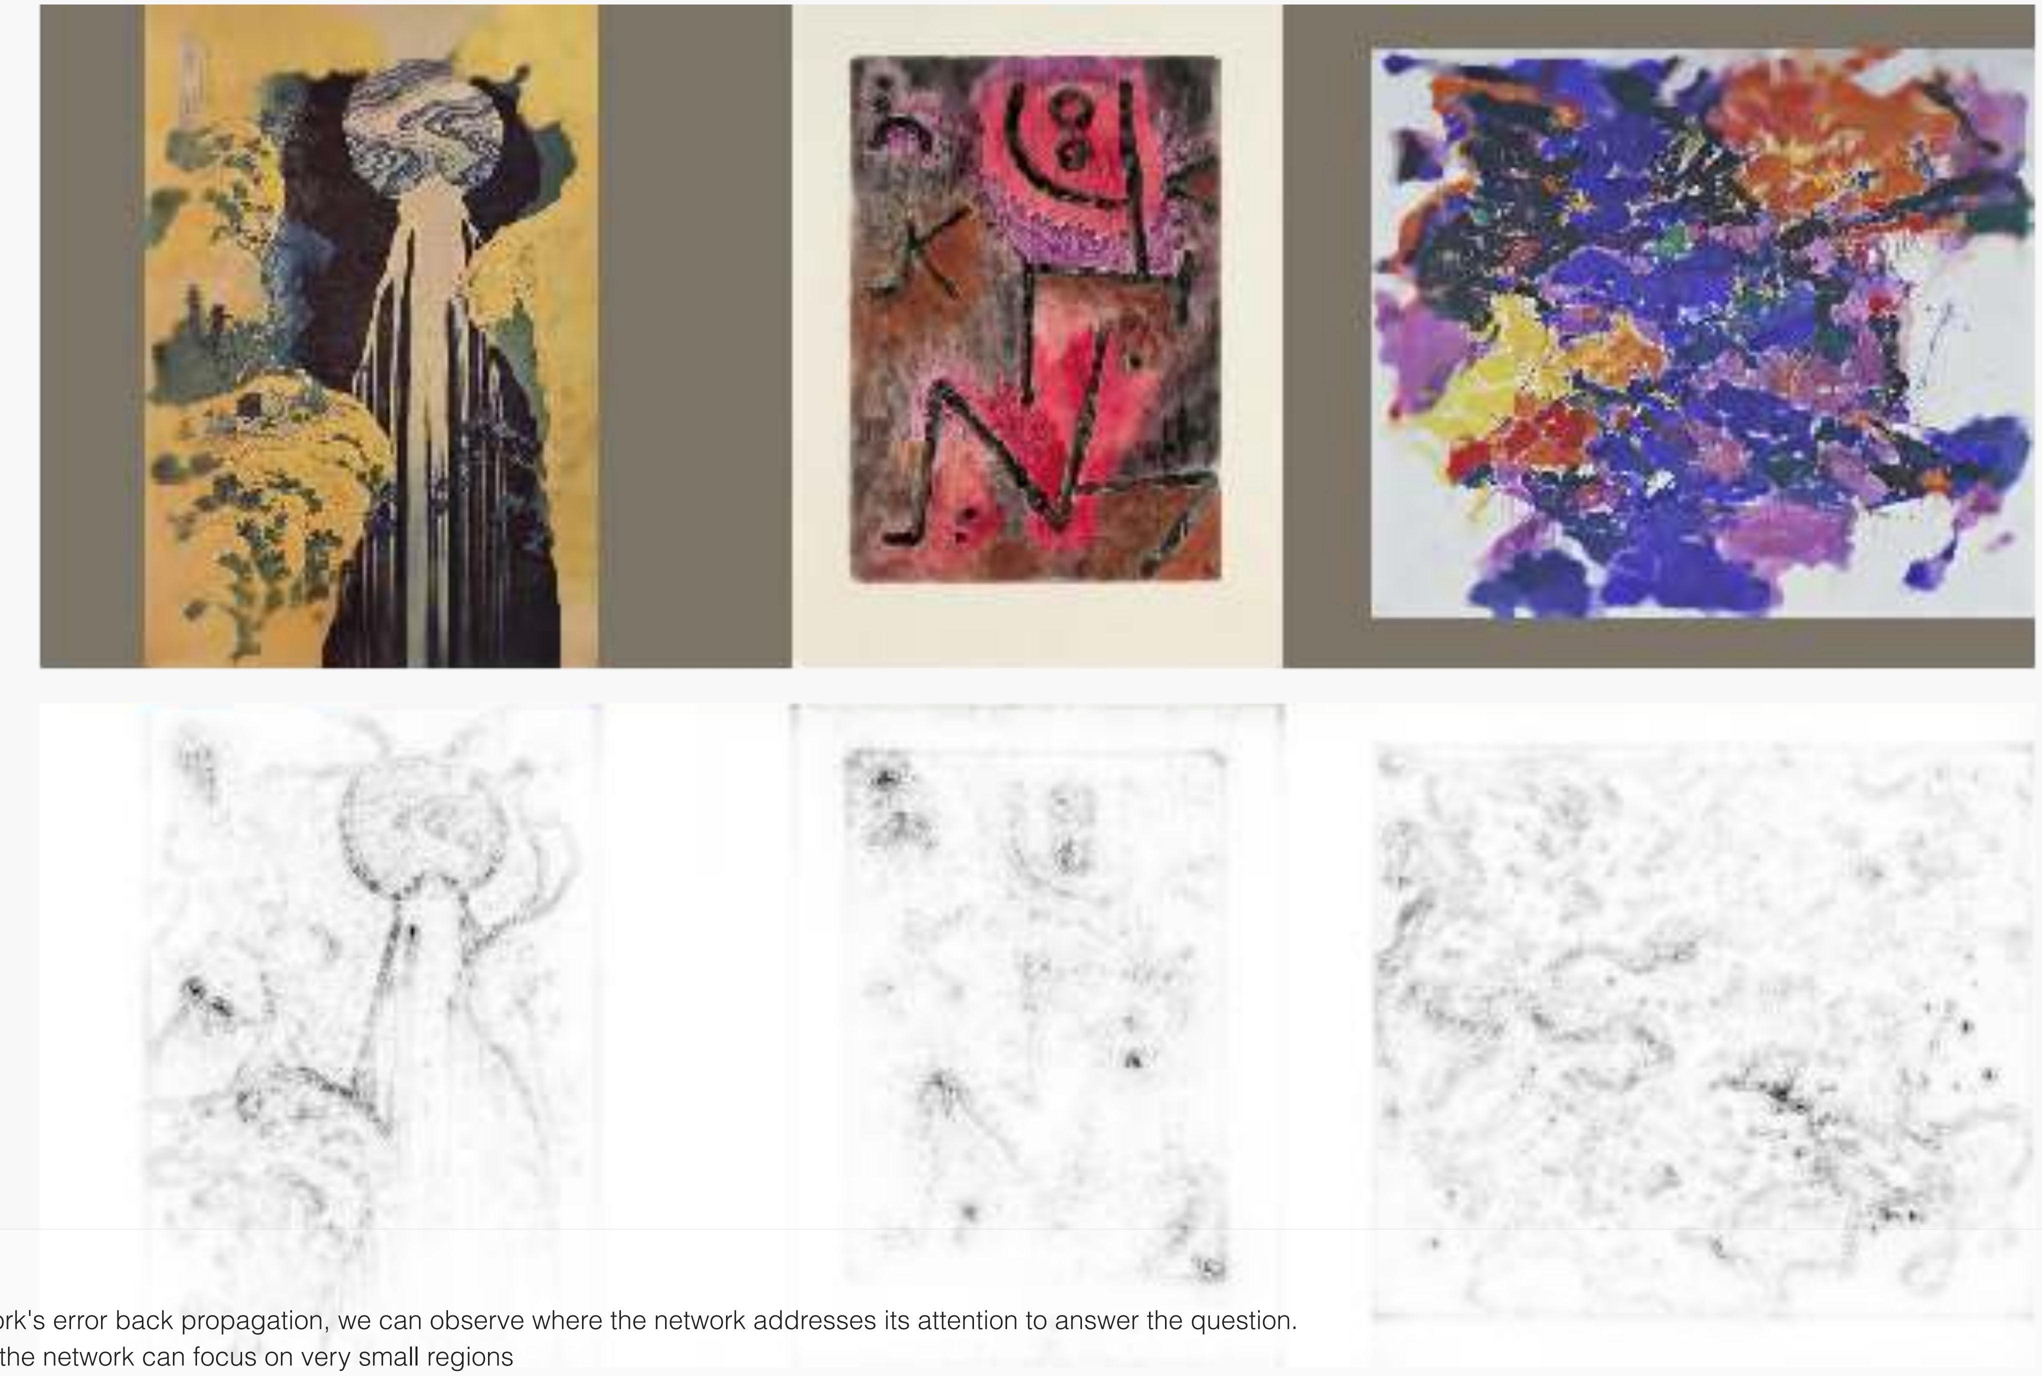

Speaker notes

- Through the visualization of the network's error back propagation, we can observe where the network addresses its attention to answer the question.
- It appears that for realistic paintings, the network can focus on very small regions
- compared to abstract one that need to look at details over the whole area.

Model Performance

Predictions per receptive field units

Speaker notes

To illustrate this phenomenon more precisely, here are the orientation predictions per receptive field units. The size of the 4 wedges in each circles indicates the prediction in that direction, and they naturally sum to 1. From a quick look at it, we can notice two things.

Model Performance

Predictions per receptive field units

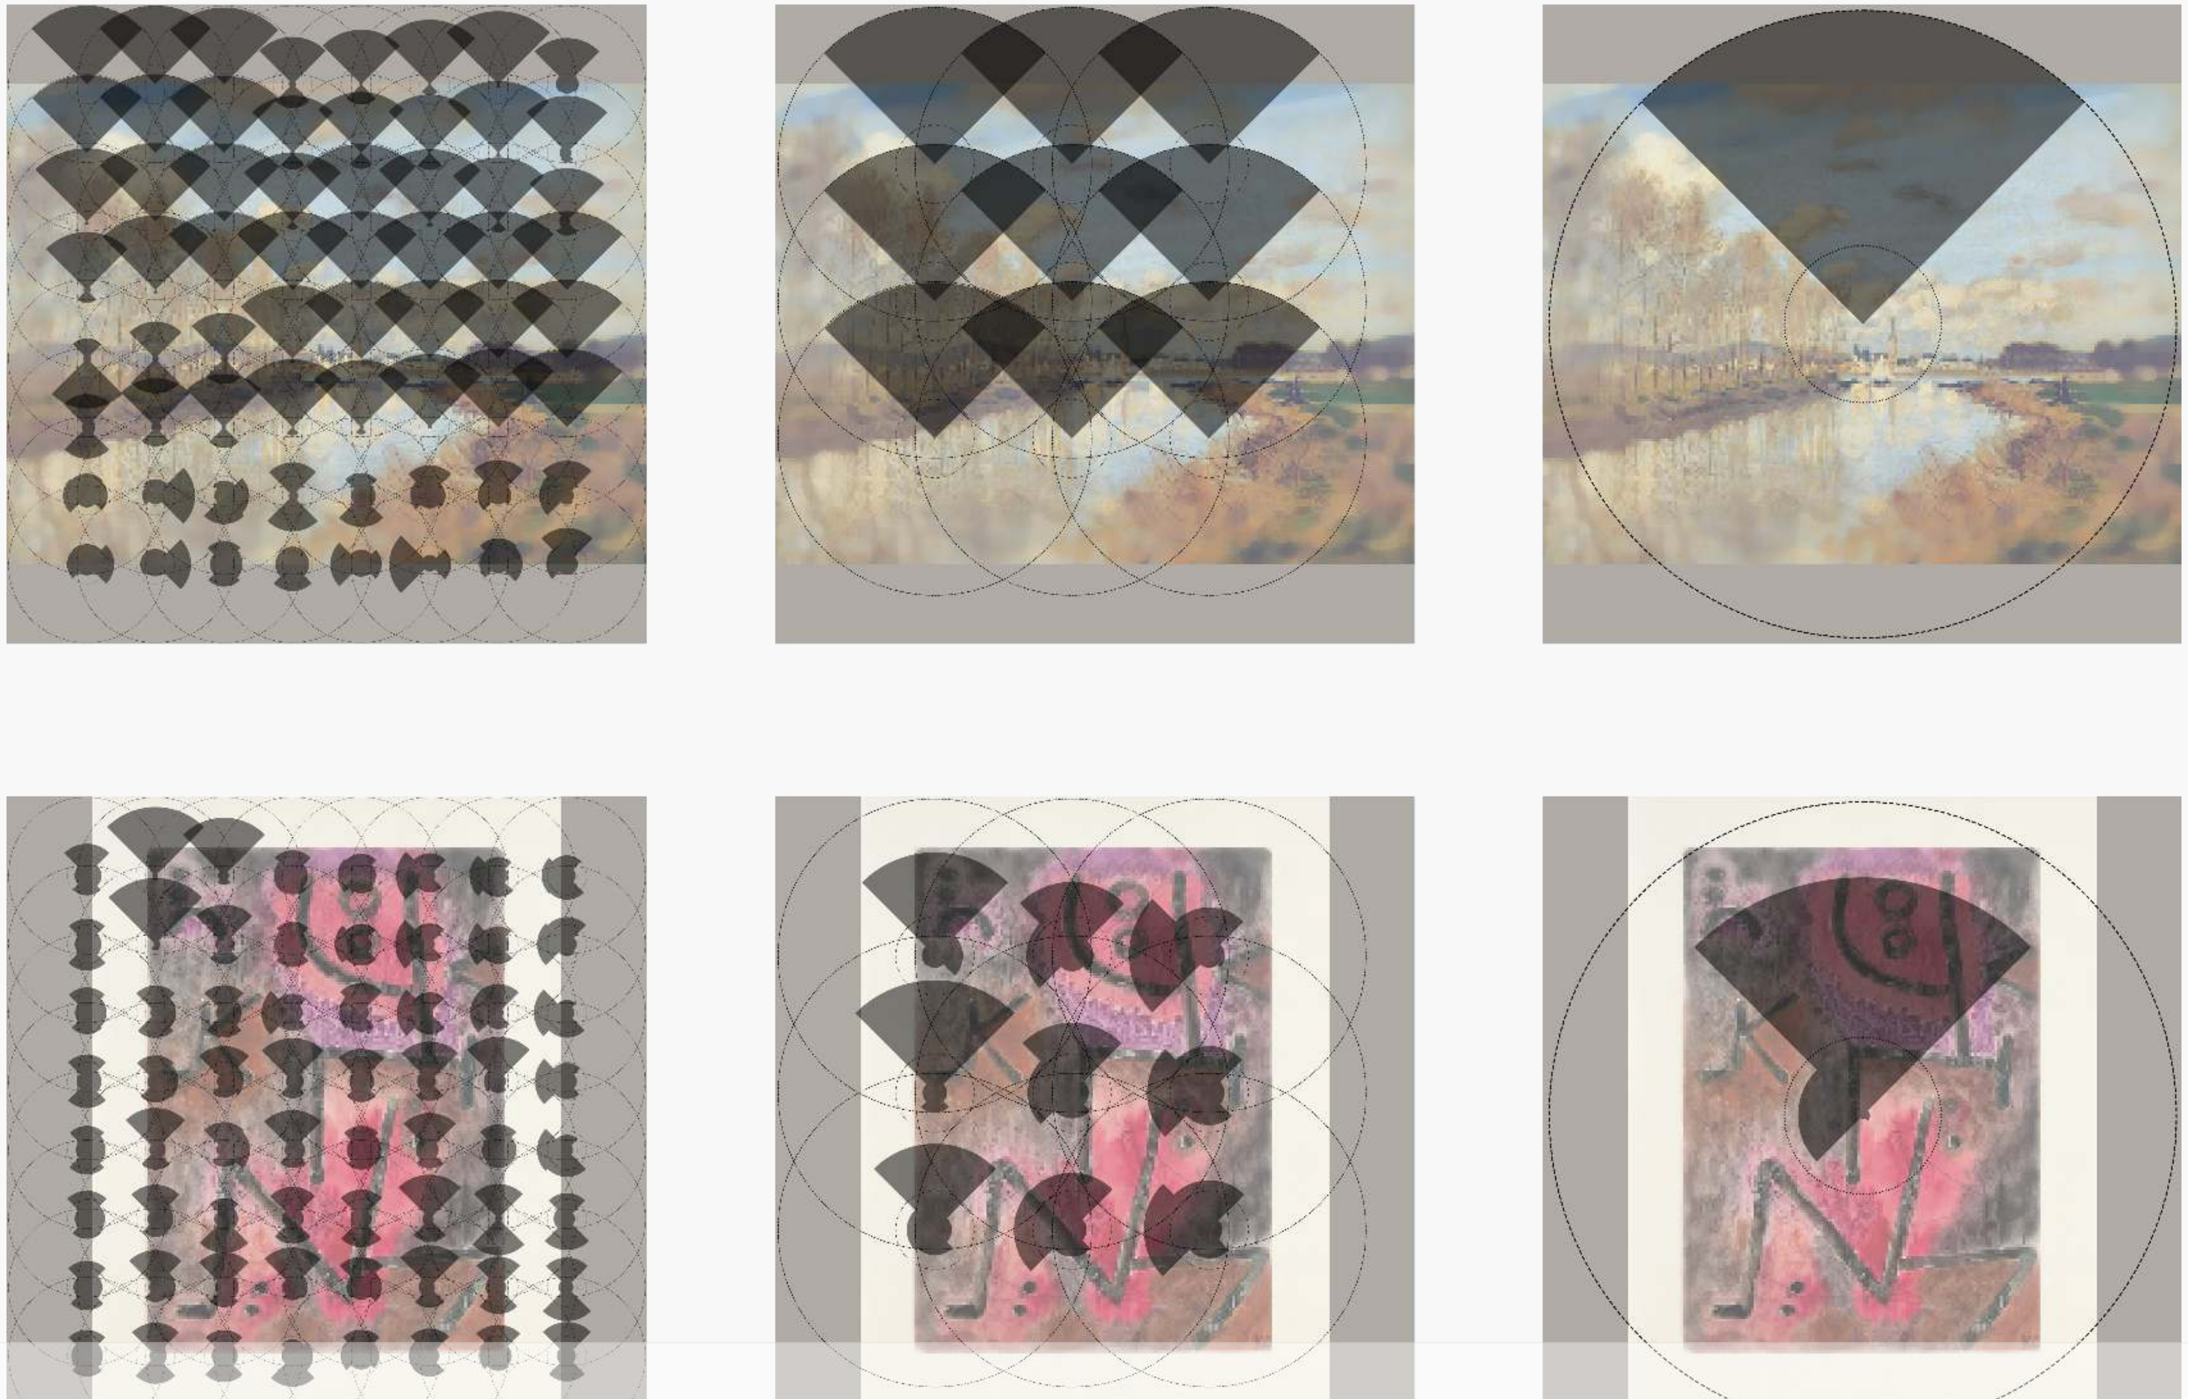

Speaker notes

- First, realistic paintings are more spatially redundant than abstract paintings. They offer more orientation cues all over the painting and also to smaller scales.
- Secondly, for abstract paintings, it seems that there is no simple local integration of cues. A fragment at a bigger scale is not an averaging of smaller related fragments.

Model Performance

Hierarchical Redundancy

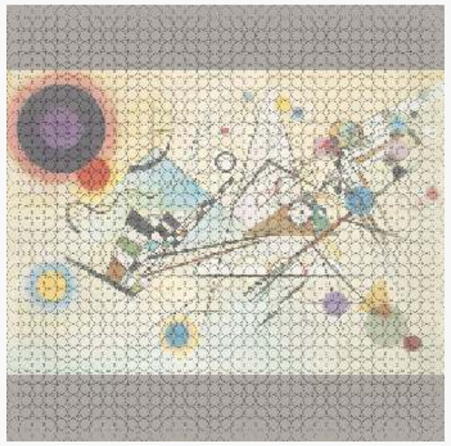

Layer 1 → 2

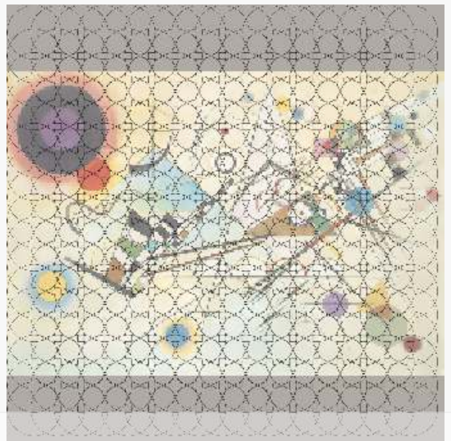

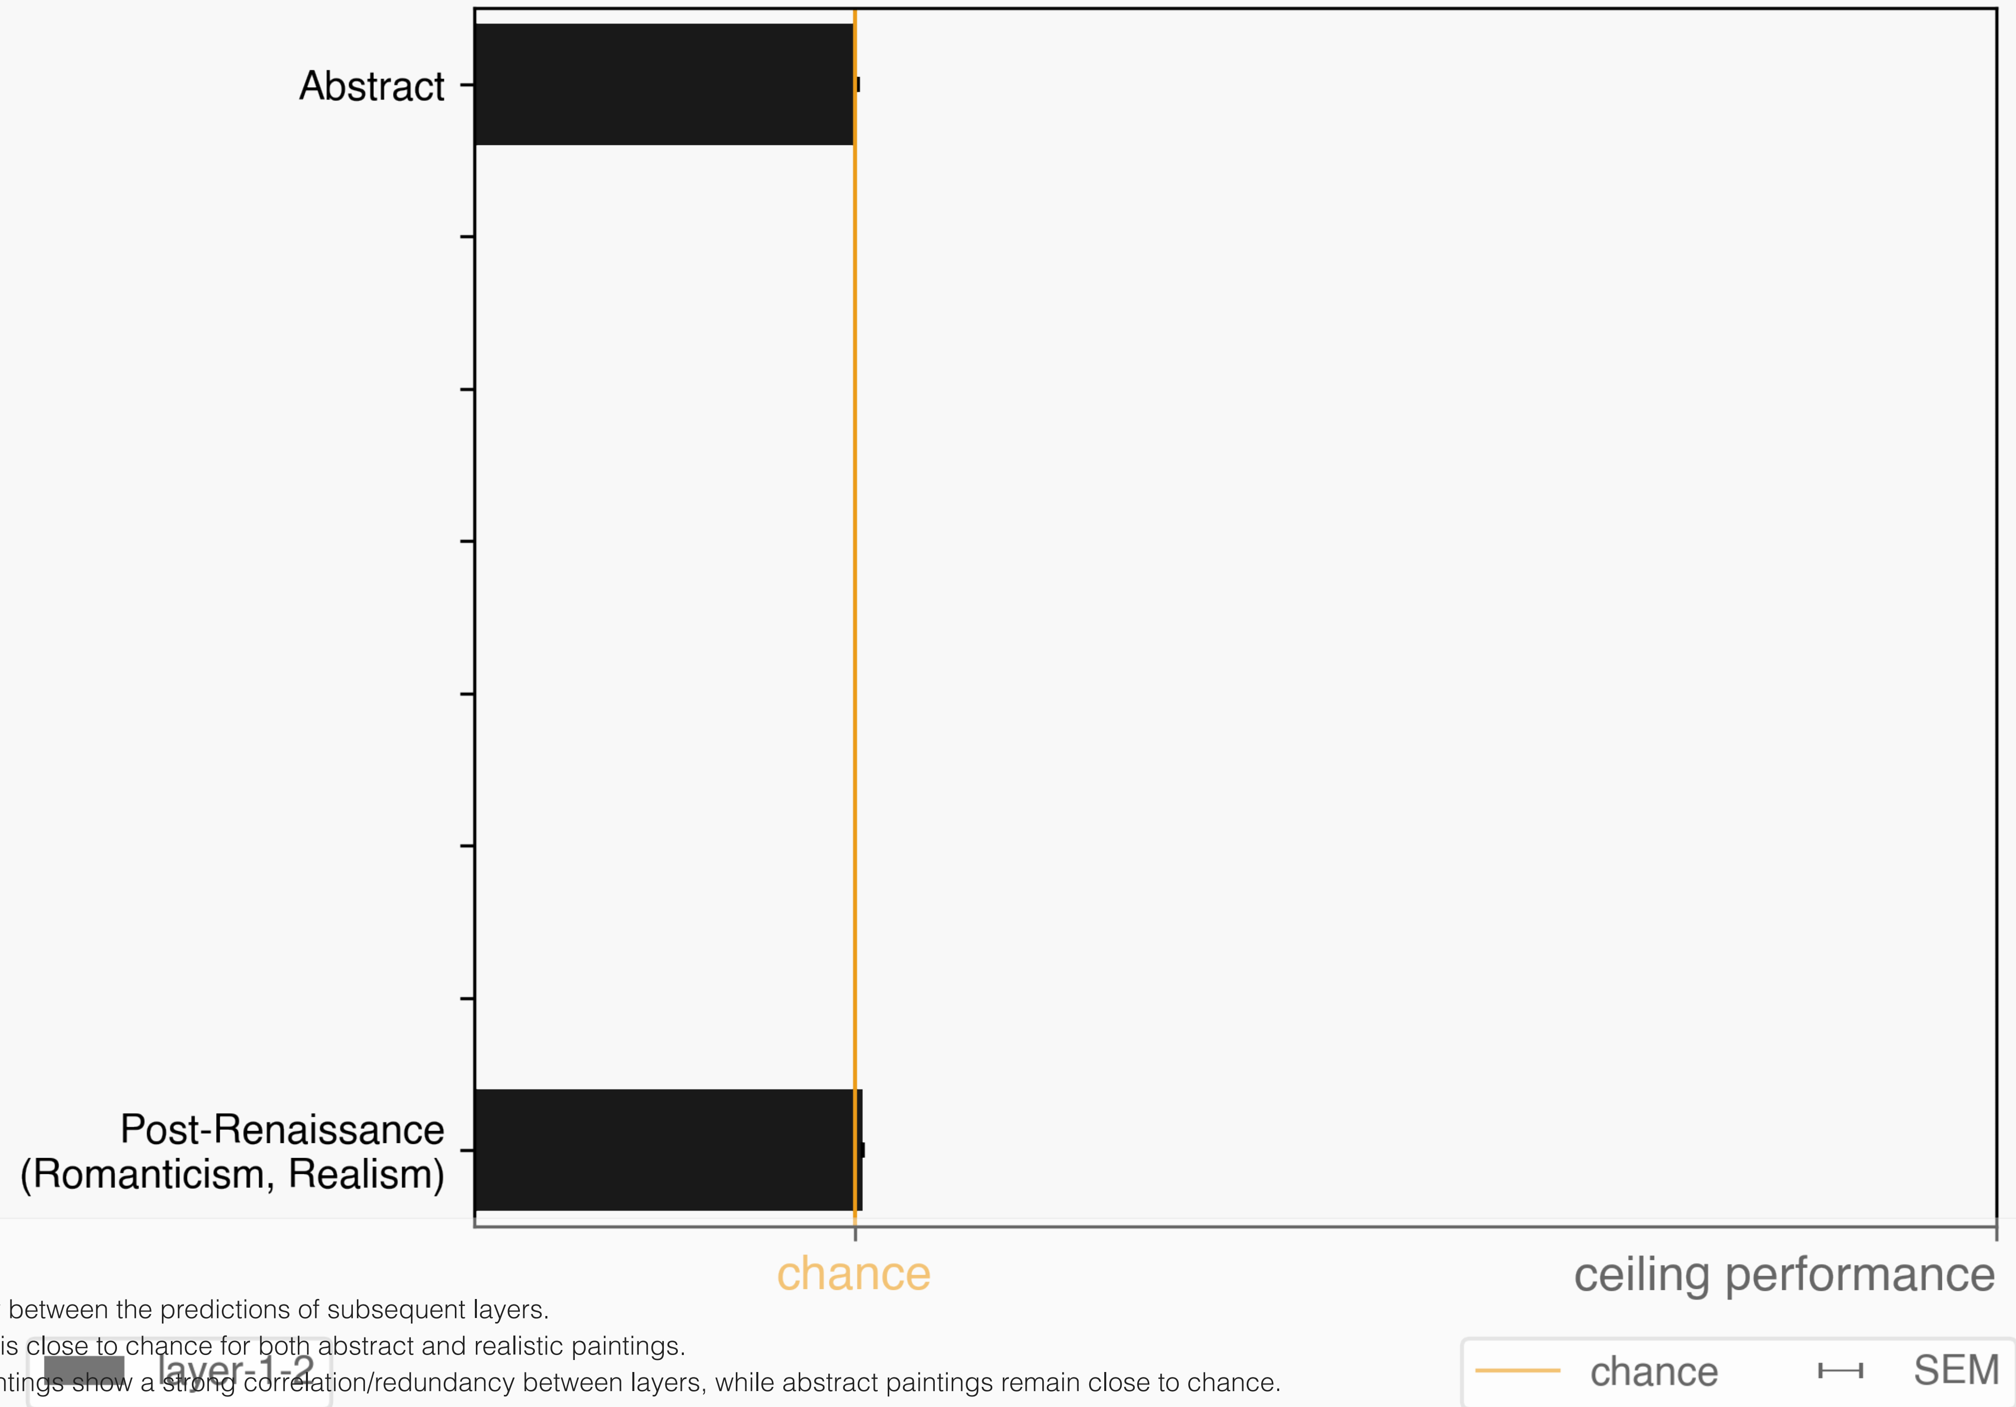

Speaker notes

- - To quantify these observations, the best measure is the cross-entropy between the predictions of subsequent layers.
  - For example, guessing predictions of layer 2 from the ones of layer 1 is close to chance for both abstract and realistic paintings.
- But as we move toward predictions of layer 5 from layer 4, realistic paintings show a strong correlation/redundancy between layers, while abstract paintings remain close to chance.
- So abstract art suffers from a higher local variability of compositional effects and needs a larger integration of orientation cues that can only make sense globally. Deeper layers must represent some emergent global property that is not necessarily carried by previous layers.

Model Performance

Hierarchical Redundancy

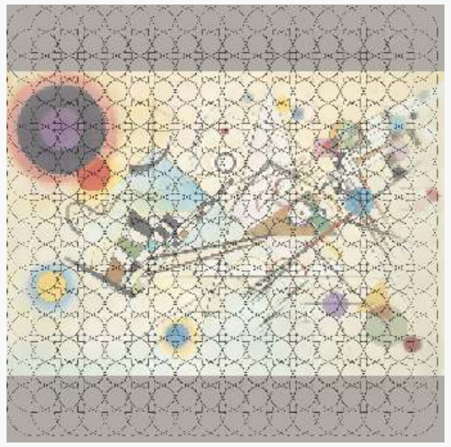

Layer 2 → 3

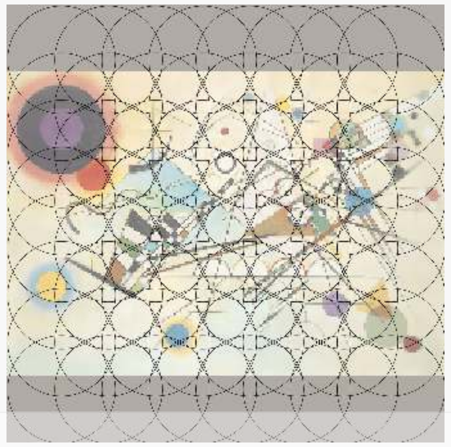

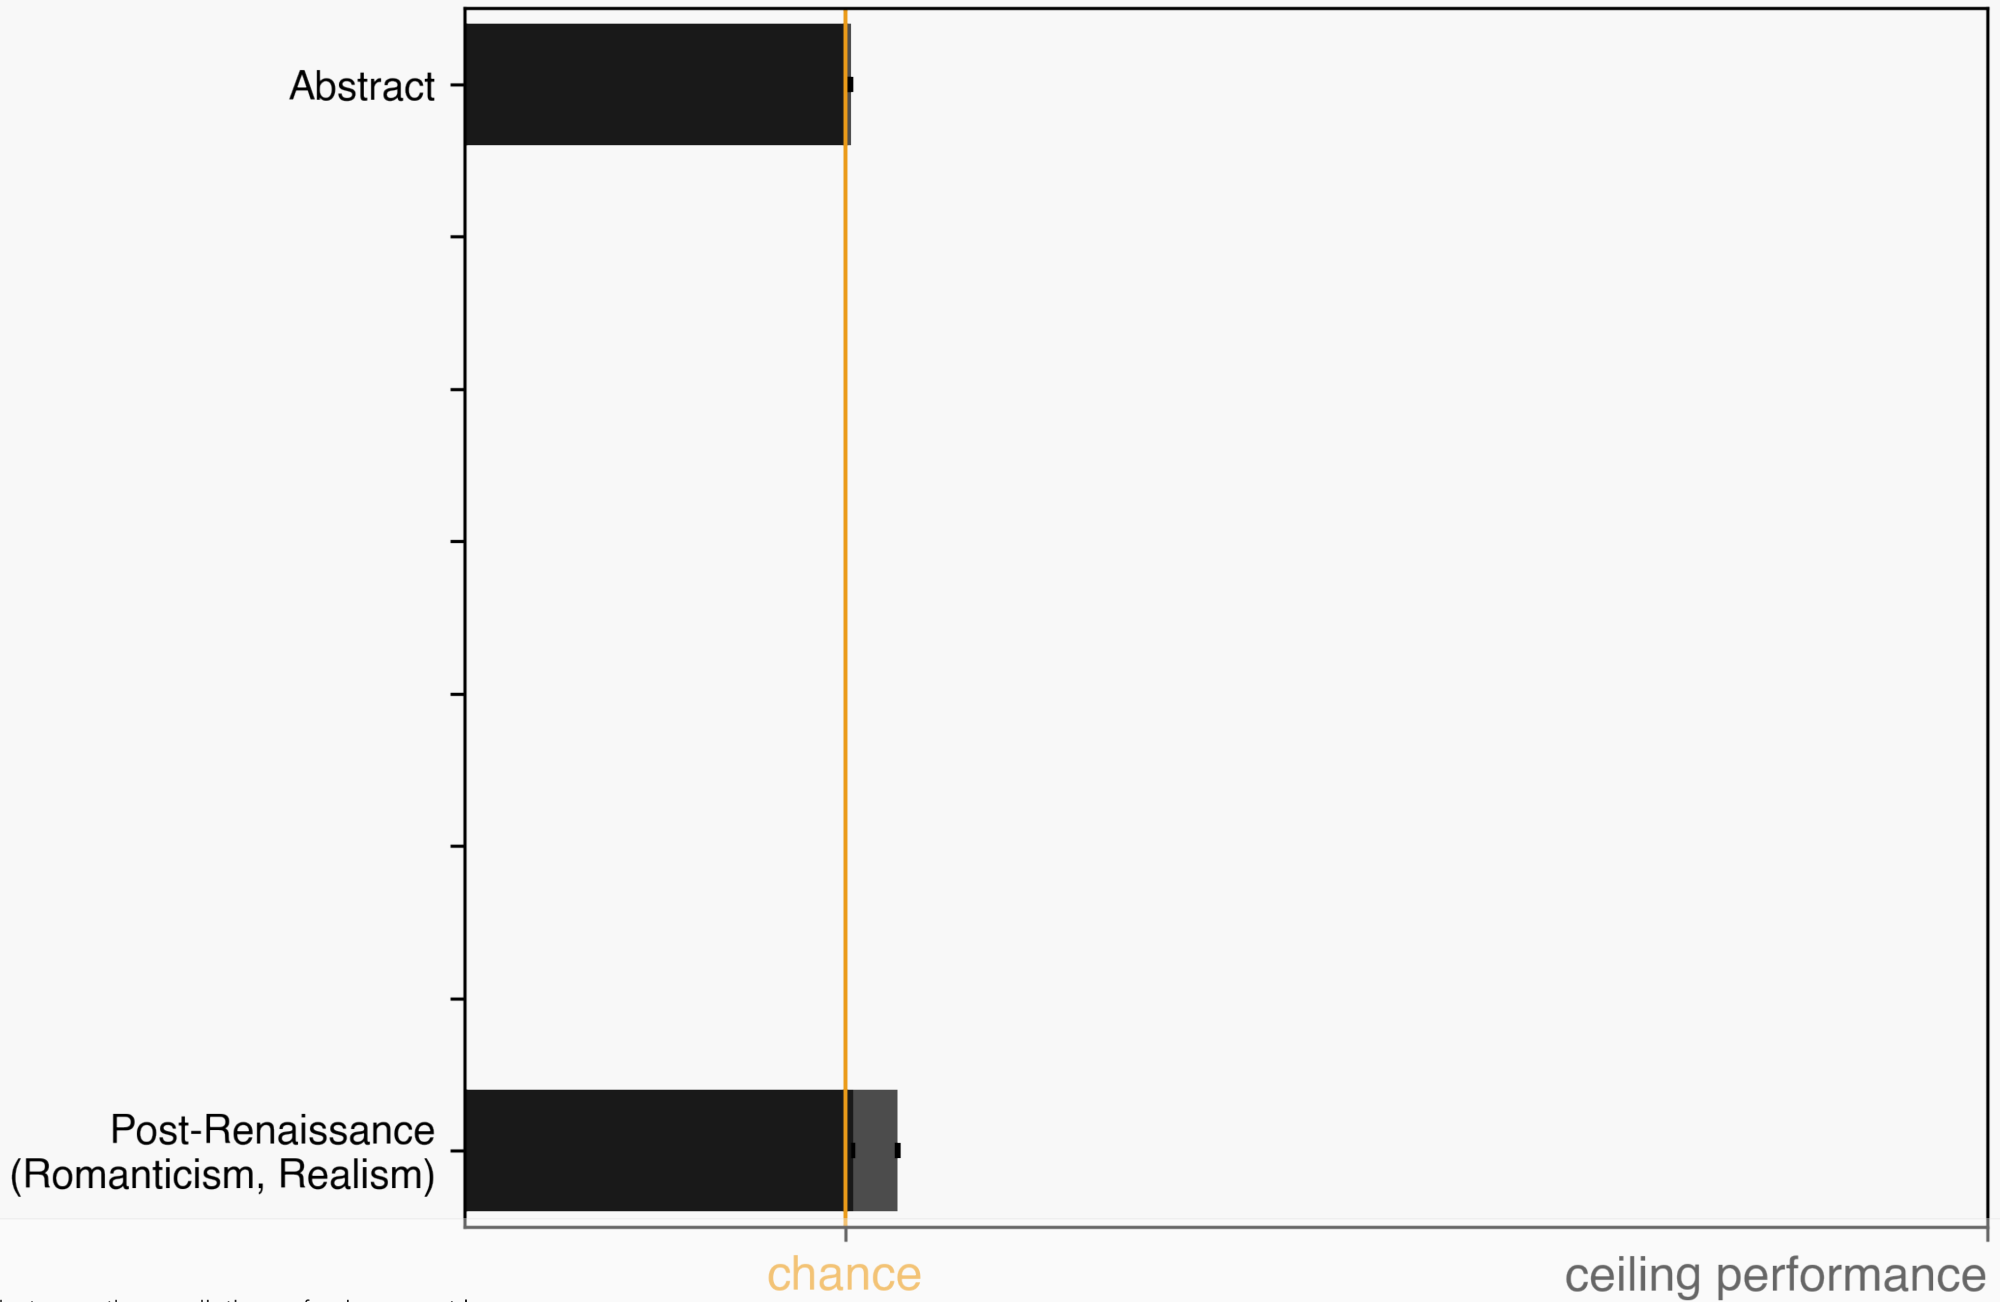

Speaker notes

- - To quantify these observations, the best measure is the cross-entropy between the predictions of subsequent layers.
  - For example, guessing predictions of layer 2 from the ones of layer 1 is close to chance for both abstract and realistic paintings.
- But as we move toward predictions of layer 5 from layer 4, realistic paintings show a strong correlation/redundancy between layers, while abstract paintings remain close to chance.
- So abstract art suffers from a higher local variability of compositional effects and needs a larger integration of orientation cues that can only make sense globally. Deeper layers must represent some emergent global property that is not necessarily carried by previous layers.

Model Performance

Hierarchical Redundancy

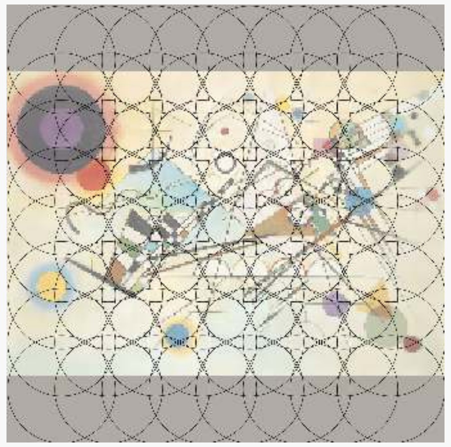

Layer 3 → 4

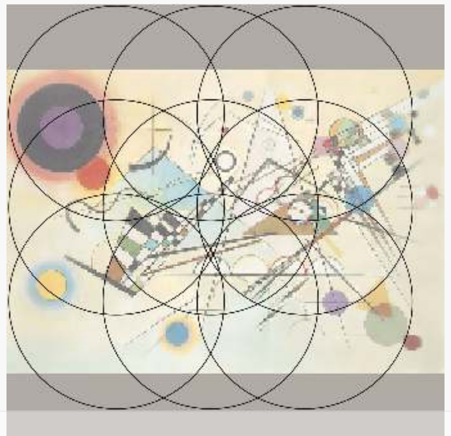

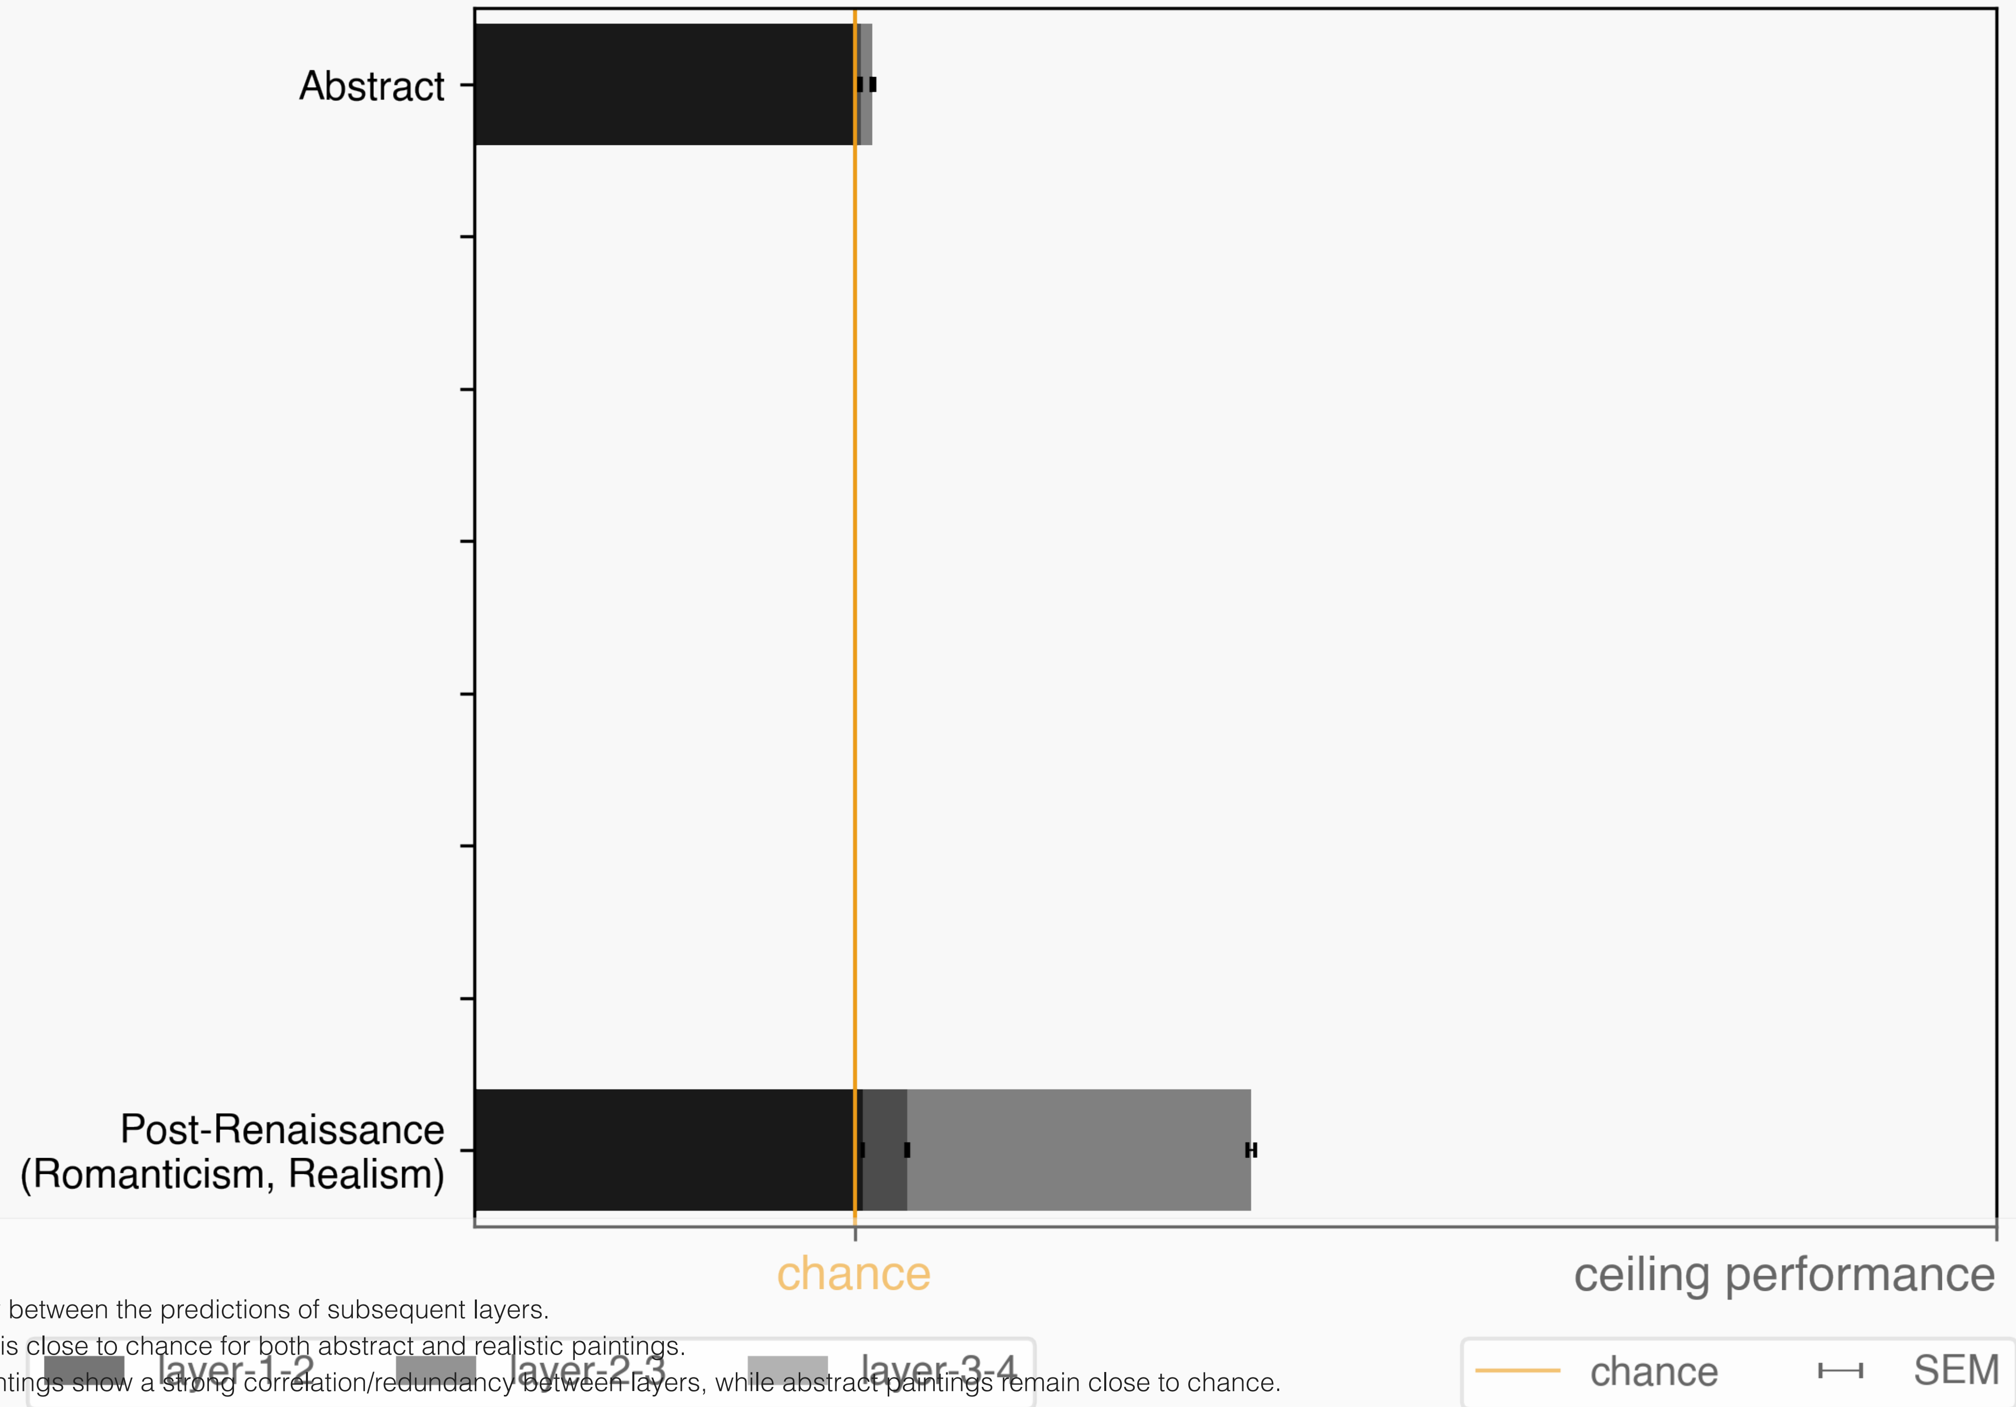

Speaker notes

- - To quantify these observations, the best measure is the cross-entropy between the predictions of subsequent layers.
  - For example, guessing predictions of layer 2 from the ones of layer 1 is close to chance for both abstract and realistic paintings.
- But as we move toward predictions of layer 5 from layer 4, realistic paintings show a strong correlation/redundancy between layers, while abstract paintings remain close to chance.
- So abstract art suffers from a higher local variability of compositional effects and needs a larger integration of orientation cues that can only make sense globally. Deeper layers must represent some emergent global property that is not necessarily carried by previous layers.

Model Performance

Hierarchical Redundancy

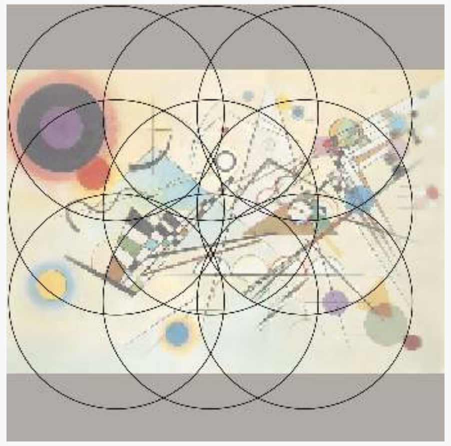

Layer 4 → 5

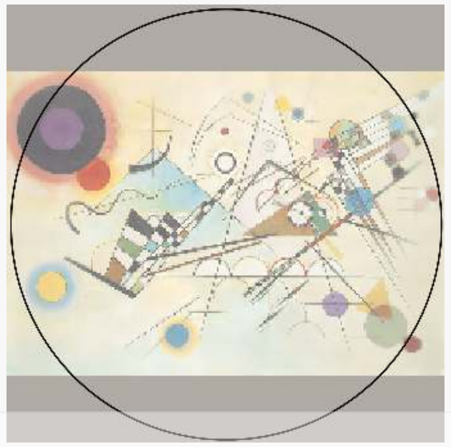

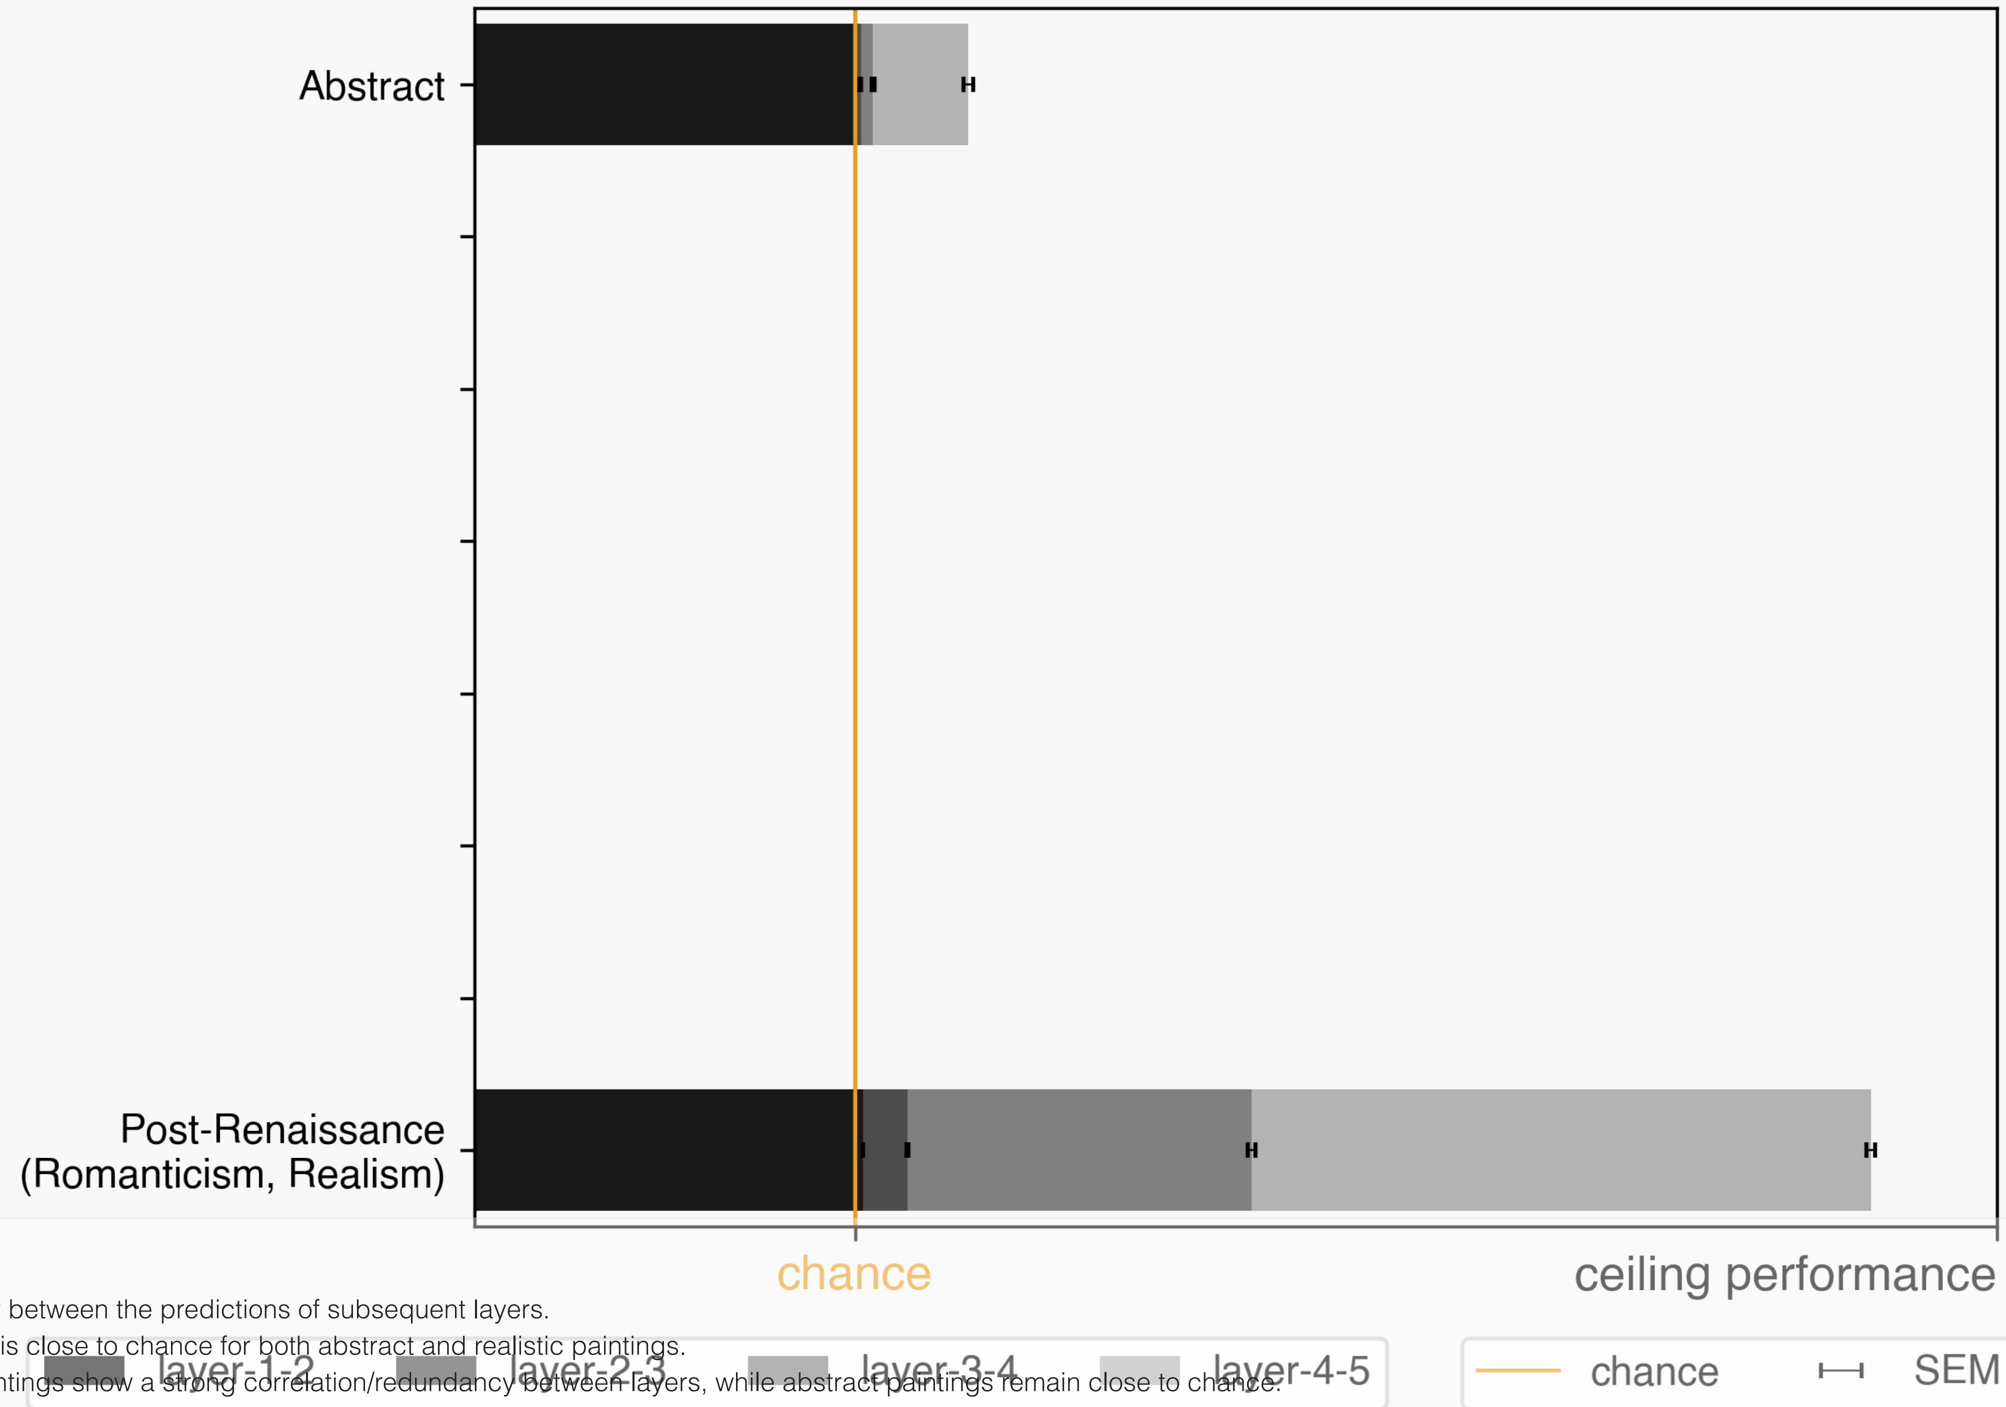

Speaker notes

- - To quantify these observations, the best measure is the cross-entropy between the predictions of subsequent layers.
  - For example, guessing predictions of layer 2 from the ones of layer 1 is close to chance for both abstract and realistic paintings.
- But as we move toward predictions of layer 5 from layer 4, realistic paintings show a strong correlation/redundancy between layers, while abstract paintings remain close to chance.
- So abstract art suffers from a higher local variability of compositional effects and needs a larger integration of orientation cues that can only make sense globally. Deeper layers must represent some emergent global property that is not necessarily carried by previous layers.

Model Performance

Hierarchical Redundancy

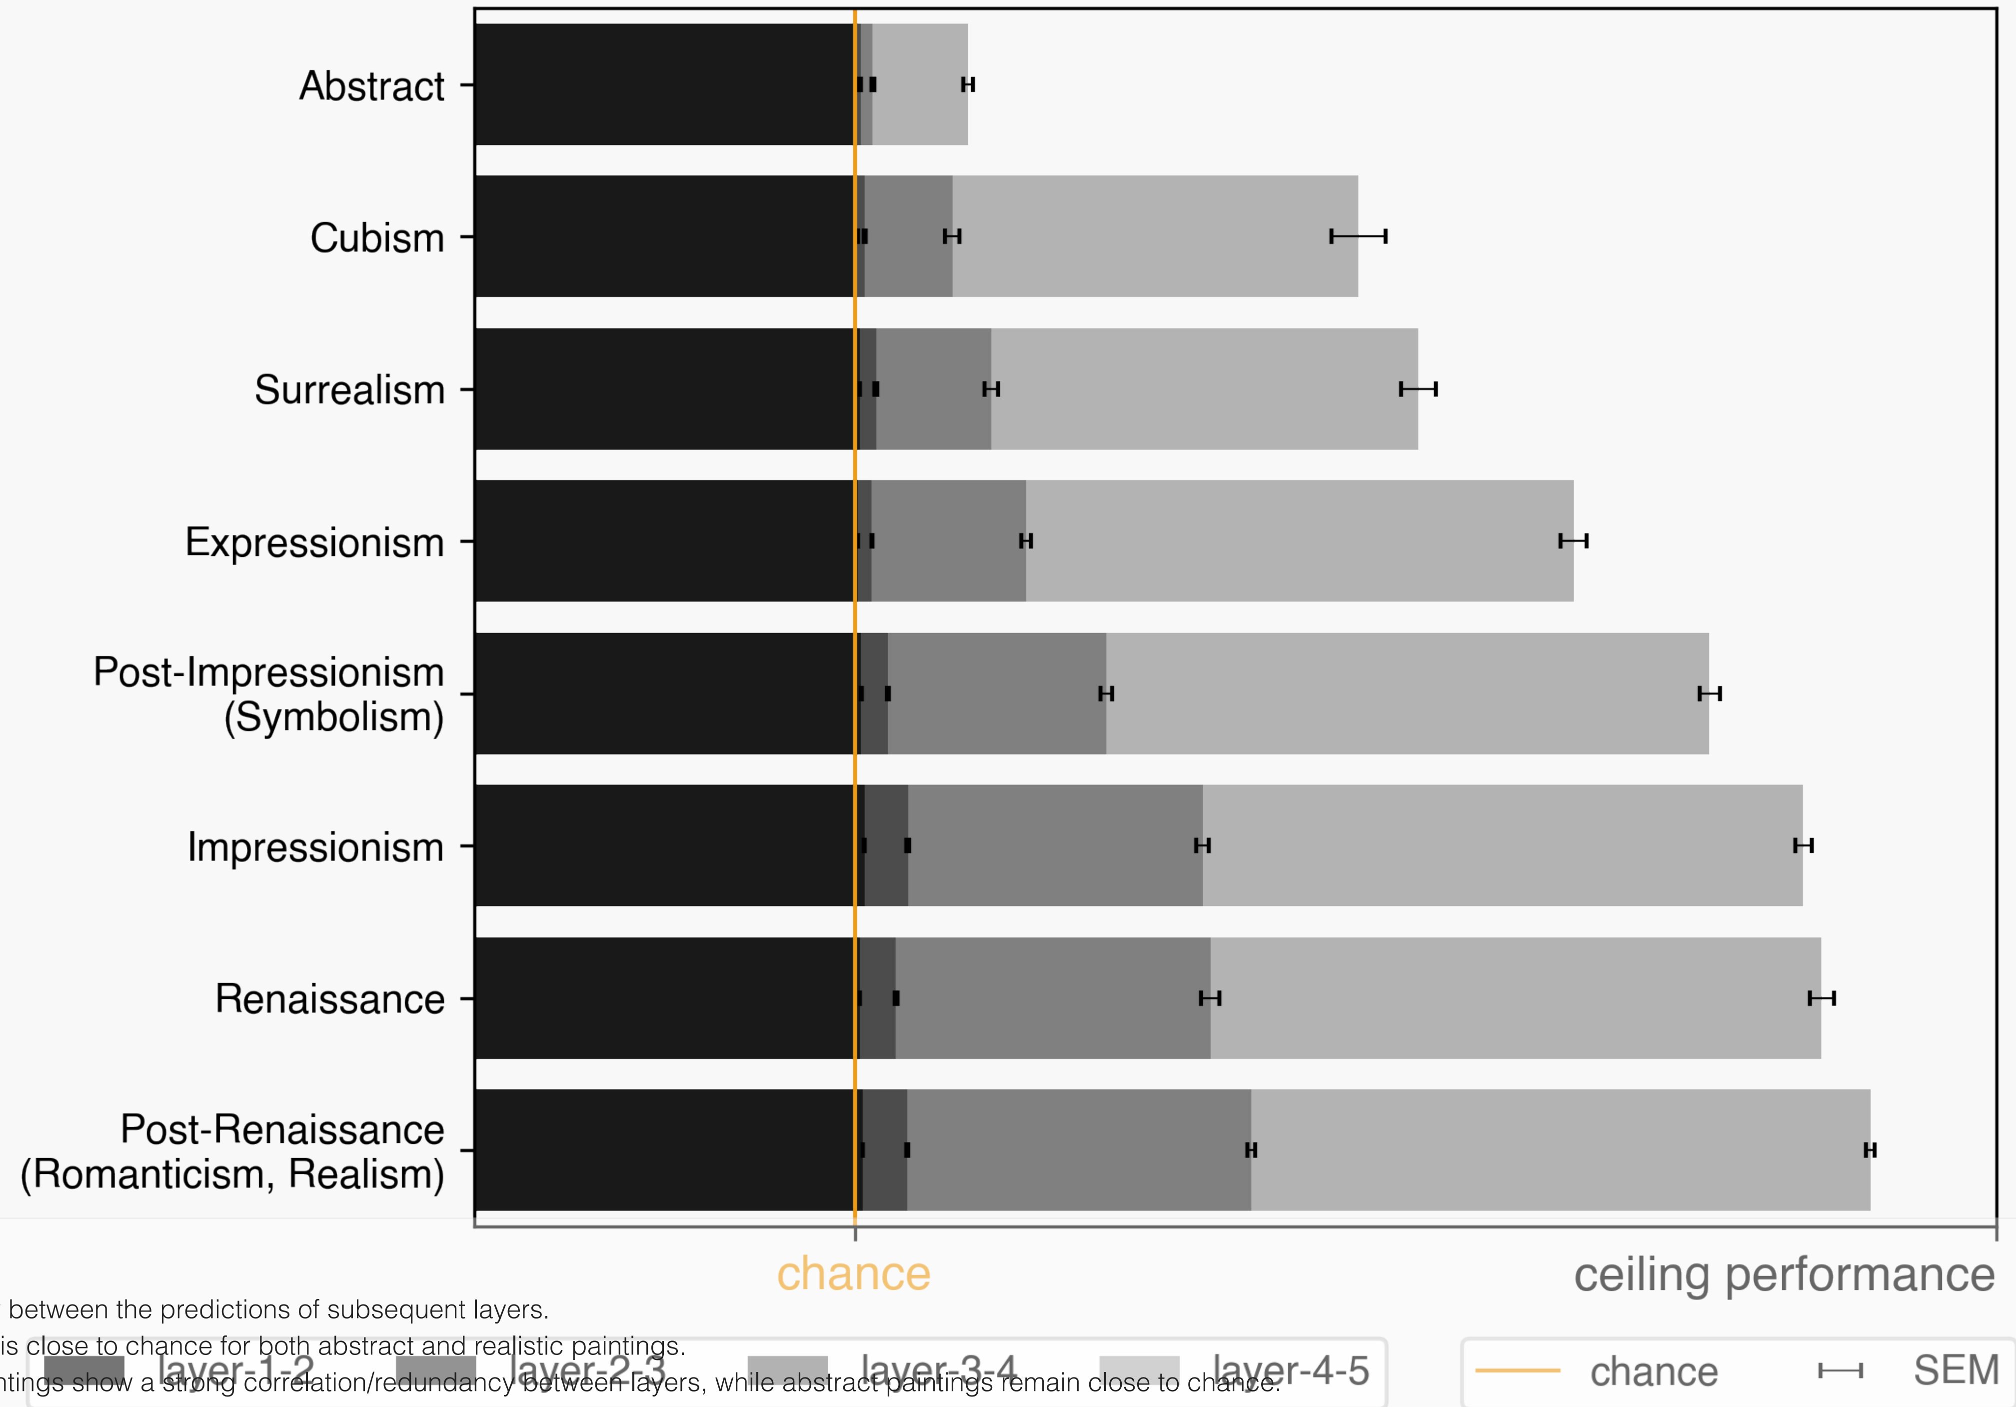

Speaker notes

- - To quantify these observations, the best measure is the cross-entropy between the predictions of subsequent layers.
  - For example, guessing predictions of layer 2 from the ones of layer 1 is close to chance for both abstract and realistic paintings.
- But as we move toward predictions of layer 5 from layer 4, realistic paintings show a strong correlation/redundancy between layers, while abstract paintings remain close to chance.
- So abstract art suffers from a higher local variability of compositional effects and needs a larger integration of orientation cues that can only make sense globally. Deeper layers must represent some emergent global property that is not necessarily carried by previous layers.

Human Performance

Speaker notes

- But how does these observations fit human performance ?

Human Performance

Accuracy per Style

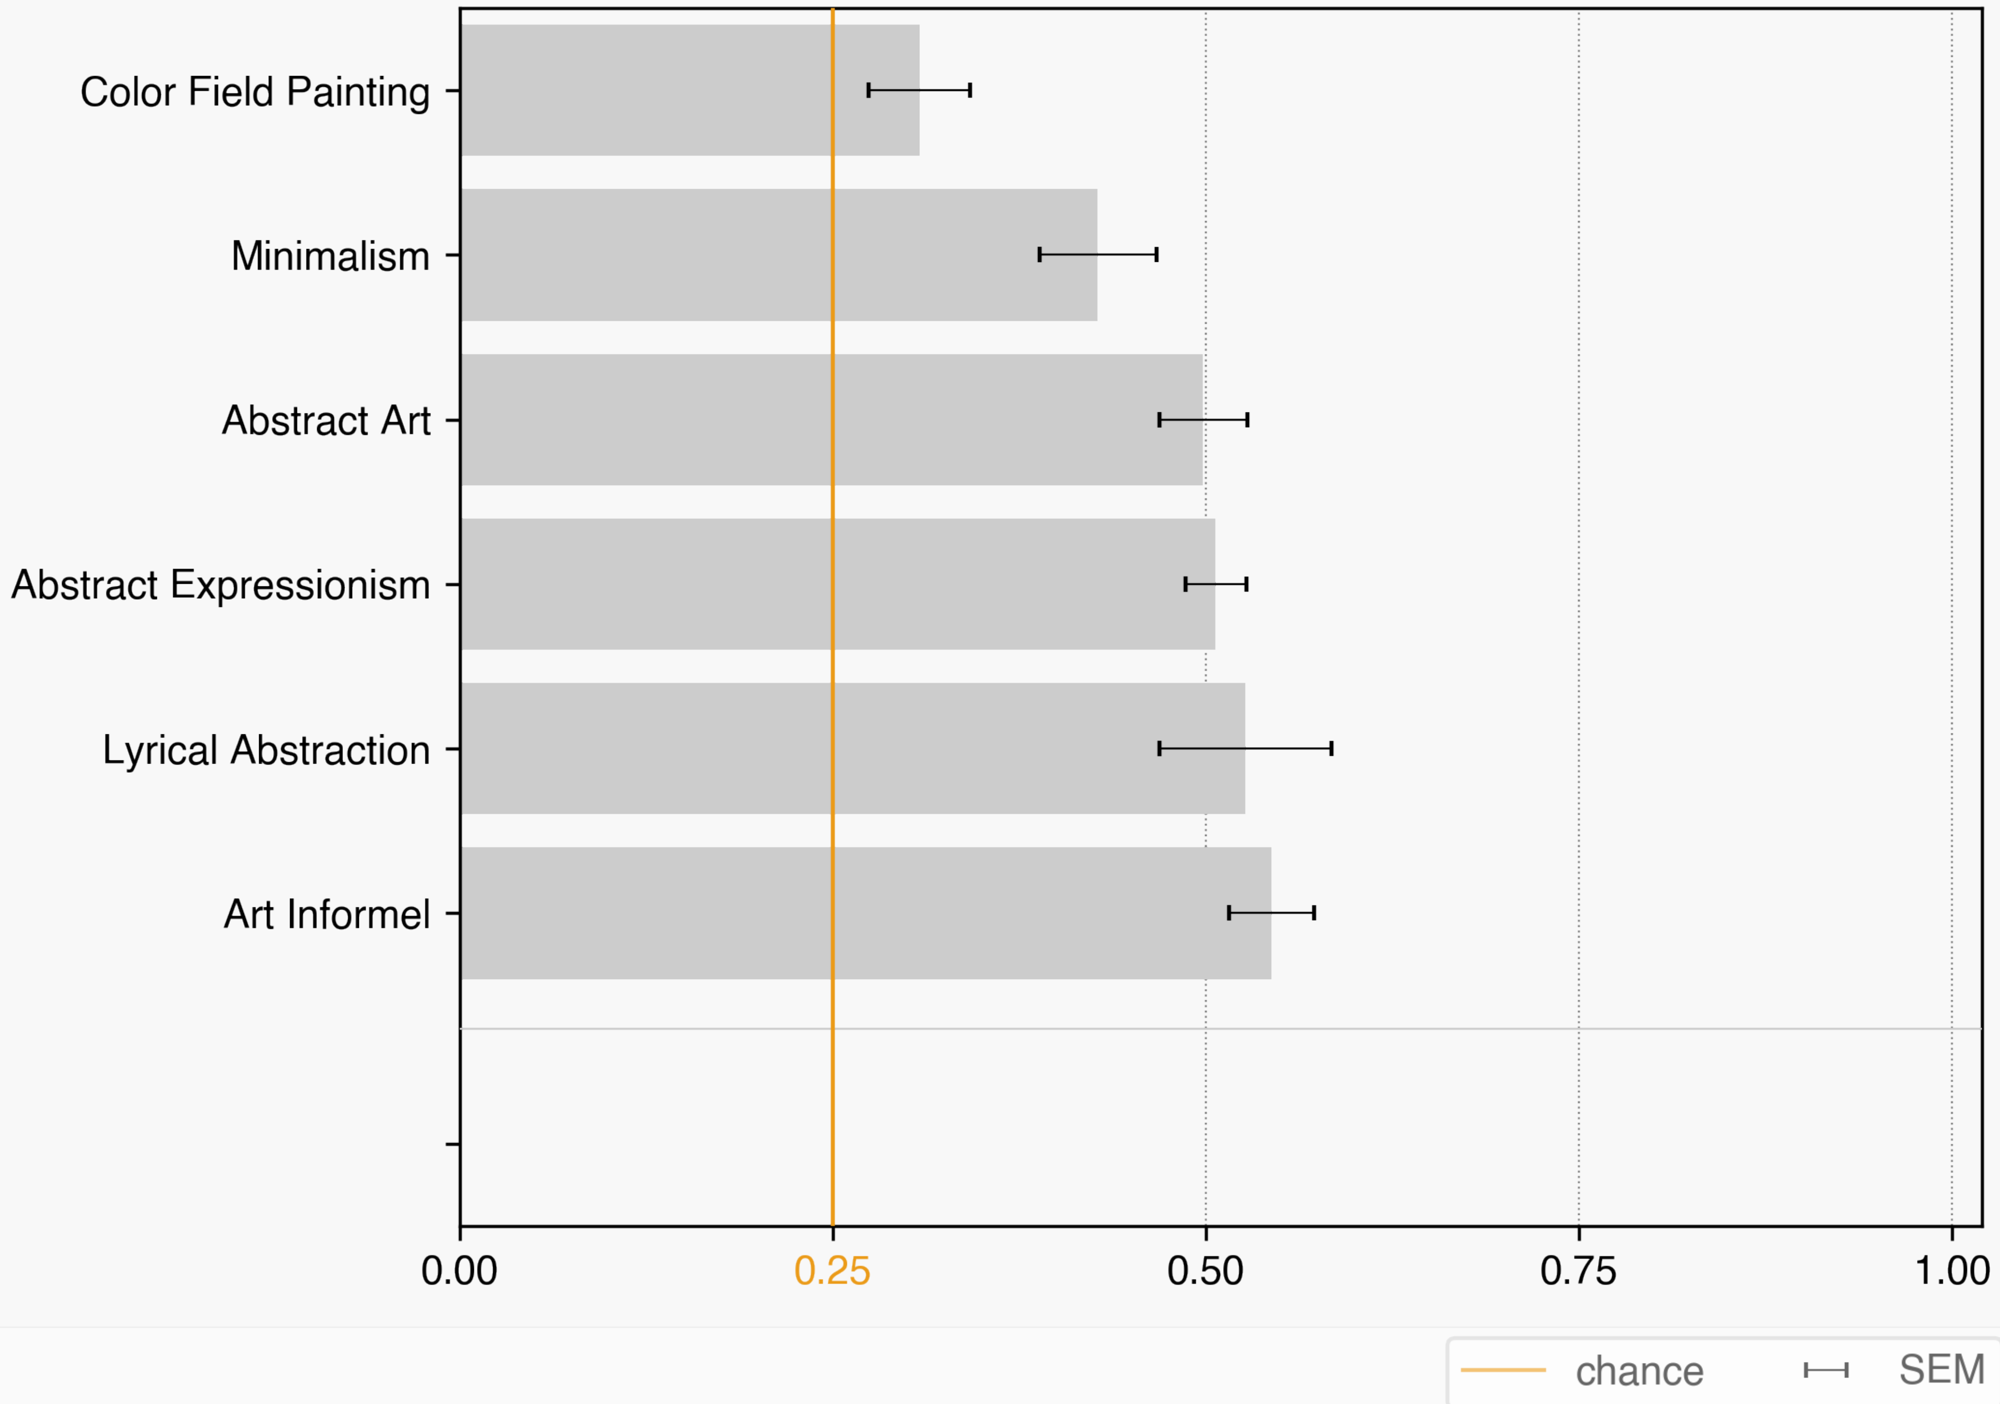

Speaker notes

- For a selection of the most representative abstract styles, the results confirm the literature performance. (Minimalism and Color Field Painting are maybe lower because they offer less pictorial content and more perfect symmetries.)
- The realistic style is also, as expected, close to perfection.

Human Performance

Accuracy per Style

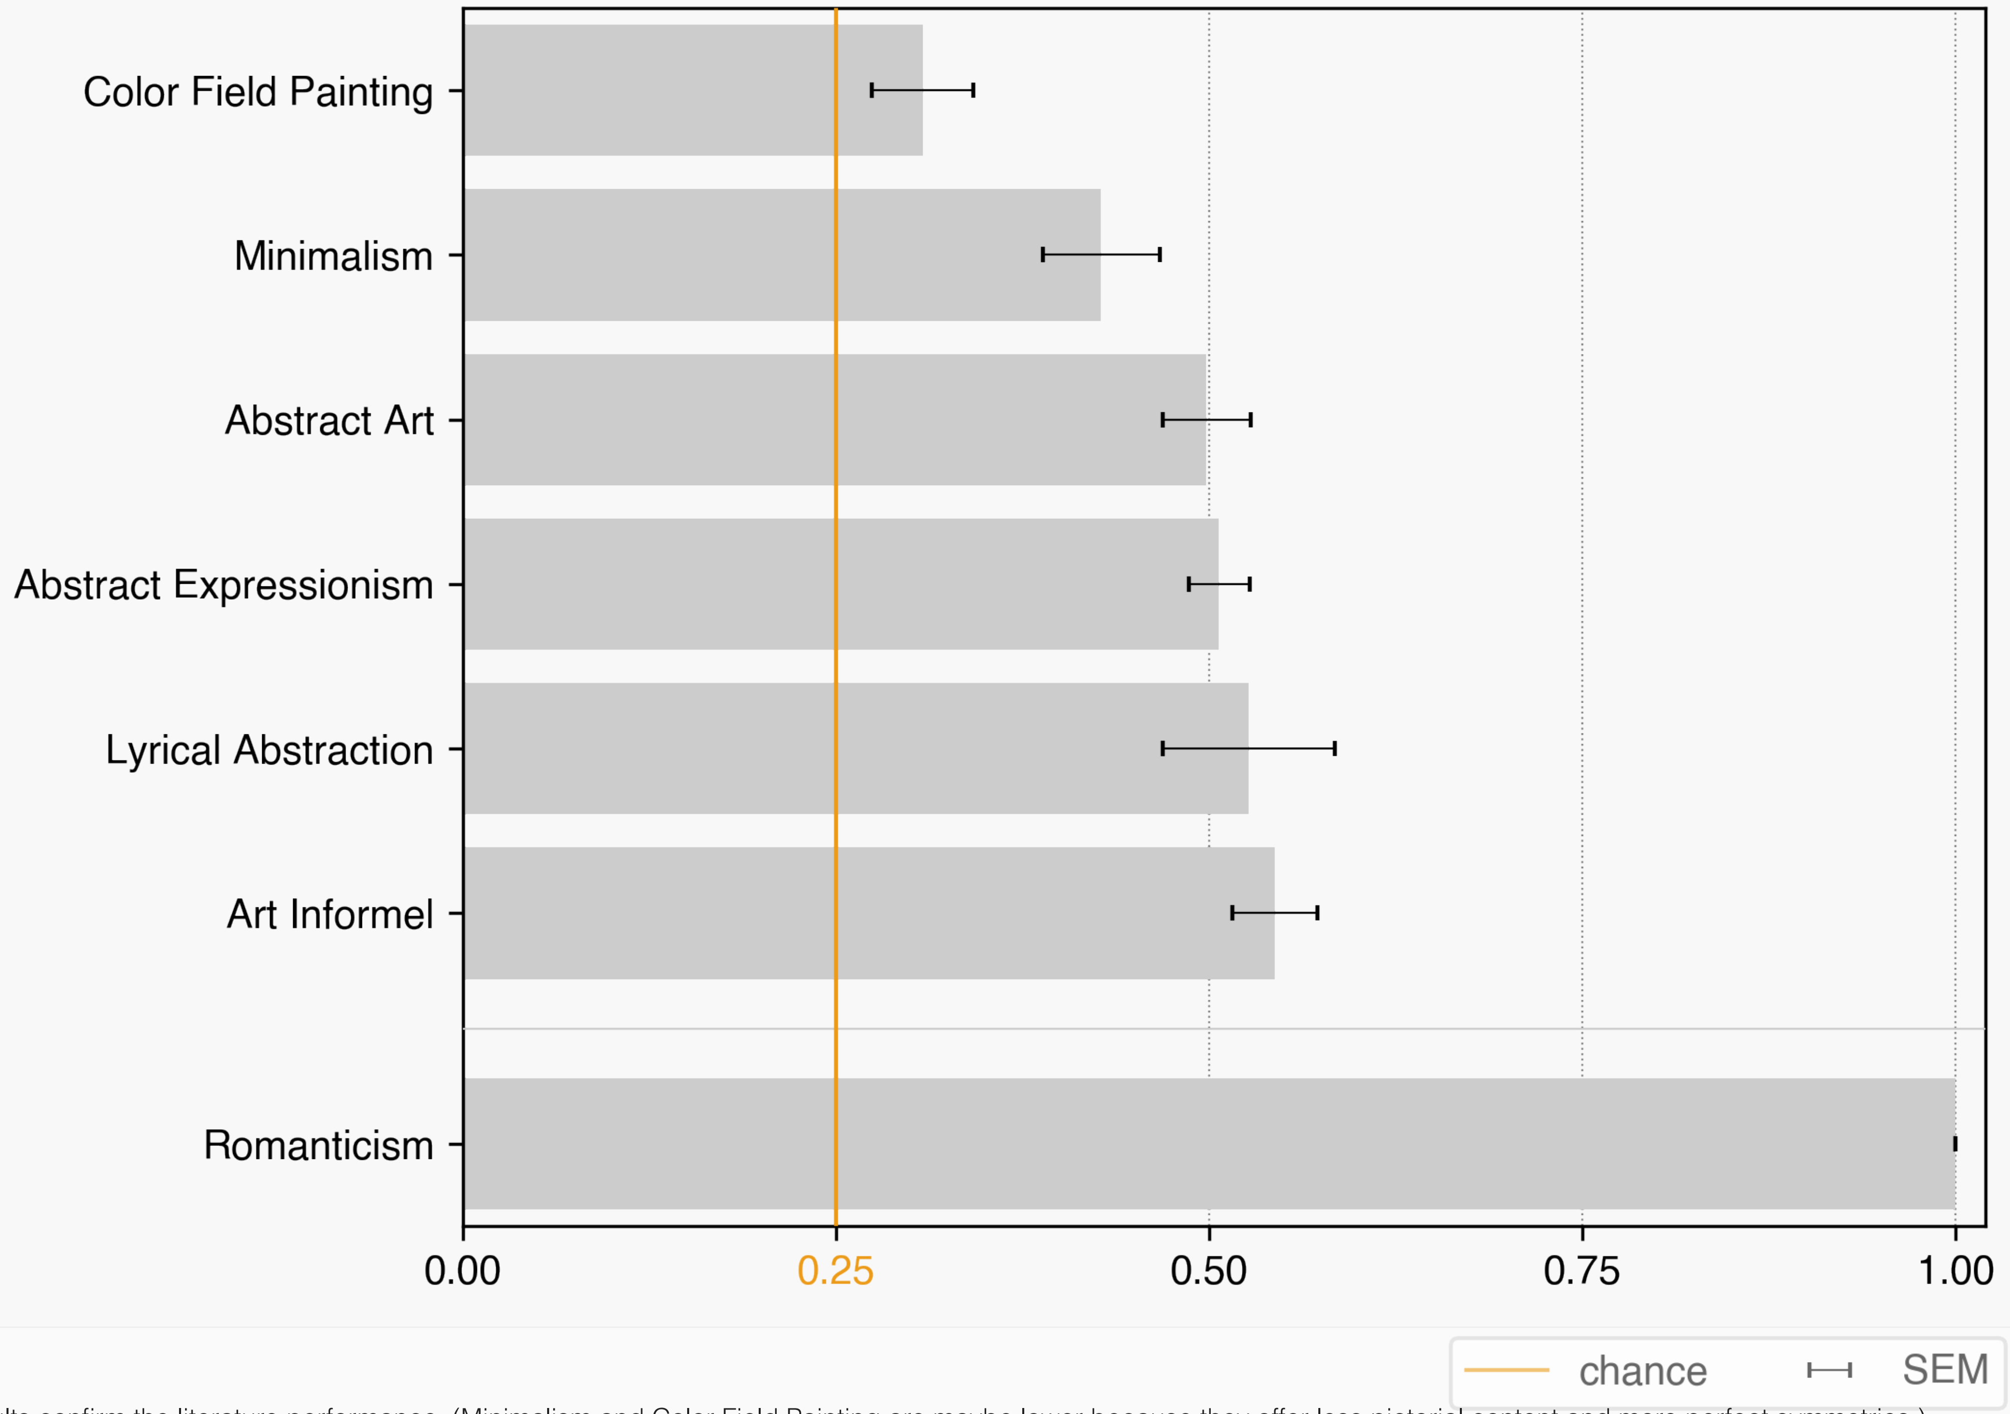

Speaker notes

- For a selection of the most representative abstract styles, the results confirm the literature performance. (Minimalism and Color Field Painting are maybe lower because they offer less pictorial content and more perfect symmetries.)
- The realistic style is also, as expected, close to perfection.

Human Performance

Accuracy per Style
+ Model

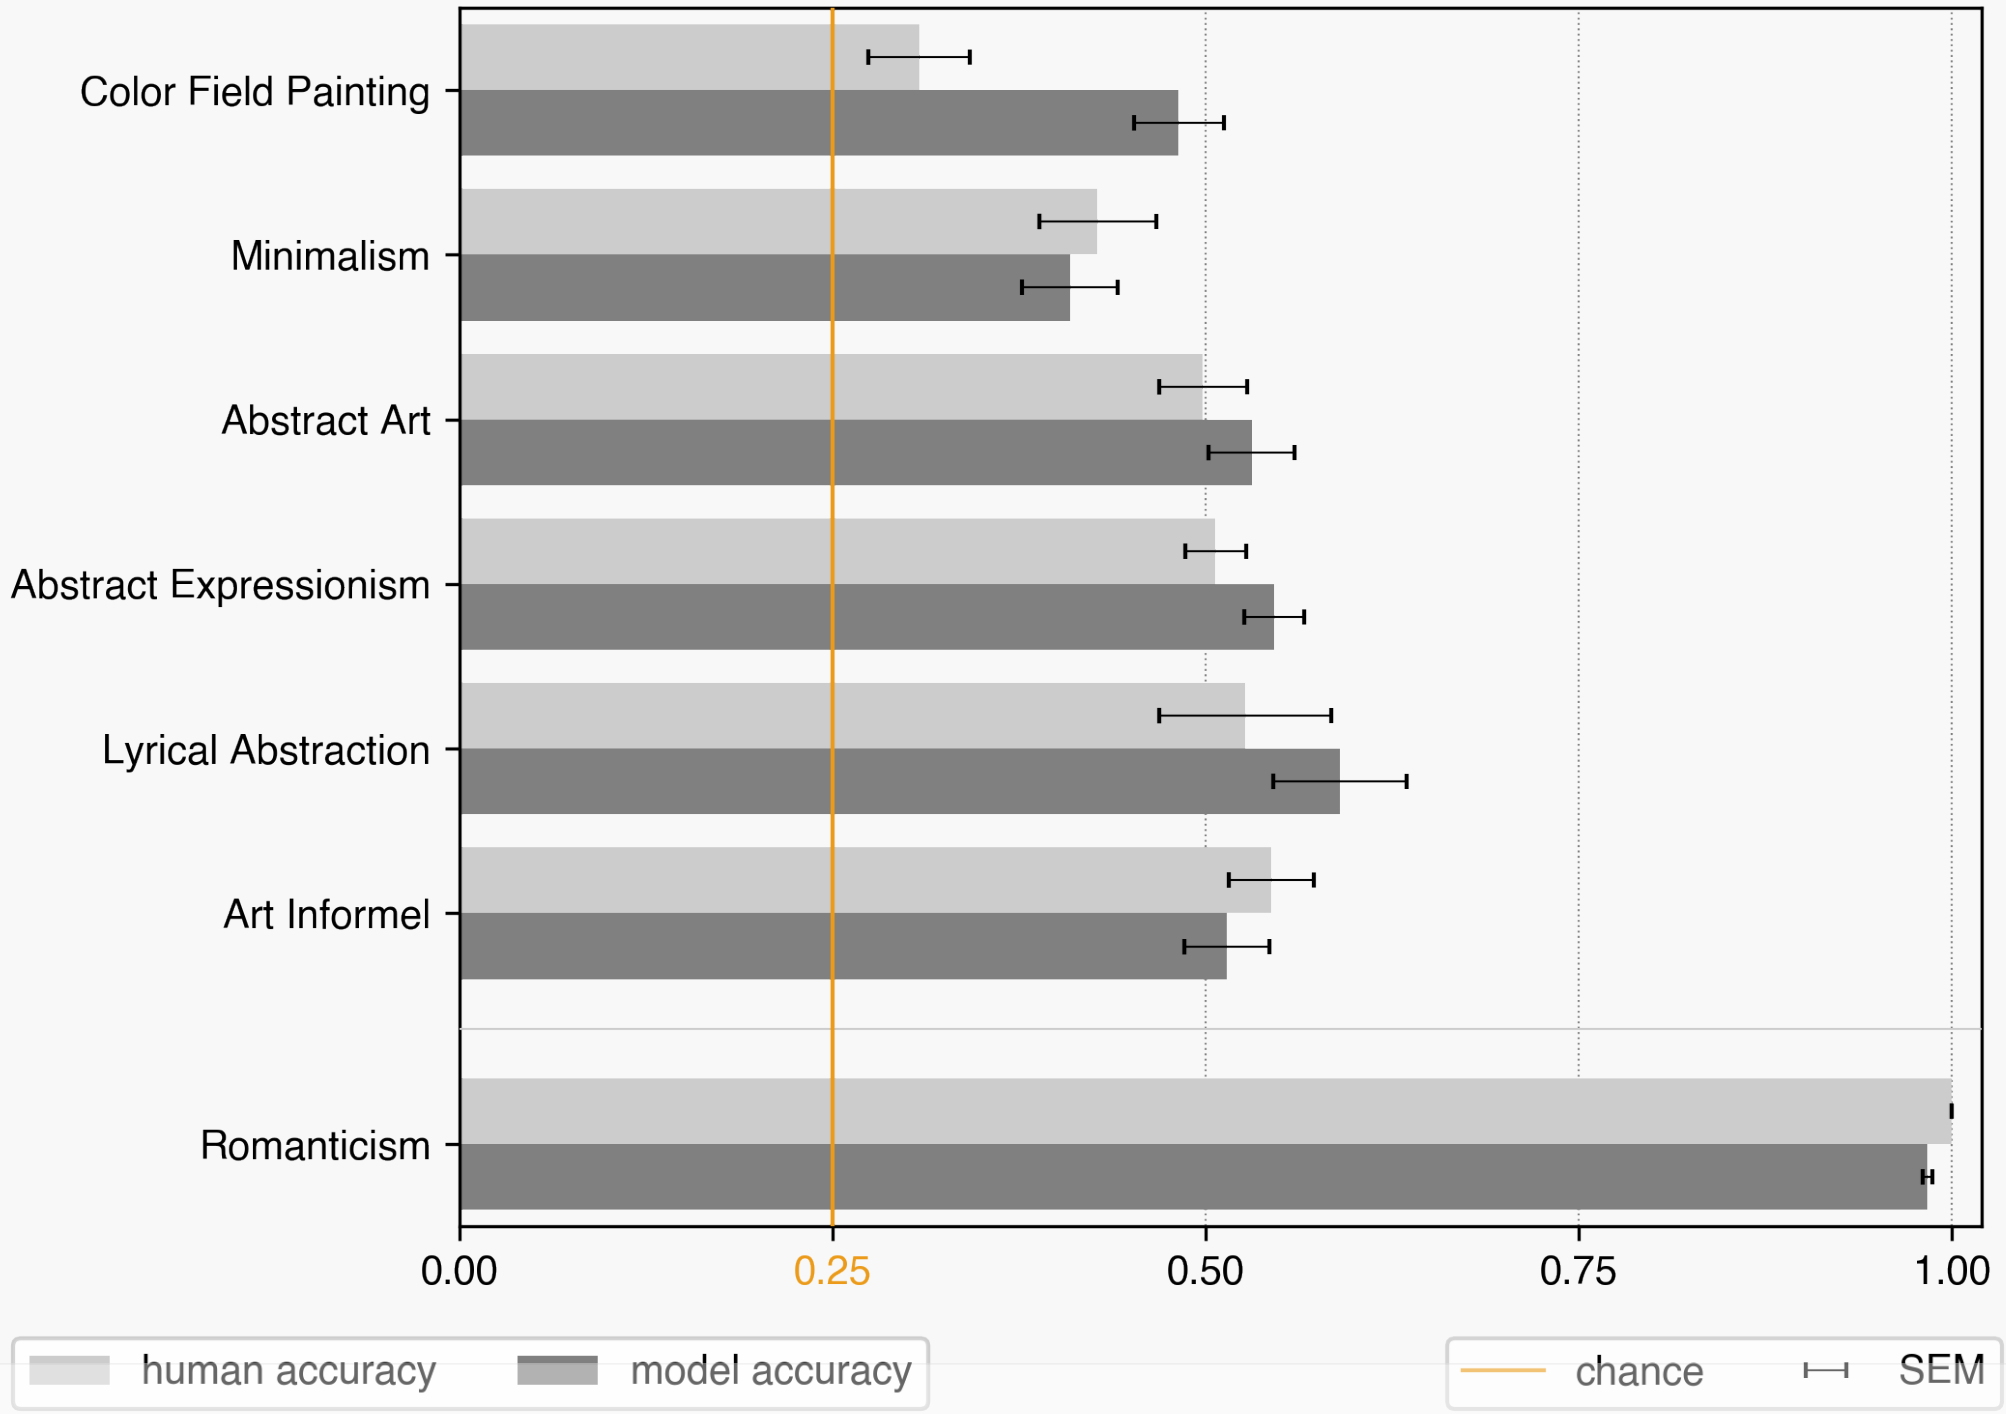

Speaker notes

If we add the model performance, its accuracy reflects quite well human's results, even if it is slightly better in average.

Human Performance

Accuracy per Layer
+ Model

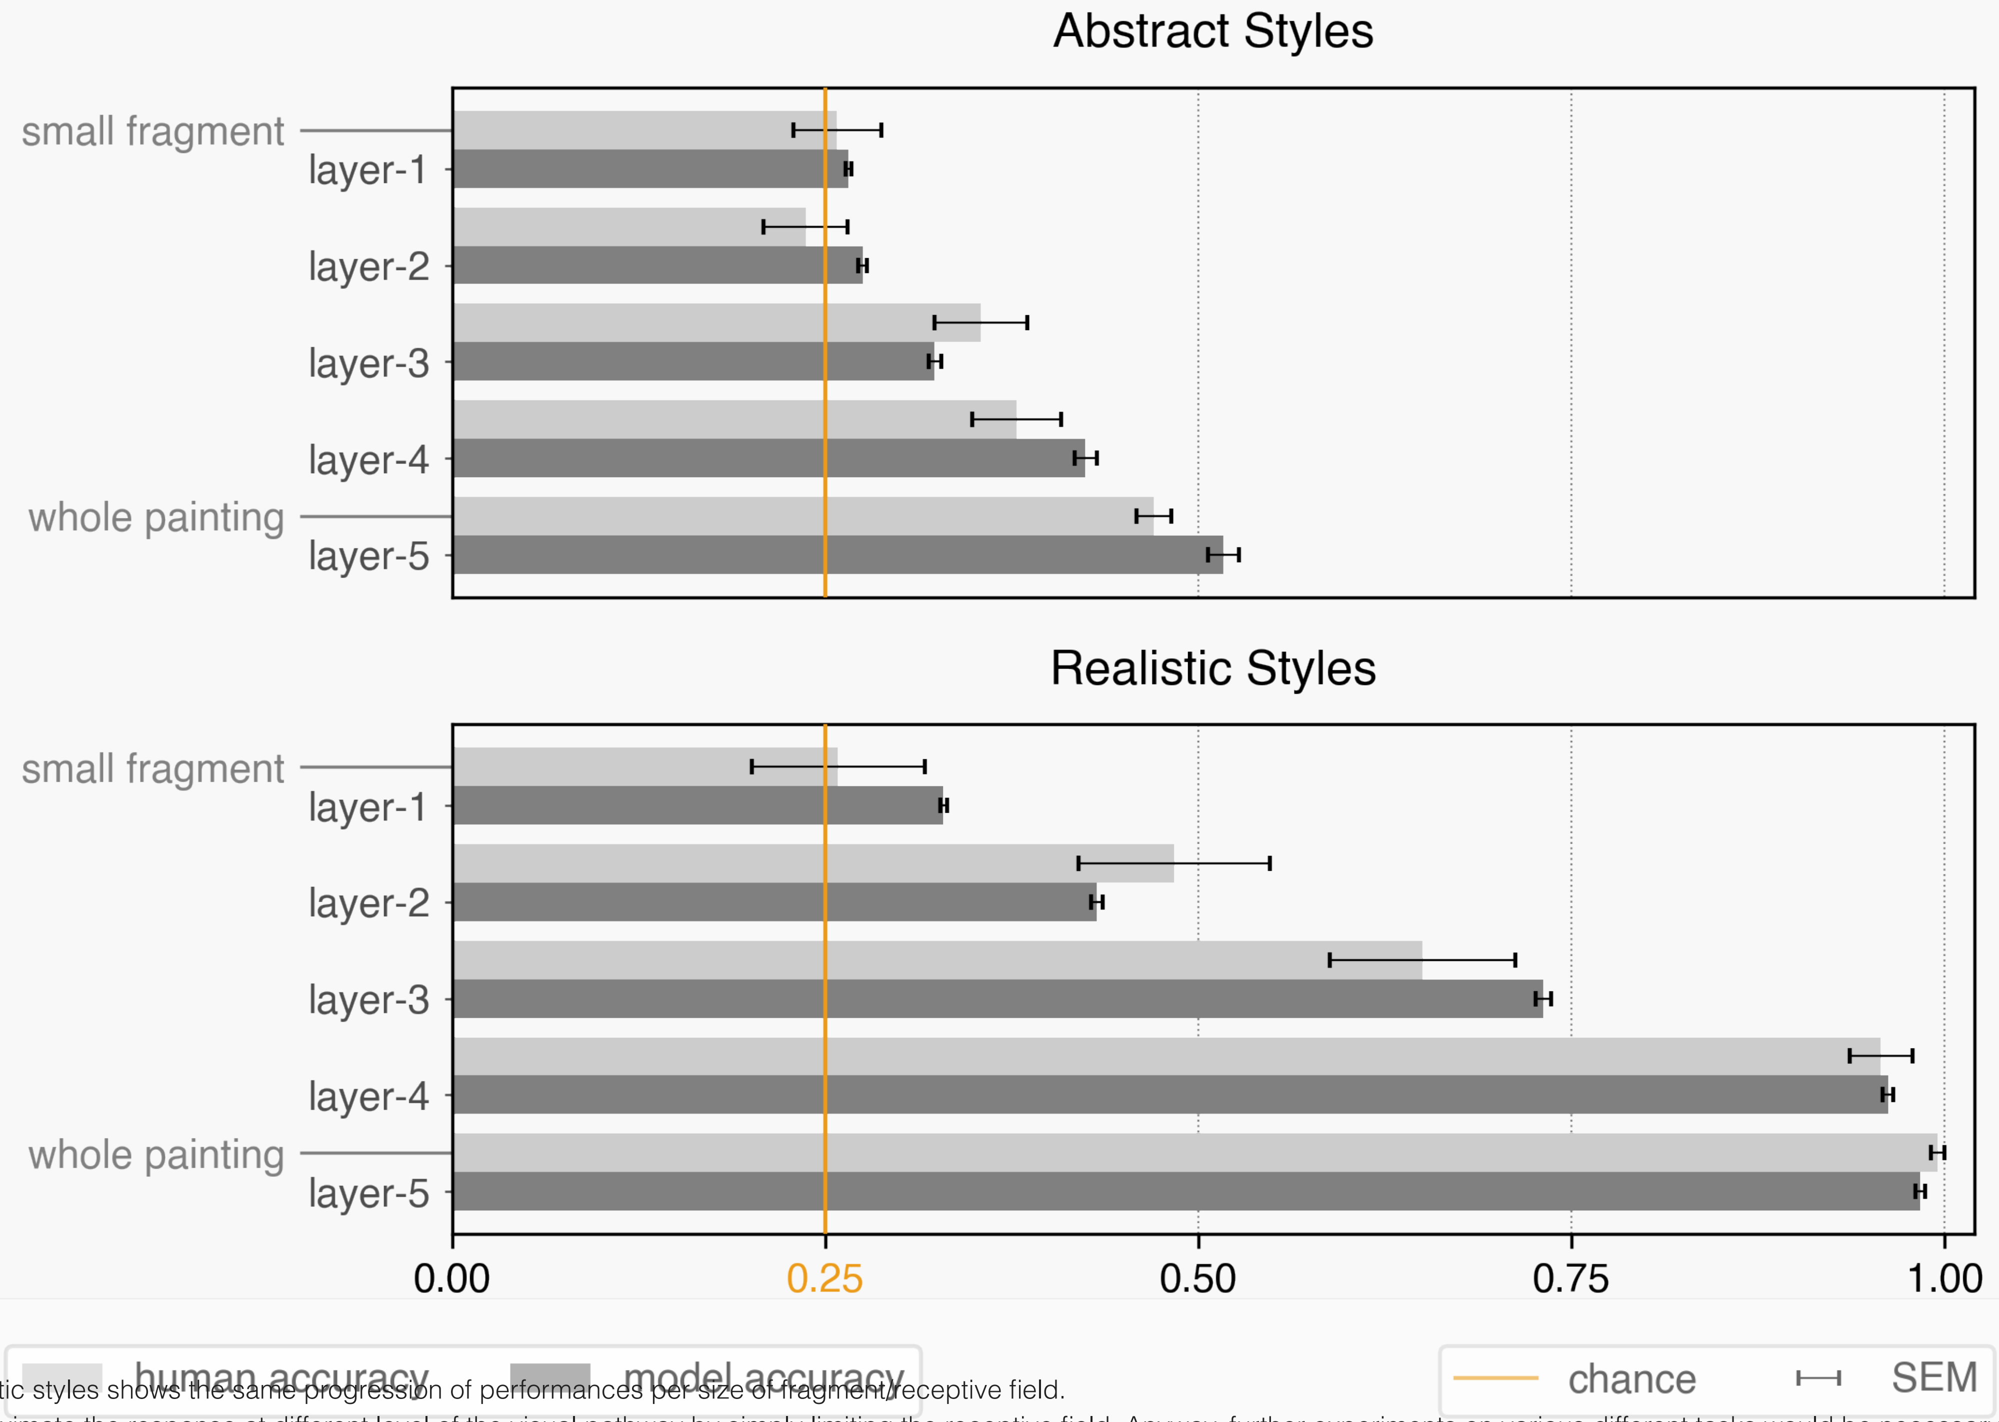

Speaker notes

- This correlation also exists on a per layer basis. Abstract and realistic styles shows the same progression of performances per size of fragment/receptive field.
- This finding is quite surprising because it seems that we can approximate the response at different level of the visual pathway by simply limiting the receptive field. Anyway, further experiments on various different tasks would be necessary to confirm this more generally.

Human Performance

Per Painting Agreement with Model

- small fragments / earlier layers
  → weak agreement

- whole painting / deeper layers
  → good agreement

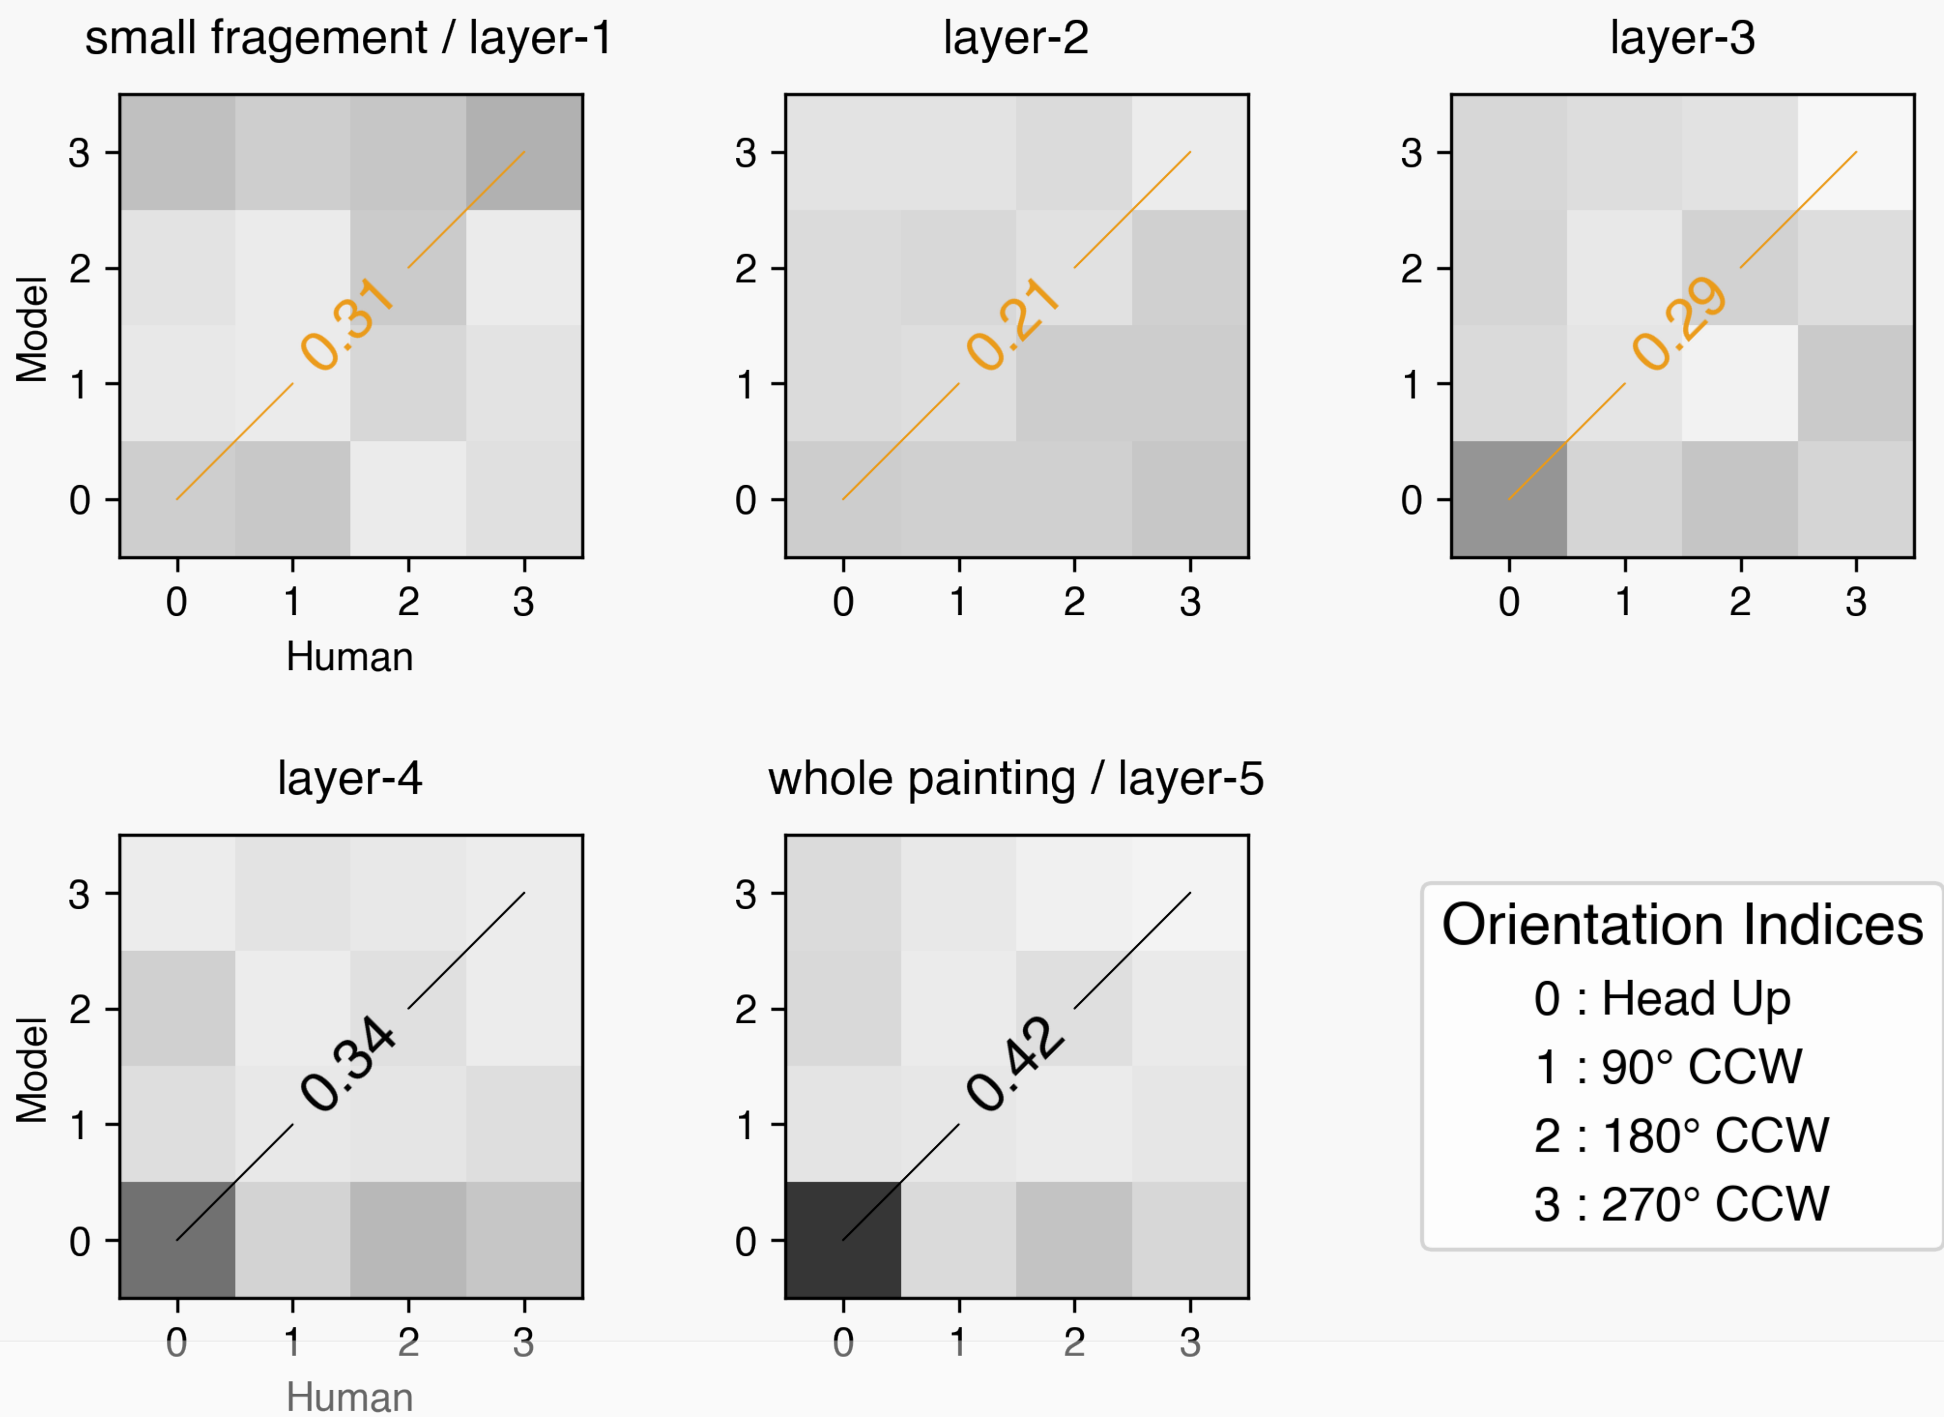

Speaker notes

But all these results were in average. What you see now, is the density of the corresponding orientations chosen by the model and humans for each abstract painting. We can observe that the total agreement, lying on the diagonal, is strong for the whole paintings only. With smaller fragments, it shifts toward chance, meaning that even if human and model perform similarly in average, their strategies should differ, especially for mistaken orientations.

Human Performance

Human Agreement

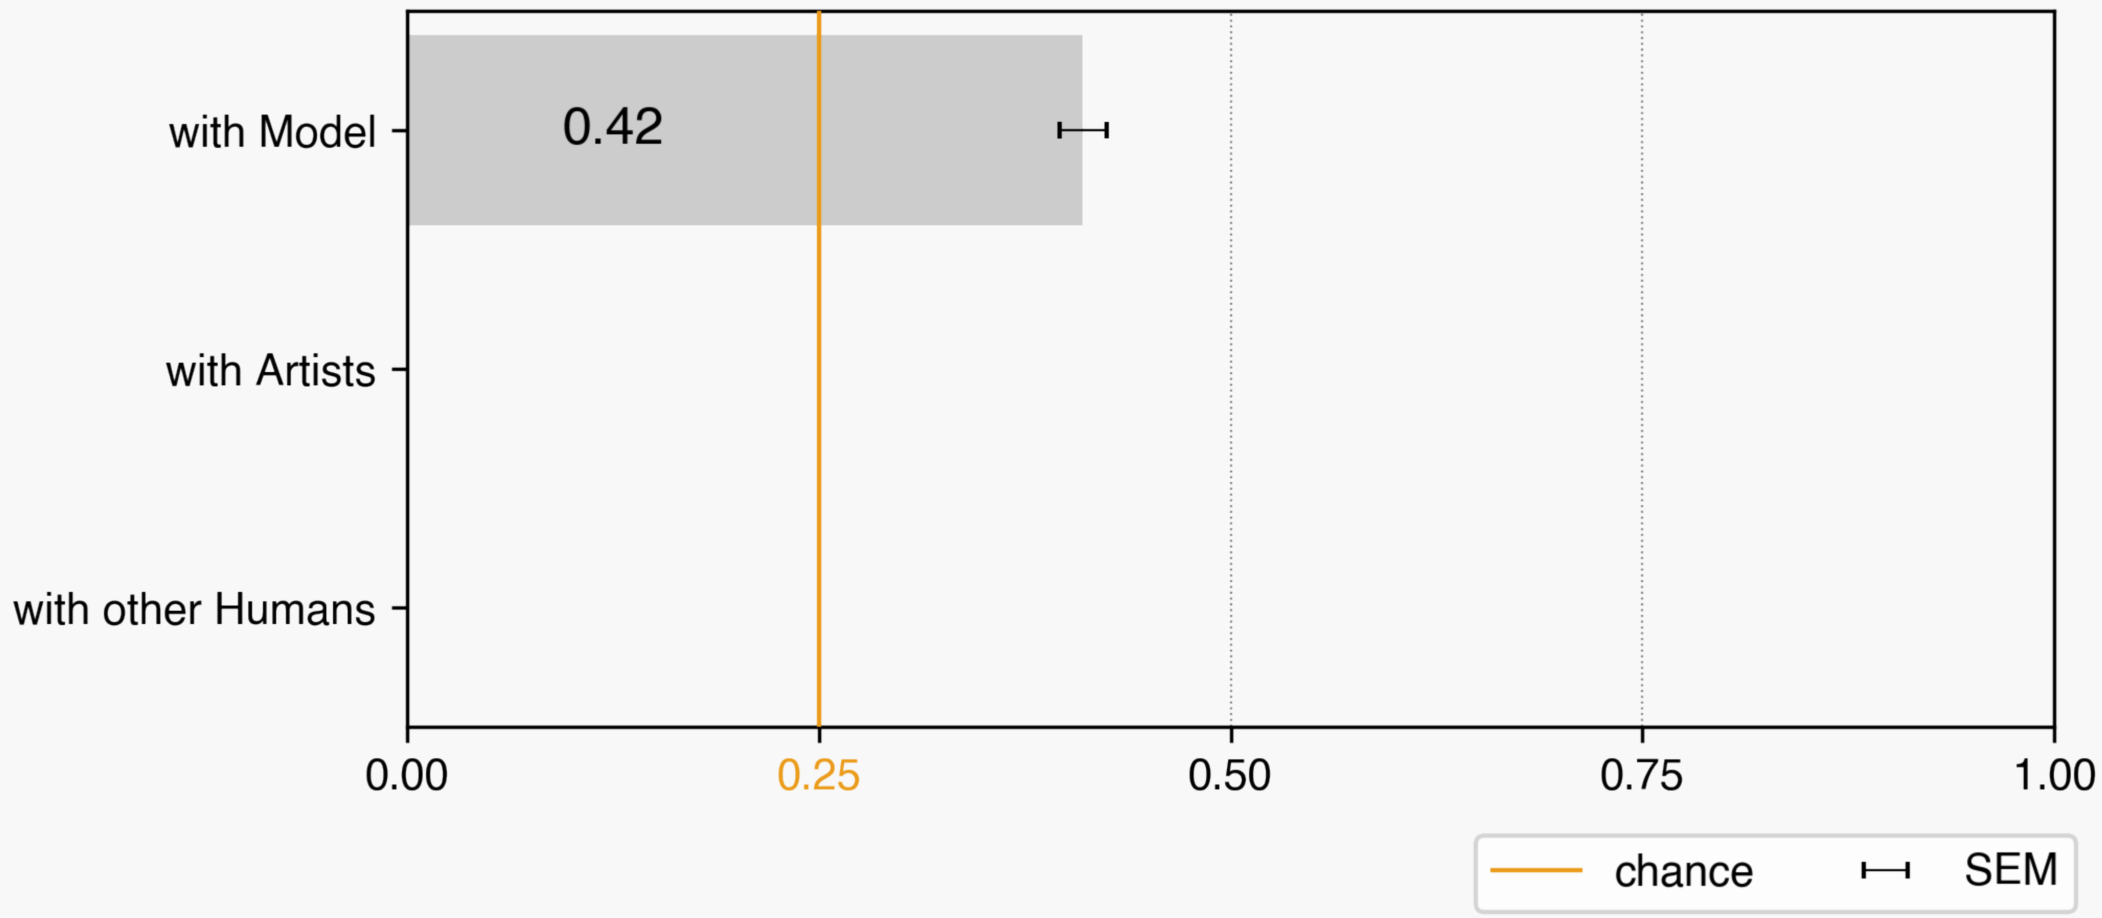

Speaker notes

- To finish, lets plot the human mean agreement with different entities. With the model we have just seen that it is around 42%.
- With artists, it is in average slightly better, around 47%.
- And for the paintings that have been evaluated multiple times by humans, it is even reaching 49%, amusingly suggesting that artists themselves may not be the most reliable reference on this task :D

Human Performance

Human Agreement

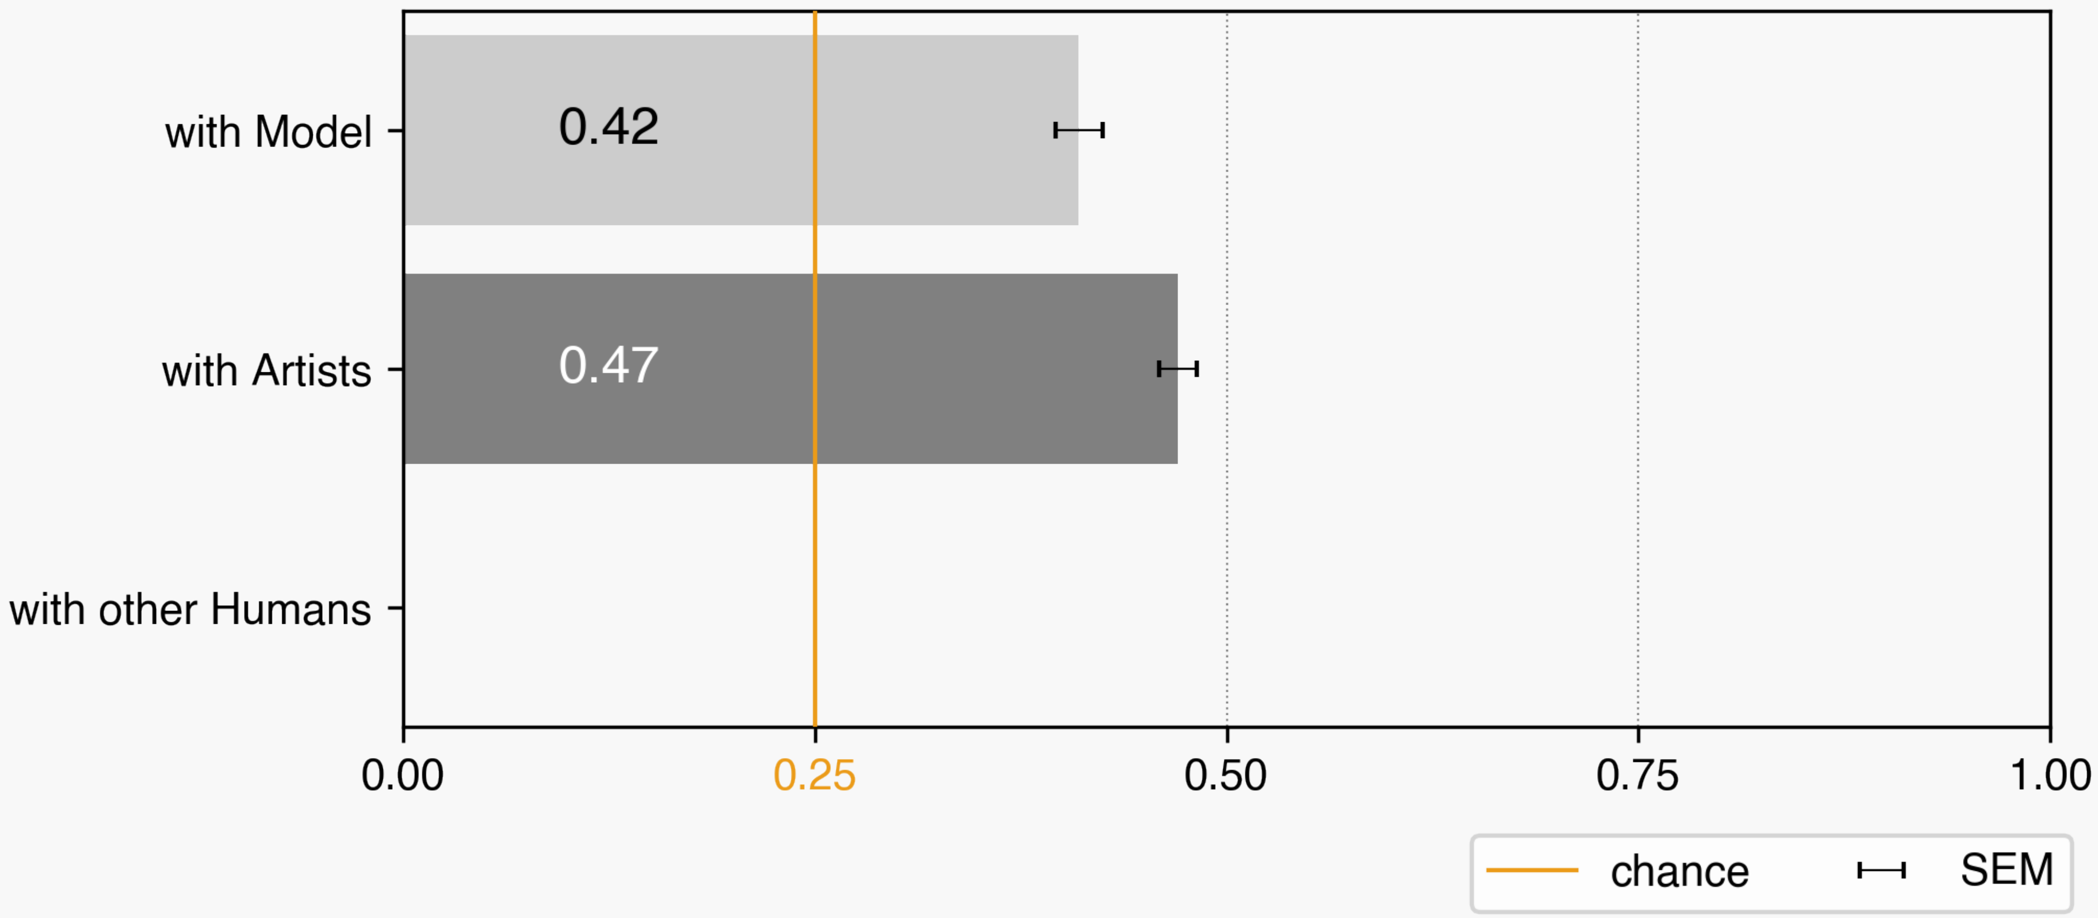

Speaker notes

- To finish, lets plot the human mean agreement with different entities. With the model we have just seen that it is around 42%.
- With artists, it is in average slightly better, around 47%.
- And for the paintings that have been evaluated multiple times by humans, it is even reaching 49%, amusingly suggesting that artists themselves may not be the most reliable reference on this task :D

Human Performance

Human Agreement

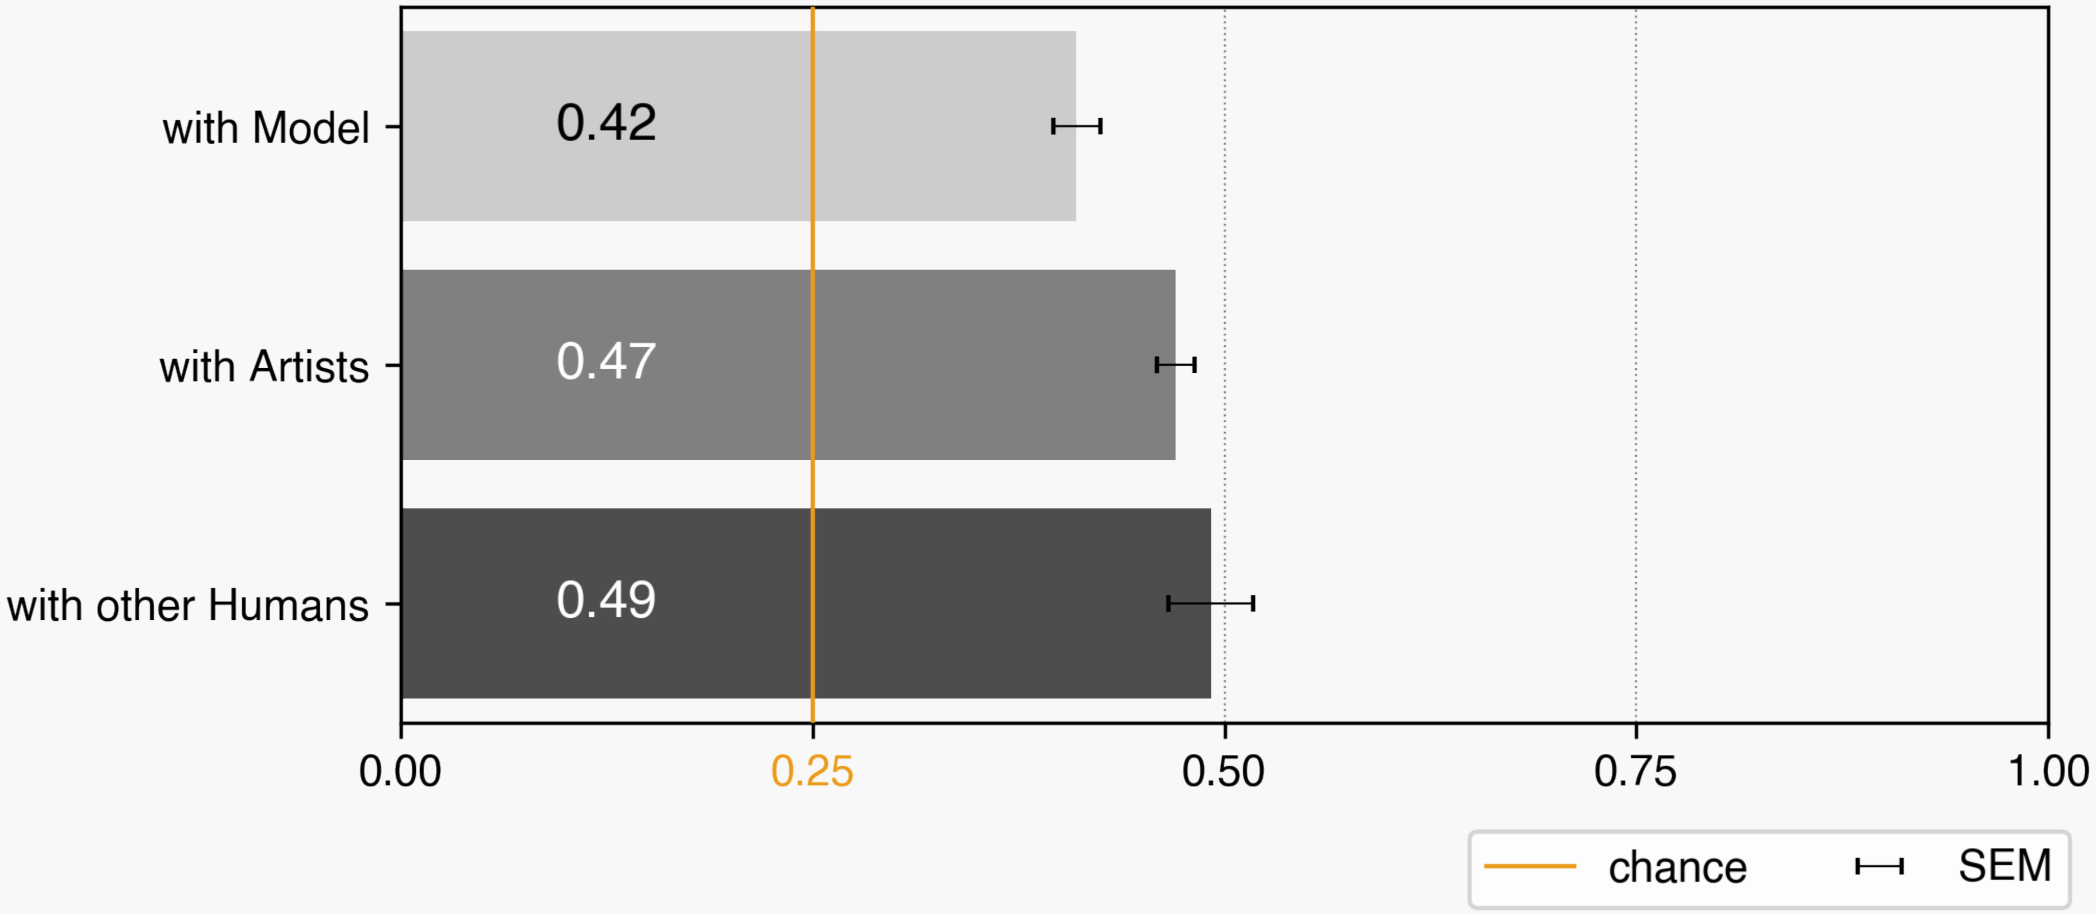

Speaker notes

- To finish, lets plot the human mean agreement with different entities. With the model we have just seen that it is around 42%.
- With artists, it is in average slightly better, around 47%.
- And for the paintings that have been evaluated multiple times by humans, it is even reaching 49%, amusingly suggesting that artists themselves may not be the most reliable reference on this task :D

Conclusion

Speaker notes

Ok, to sum up...

Conclusion

- Our deep learning model captures human orientation perception across granularities and styles.
- Abstract art, more than other styles, relies on large spatial integration of cues, corresponding to larger receptive field of deeper layers.
- Fragmented stimuli appear to probe human mechanisms corresponding to superficial layers in the model, however
- the detailed operation of human mechanisms is not identical to the modeled components for small fragments.

Speaker notes

- First, our deep learning model captures human orientation perception across granularities and styles.
- Secondly, abstract art, more than other styles, relies on large spatial integration of cues, corresponding to larger receptive field of deeper layers.
- Fragmented stimuli appear to probe human mechanisms corresponding to superficial layers in the model, however
- the detailed operation of human mechanisms is not identical to the modeled components for small fragments.

I was coming home with my paint box after a study, still lost in my dreams, when I suddenly saw a painting of an indescribable beauty. [...] I was first dumbfounded [but] I finally solved the enigma : it was one of my painting put on its side.

Wassily Kandinsky

---

contact : pierre.lelievre@psl.eu website : experiments.plelievre.com

PSL UNIVERSITÉ PARIS | ENS | SACRe | LSP Laboratoire des Systèmes Perceptifs

Speaker notes

(Read quote if time)
Thank you for your attention.